2013 Workshop on Computational Models of Narrative

CMN'13, 4–6 August, 2013, Hamburg, Germany

Edited by

Mark A. Finlayson
Bernhard Fisseni
Benedikt Löwe
Jan Christoph Meister

OASICS

Editors

Mark A. Finlayson
Computer Science and
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA, USA
`markaf@mit.edu`

Bernhard Fisseni
Fachbereich Mathematik
Universität Hamburg
Hamburg, Germany
`bernhard.fisseni@uni-due.de`

Institut für Germanistik
Universität Duisburg-Essen
Essen, Germany

Benedikt Löwe
Institute for Logic, Language
and Computation
Universiteit van Amsterdam
Amsterdam, The Netherlands
`b.loewe@uva.nl`

Fachbereich Mathematik
Universität Hamburg
Hamburg, Germany

Jan Christoph Meister
Institut für Germanistik II
Universität Hamburg
Hamburg, Germany
`jan-c-meister@uni-hamburg.de`

OASIcs – OpenAccess Series in Informatics

OASIcs aims at a suitable publication venue to publish peer-reviewed collections of papers emerging from a scientific event. OASIcs volumes are published according to the principle of Open Access, i.e., they are available online and free of charge.

Contents

Invited Talks

Contributed Papers

◼ Preface

The workshop series *Computational Models of Narrative* (CMN) is dedicated to advancing a nascent field: the computationally-grounded science of narrative; we believe that a true science of narrative must adhere to the principle espoused by Herbert Simon in his book *The Sciences of the Artificial*: that without computational modeling, the science of a complex human phenomenon such as narrative will never be successful. This expands the workshop's purview beyond the limited body of effort that directly incorporates computer simulation and gives us a broad mandate to include a great deal of cognitive, linguistic, neurobiological, social scientific, and literary work—indeed, any research where the researchers have successfully applied their field's unique insights to narrative in a way that is compatible with a computational frame of mind. We seek work whose results are thought out carefully enough, and specified precisely enough, that they could eventually inform computational modeling of narrative.

In keeping with interdisciplinary nature of the field, the workshop series moves between different communities in order to enhance engagement, cross-pollination, and visibility. CMN'10 was hosted as one of Fall Symposia of the *Association for the Advancement of Artificial Intelligence* (AAAI) in Arlington, VA; CMN'12 was hosted by the *Language Resources and Evaluation Conference* (LREC) in İstanbul. This year, we were associated to the *Annual Meeting of the Cognitive Science Society* held in Berlin. In future years we plan to continue our peripatetic schedule by co-locating with neuroscience conferences and humanities conferences.

This workshop, CMN'13, was a satellite meeting of the *Annual Meeting of the Cognitive Science Society* held in Berlin (31 July to 3 August 2013), immediately preceding our workshop in Hamburg. The workshop organizers also organized a symposium entitled "Computational Aspects of Narratives" at the Berlin conference with Richard Gerrig, Jeffrey Loewenstein, Inderjeet Mani, Jan Christoph Meister, and Richard Young as speakers. The links to cognitive science were emphasized in our call for papers, and a number of papers linking traditional CMN topics to questions of cognitive science are represented in these proceedings.

This year we are proud to offer two best paper awards. First, the *Award for Best Student Paper on a Cognitive Science Topic* goes to Angela Nyhout of the University of Waterloo for her paper "Constructing spatial representations from narratives and non-narrative descriptions: Evidence from 7-year-olds" which was co-authored with Daniela O'Neill. Second, the *Award for Best Student Paper* goes to Graham Sack of Columbia University for his paper "Character Networks for Narrative Generation: Structural Balance Theory and the Emergence of Proto-Narratives."

Many thanks to our generous sponsors without whom this workshop would not have been possible: The John Templeton Foundation funds the project "What makes stories similar?" at the Universität Hamburg (grant number 20565) which in turn covered the larger part of the infrastructural costs of the workshop; the Minerva Program at the United States Office of the Secretary of Defense and the United States Office of Naval Research Global (ONR-G) provided money for travel grants; and the Cognitive Science Society funded one of our best paper awards.

The host organisations in Hamburg were the *Fachbereich Mathematik* and the *Interdisciplinary Center for Narratology*; the *Faculteit der Natuurwetenschappen, Wiskunde en Informatica* of the *Universiteit van Amsterdam* provided important infrastructural support.

We also thank our student assistants which helped us in preparing proceedings and the conference: Varun B. Dwarakanathan for his help in consolidating the bibliographies, and Tanja Auge, Alexander Block and Mira Viehstädt for local support during the conference.

Mark A. Finlayson
Cambridge, Massachusetts

Bernhard Fisseni
Benedikt Löwe
Jan Christoph Meister
Hamburg, Germany

List of Authors

Ines Adornetti
Department of Philosophical Researches
University of Roma "Tor Vergata"
via Columbia 1
00133 Rome
Italy
`inesado@yahoo.it`

Thijs Alofs
Human Media Interaction
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands
`t.alofs@gmail.com`

Julio César Bahamón
Department of Computer Science
Liquid Narrative Research Group
North Carolina State University
Suite 2246, Engineering Building II
Campus Box 8206
Raleigh NC 27695, USA
`jcbahamo@ncsu.edu`

B. Hunter Ball
Department of Psychology
Arizona State University
P.O. Box 871104
Tempe, AZ 85287-1104, USA
`hunter.ball@asu.edu`

Mehul Bhatt
Cognitive Systems and
SFB/TR 8: Spatial Cognition
University of Bremen
Enrique-Schmidt-Str. 5
28334 Bremen
Germany
`bhatt@informatik.uni-bremen.de`

Josep Blat
Interactive Technologies Group, Information
Technologies Department, Pompeu Fabra
University
Tànger 122–140, 08018 Barcelona, Spain
`josep.blat@upf.edu`

Andrea Bolioli
Cross Library Services srl
Via Sommarive 18
38123 Trento, Italy
`bolioli@cross-library.com`

Fritz Breithaupt
Department of Germanic Studies
Indiana University
1020 E. Kirkwood Avenue
Bloomington, IN 47405-6601, USA
`fbreitha@indiana.edu`

Gene A. Brewer
Department of Psychology
Arizona State University
P.O. Box 871104
Tempe, AZ 85287-1104, USA
`gene.brewer@asu.edu`

David Andre Broniatowski
Center for Advanced Modeling in the Social,
Behavioral, and Health Sciences
Johns Hopkins University
5801 Smith Avenue
Baltimore, MD 21209, USA
`broniatowski@jhu.edu`

Matteo Casu
CELI srl
via San Quintino 31
10121 Torino, Italy
`casu@celi.it`

Joaquim Colàs Alvarez
Interactive Technologies Group
Information Technologies Department
Pompeu Fabra University
Tànger 122–140, 08018 Barcelona, Spain
`joaquim.colas@upf.edu`

Steven R. Corman
Arizona State University
P.O. Box 871205
Tempe, AZ 85287-1205, USA
`steve.corman@asu.edu`

Erica Cosentino
Department of Philosophy
University of Calabria
via P. Bucci 18/c
87036 Rende (CS)
Italy
`ericacosentino@libero.it`

Rossana Damiano
Dipartimento di Informatica and CIRMA
Università di Torino
C.so Svizzera 185
10149 Torino, Italy
`rossana@di.unito.it`

Sándor Darányi
Swedish School of Library and Information
Science
University of Borås
Allégatan 1
50190 Borås, Sweden
`Sandor.Daranyi@hb.se`

Morteza Dehghani
Institute for Creative Technologies
University of Southern California
12015 Waterfront Dr
Los Angeles, CA 90094, USA
`morteza@ict.usc.edu`

Louis Dessalles
INFRES, Network and Computer Science
Department
Télécom ParisTech
46 Rue Barrault
75013 Paris, France
`jean-louis.dessalles@`
`telecom-paristech.fr`

Francesco Ferretti
Department of Philosophy, Communication
and Visual Arts
University of Roma Tre
via Ostiense 234
00146 Rome
Italy
`francesco.ferretti@uniroma3.it`

Bernhard Fisseni
Fachbereich Mathematik
Universität Hamburg
Bundesstr. 55
20146 Hamburg, Germany
Universität Duisburg-Essen
Universitätsstr. 12
45117 Essen, Germany
`Bernhard.fisseni@uni-due.de`

Kevin Gardner
Hutton Honors College
Indiana University
1020 E. Kirkwood Avenue
Bloomington, IN 47405-6601, USA
`kmgardne@indiana.edu`

Richard Gerrig
Department of Psychology
Stony Brook University
Stony Brook, NY 11794-2500, USA
`richard.gerrig@stonybrook.edu`

Pablo Gervás
Instituto de Tecnología del Conocimiento
Universidad Complutense de Madrid
c/ Profesor José García Santesmases s/n
28040 Madrid, Spain
`pgervas@sip.ucm.es`

Sarah I. Gimbel
Brain and Creativity Institute
University of Southern California
3620A McClintock Avenue
Los Angeles, CA 90089, USA
`sgimbel@usc.edu`

Andrew S. Gordon
Institute for Creative Technologies
University of Southern California
12015 Waterfront Dr
Los Angeles, CA 90094, USA
`gordon@ict.usc.edu`

D. Fox Harrell
Comparative Media Studies Program and
Computer Science and Artificial Intelligence
Laboratory
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139, USA
`fox.harrell@mit.edu`

Mary Helen Immordino-Yang
Brain and Creativity Institute
University of Southern California
3620A McClintock Avenue
Los Angeles, CA 90089, USA
`immordin@ict.usc.edu`

Franciska de Jong
Human Media Interaction
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands
Erasmus Studio
Erasmus University Rotterdam
P.O. Box 1738
3000 DR Rotterdam, The Netherlands
`f.m.g.dejong@utwente.nl`

Dominic Kao
Computer Science and Artificial Intelligence
Laboratory
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139, USA
`dkao@mit.edu`

Jonas Kaplan
Brain and Creativity Institute
University of Southern California
3620A McClintock Avenue
Los Angeles, CA 90089, USA
`jtkaplan@ict.usc.edu`

Jackie S. Kim
Institute for Creative Technologies
University of Southern California
12015 Waterfront Dr
Los Angeles, CA 90094, USA
`skim@ict.usc.edu`

Sarit Kraus
Department of Computer Science
Bar-Ilan University
Ramat Gan, 5290002, Israel
`sarit@cs.biu.ac.il`

John K. Kruschke
Department of Psychological and Brain
Sciences
Indiana University
1020 E. Kirkwood Avenue
Bloomington, IN 47405-6601, USA
`kruschke@indiana.edu`

Aadil Kurji
Department of Philosophy
University of Bristol
Woodland Rd
Bristol BS8 1TB, UK
`ak12004@bristol.ac.uk`

Elektra Kypridemou
School of Pure and Applied Sciences
Open University of Cyprus
P.O. Box 12794
2252, Latsia, Cyprus
`elektra.kypridemou@st.ouc.ac.cy`

Sanne M.A. Lamers
Department of Psychology, Health and
Technology
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands
`s.m.a.lamers@utwente.nl`

Maurizio Lana
Dipartimento di Studi Umanistici
Università degli Studi del Piemonte
Orientale
via Manzoni 10
13100 Vercelli, Italy
`m.lana@lett.unipmn.it`

Faith Lawrence
Digital Humanities Deparment
King's College, London
26–29 Drury Lane
London, WC2B 5RL, UK
`faith.lawrence@kcl.ac.uk`

Greg Lessard
French Studies
Queen's University
University Avenue
Kingston, Ontario, K7L 3N6, Canada
`greg.lessard@queensu.ca`

Michael Levison
School of Computing
Queen's University
University Avenue
Kingston, Ontario, K7L 3N6, Canada
`levison@cs.queensu.ca`

Torrin Liddell
Department of Psychological and Brain
Sciences
Indiana University
1020 E. Kirkwood Avenue
Bloomington, IN 47405-6601, USA
`tliddell@indiana.edu`

Antonio Lieto
Dipartimento di Informatica
Università di Torino
C.so Svizzera 185
10149 Torino, Italy
`lieto@di.unito.it`

Chong-U Lim
Computer Science and Artificial Intelligence
Laboratory
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139, USA
`culim@mit.edu`

Jeroen Linssen
Human Media Interaction
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands
`j.m.linssen@utwente.nl`

Inderjeet Mani
Yahoo! Labs
701 First Avenue
Sunnyvale, CA 94089, USA
`inderjeet.mani@gmail.com`

Mike Metke
Institute for Creative Technologies
University of Southern California
12015 Waterfront Dr
Los Angeles, CA 90094, USA
`mmetke@ict.usc.edu`

Loizos Michael
School of Pure and Applied Sciences
Open University of Cyprus
P.O. Box 12794
2252, Latsia, Cyprus
`loizos@ouc.ac.cy`

Ayman Moghnieh
Interactive Technologies Group
Information Technologies Department,
Pompeu Fabra University
Tànger 122–140, 08018 Barcelona, Spain
`ayman.moghnieh@upf.edu`

Olufemi Deborah Ninan
Department of Computer Science &
Engineering
Faculty of Technology
Obafemi Awolowo University
220005 Ile-Ife, Nigeria
`jninan@oauife.edu.ng`

Angela Nyhout
Department of Psychology
University of Waterloo
200 University Ave W.
Waterloo, Ontario, N2L 3G1, Canada
`aknyhout@uwaterloo.ca`

Daniela K. O'Neill
Department of Psychology
University of Waterloo
200 University Ave W.
Waterloo, Ontario, N2L 3G1, Canada
`doneill@uwaterloo.ca`

Odetunji Ajadi Ọdéjọbí
Department of Computer Science &
Engineering
Faculty of Technology
Obafemi Awolowo University
220005 Ile-Ife, Nigeria
`oodejobi@oauife.edu.ng`

Nir Ofek
Department of Information Systems
Engineering
Ben-Gurion University of the Negev
P.O.B. 653
Beer-Sheva, 84105, Israel
`nirofek@post.bgu.ac.il`

Henry Prakken
Faculty of Law
University of Groningen
Department of Information and Computing
Sciences
Utrecht University
Princetonplein 5
3584 CC Utrecht, The Netherlands
`h.prakken@uu.nl`

Silja Renooij
Department of Information and Computing
Sciences
Utrecht University
Princetonplein 5
3584 CC Utrecht, The Netherlands
`s.renooij@uu.nl`

Valerie F. Reyna
Departments of Human Development and
Psychology, Center for Behavioral Economics
and Decision Research and Cornell Magnetic
Resonance Imaging Facility
Cornell University
B44, Martha Van Rensselaer Hall
Ithaca, NY 14853, USA
`vr53@cornell.edu`

Urs Richle
TECFA-FPSE
University of Geneva
Bd du Pont-d'Arve, 40
1211 Genève 4, Switzerland
`urs.richle@unige.ch`

Renato Roda
Dipartimento di Culture, Politiche, Società
Università degli Studi di Torino
Lungo Dora Siena 100 A
10153 Torino, Italy
`renato.roda@unito.it`

Lior Rokach
Department of Information Systems
Engineering
Ben-Gurion University of the Negev
P.O.B. 653
Beer-Sheva, 84105, Israel
`liorrk@bgu.ac.il`

Avi Rosenfeld
Department of Industrial Engineering
Jerusalem College of Technology
Jerusalem, 91160, Israel
`rosenfa@jct.ac.il`

Graham Alexander Sack
English & Comparative Literature
Department
Columbia University
602 Philosophy Hall
New York, NY 10027, USA
`gas2117@columbia.edu`

Kenji Sagae
Institute for Creative Technologies
University of Southern California
12015 Waterfront Dr
Los Angeles, CA 90094, USA
`sagae@ict.usc.edu`

Antoine Saillenfest
INFRES, Network and Computer Science
Department
Télécom ParisTech
46 Rue Barrault
75013 Paris, France
`antoine.saillenfest@`
`telecom-paristech.fr`

Deniz Sarikaya
Fachbereich Mathematik, ML
Universität Hamburg
Bundesstr. 55
20146 Hamburg, Germany
`deniz.sarikaya@stud.uni-hamburg.de`

Carl Schultz
Cognitive Systems and
SFB/TR 8: Spatial Cognition
University of Bremen
Enrique-Schmidt-Str. 5
28334 Bremen
Germany
`cschultz@informatik.uni-bremen.de`

Moshe Simon-Shoshan
Rothberg International School
Hebrew University
Mount Scopus Campus
Jerusalem, 91905, Israel
`mdshoshan@gmail.com`

Sigal Sina
Department of Computer Science
Bar-Ilan University
Ramat Gan, 5290002, Israel
`sinasi@macs.biu.ac.il`

Anneke Sools
Department of Psychology, Health and
Technology
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands
`a.m.sools@utwente.nl`

Jakob Suchan
Cognitive Systems and
SFB/TR 8: Spatial Cognition
University of Bremen
Enrique-Schmidt-Str. 5
28334 Bremen
Germany
`jsuchan@informatik.uni-bremen.de`

Ivo Swartjes
Ranj Serious Games
Lloydstraat 21m
3024 EA Rotterdam, The Netherlands
`ivo@ranj.nl`

Nicolas Szilas
TECFA-FPSE
University of Geneva
Bd du Pont-d'Arve, 40
1211 Genève 4, Switzerland
`nicolas.szilas@unige.ch`

Kimberly M. Talboom
Department of Psychology
Arizona State University
P.O. Box 871104
Tempe, AZ 85287-1104, USA
`kimberly.talboom@asu.edu`,

Alan Tapscott Baltar
Interactive Technologies Group
Information Technologies Department
Pompeu Fabra University
C/Roc Boronat 138, 08018 Barcelona, Spain
`alan.tapscott@upf.edu`

Mariët Theune
Human Media Interaction
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands
`m.theune@utwente.nl`

Christine Tipper
Brain and Creativity Institute
University of Southern California
3620A McClintock Avenue
Los Angeles, CA 90089, USA
`tipper@usc.edu`

Khiet P. Truong
Human Media Interaction
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands
`k.p.truong@utwente.nl`

Bart Verheij
Institute of Artificial Intelligence
University of Groningen
Nijenborgh 9
9747 AG Groningen, The Netherlands
`b.verheij@ai.rug.nl`

Mira Viehstädt
Fachbereich Mathematik, ML
Universität Hamburg
Bundesstr. 55
20146 Hamburg, Germany
`mira.viehstaedt@stud.uni-hamburg.de`

Charlotte Vlek
Institute of Artificial Intelligence
University of Groningen
Nijenborgh 9
9747 AG Groningen, The Netherlands
`c.s.vlek@rug.nl`

Gerben J. Westerhof
Department of Psychology, Health and
Technology
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands
`g.j.westerhof@utwente.nl`

R. Michael Young
Department of Computer Science
Liquid Narrative Research Group
North Carolina State University
Suite 2246, Engineering Building II
Campus Box 8206
Raleigh NC 27695, USA
young@csc.ncsu.edu

Sam Zorowitz
Department of Psychological and
Johns Hopkins University
3400 N. Charles St. Postcode and
Baltimore, MD 21218, USA
szorowi1@gmail.com

A Participatory Perspective on the Experience of Narrative Worlds

Richard Gerrig

Department of Psychology
Stony Brook University
Stony Brook, NY, USA
`richard.gerrig@stonybrook.edu`

──── **Abstract** ──

As people experience narratives, they often behave as if they are participants in the narrative world. This talk embraces that claim to develop a participatory perspective on readers' and viewers' narrative experiences. This perspective asserts, for example, that readers encode participatory responses as reactions to characters' utterances and actions. The talk will review three areas of empirical research that have emerged from this perspective. The first area will be readers' experiences of narrative mysteries—circumstances in which a text raises questions that are not immediately settled. The second area will be the consequences of readers' participation as they weigh in on characters' actions and decisions. The third area will be the potential for changes in people's beliefs and attitudes as a product of their narrative experiences.

1 Summary

As people experience narratives, they often behave as if they are participants in the narrative world. The core claim of the participatory perspective on readers' and viewers' narrative experiences is that readers regularly encode the types of mental contents they would encode were they really participants in the narrative's events (for a review, see [3]). For example, as people watch real-world events unfold, they are likely to encode preferences for how their acquaintances should behave and express approval or disapproval once those acquaintances have made behavioral decisions. People make parallel types of participatory responses—they encode preferences and evaluations—when they are transported to narrative worlds [2].

Several areas of empirical research have emerged from this perspective. An initial type of research has focused on readers' experiences of narrative mysteries—circumstances in which a text raises questions that are not immediately settled. How might readers respond to such mysteries? One possibility is that readers represent the information in the text without taking note of the mysteries it presents. However, the participatory perspective suggests that readers would attend to these mysteries as if they were real-world participants. By formally encoding such mysteries when they first arise, participants (in real life and as readers) are prepared to assimilate the mysteries' solutions. Research demonstrates several ways in which mysteries structure readers' narrative experiences (e.g., [4, 7]). For example,

information related to mysteries remains prominent in readers' discourse representations when the mystery remains unresolved.

The participatory perspective has also generated research that documents the impact of participatory responses on the time course with which readers assimilate narrative outcomes. For example, Jacovina and Gerrig [6] suggested that readers encode mental preferences when characters make decisions; those preferences affect how readers comprehend narrative outcomes. To test these predictions, Jacovina and Gerrig wrote short narratives in which characters had to make a decision (e.g., whether Sandy should invest her holiday bonus in a safe stock or a risky stock recommended on a blog). Jacovina and Gerrig found that participants were slower to read outcomes sentences that presented mismatches between their preferred action and the outcome. These data suggest that readers encode participatory responses regarding characters' actions as they read narratives, and that these responses influence the time course with which they comprehend narrative events.

The participatory perspective also casts light on circumstances in which readers' narrative experiences change their beliefs, attitudes, and behaviors. Researchers have provided a broad range of demonstrations that narrative experiences bring about persuasion (e.g., [8, 5]). The participatory perspective expands theoretical analysis of narrative persuasion by drawing on concepts from social learning theory [1]. Social learning theory suggests that people often learn by observing other individuals who obtain punishment or reinforcement. Because readers and viewers function as participants in narrative worlds, they also experience the vicarious influence of characters' outcomes. Thus, readers' beliefs, attitudes, and behaviors may change as they observe characters' fates.

References
---

1   A. Bandura. *Social learning theory*. Prentice-Hall, Englewood Cliffs, NJ, 1977.

2   M. A. Bezdek, J. E. Foy, and R. J. Gerrig. *"Run for it!". Viewers' participatory responses to film narratives*. Submitted for publication, 2013.

3   R. J. Gerrig and M. E. Jacovina. Reader participation in the experience of narrative. In B. H. Ross, editor, *The Psychology of Learning and Motivation. Advances in Research and Theory. Volume Fifty-One*, pages 223–254, Burlington, MA, 2009. Academic Press.

4   R. J. Gerrig, J. Love, and G. McKoon. Waiting for Brandon: How readers respond to small mysteries. *Journal of Memory and Language*, 60:144–153, 2009.

5   M. C. Green and T. C. Brock. The role of transportation in the persuasiveness of public narratives. *Journal of Personality and Social Psychology*, 79:701–721, 2000.

6   M. E. Jacovina and R. J. Gerrig. How readers experience characters' decisions. *Memory & Cognition*, 38:753–761, 2010.

7   J. Love, G. McKoon, and R. J. Gerrig. Searching for Judy: How small mysteries affect narrative processes and memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36:790–796, 2010.

8   D. A. Prentice, R. J. Gerrig, and D. S. Bailis. What readers bring to the processing of fictional texts. *Psychonomic Bulletin & Review*, 4:416–420, 1997.

Plots as Summaries of Event Chains

Inderjeet Mani

Yahoo! Labs
Sunnyvale, CA, USA
`inderjeet.mani@gmail.com`

Abstract

The plot of a narrative addresses what happened, and why. While a number of interesting theories of plot have been explored, it has proved hard in narrative interpretation to automatically compute a representation of the plot. This talk describes how to build a representation of what happened by summarizing temporal chains of events that involve a particular protagonist. These chains, which are based on the work of Chambers, can be summarized by various methods, including pruning subgraphs in the representation. Linguistic challenges include habitual expressions and non-literal language. The talk concludes with suggestions for how to layer causal information on top of the representation of what happened.

CB-POCL: A Choice-Based Algorithm for Character Personality in Planning-based Narrative Generation*

Julio César Bahamón and R. Michael Young

Liquid Narrative Research Group
North Carolina State University
Raleigh, NC, USA
`{jcbahamo,young}@csc.ncsu.edu`

─── **Abstract** ───

The quality and believability of a story can be significantly enhanced by the presence of compelling characters. Characters can be made more compelling by the portrayal of a distinguishable personality. This paper presents an algorithm that formalizes an approach previously described for the incorporation of character personality in narrative that is automatically generated. The approach is based on a computational model that operationalizes personality as behavior that results from the choices made by characters in the course of a story. This operationalization is based on the Big Five personality structure and results from behavioral psychology studies that link behavior to personality traits.

1 Introduction

The presence of interesting and compelling characters is an essential element of narrative composition [14, 6]. Effective characters enable the audience to form a clear mental model of their beliefs, desires, intentions, and morality. This understanding of the characters can lead to a better understanding of the entire story and to a more effective delivery of its content.

Well-defined characters add to the complexity of a story, enhance its discourse, and are vital for the realization of crucial story elements, such as events and dialog [22, 14, 6]. Among the factors that contribute to the effective definition of a character we include: physical attributes, talents, emotions, beliefs, and personality. Characters should portray these factors in an interesting and believable manner. In particular, the depiction of a distinguishable personality is one of the key features that makes a character interesting and compelling.

Characters also play an essential narrative role as agents of change. Change can result both from a character's actions and from its reaction to the actions of others [6, 22]. Therefore, actions are one of the main techniques used by creative writers to define and describe fictional characters [14, 5, 22]. Further consideration of narrative structure, specifically plot points where branching occurs [4], indicates that choices made by characters can have a direct impact in determining the actions they perform. Additionally, choices may be linked to

---

specific personality traits, an idea supported by research in behavioral psychology that has found correlation between people's actions and their personality [21, 10].

We posit that the link between choice and personality can be used in narrative to enable the perception of specific personality traits. An audience that is made aware of the existence of multiple choices that are available to a character will form an opinion of such character's personality based on the available choices, the choice selection, and the events that provide a context to the choice. Our approach is thus based on a computational model that operationalizes personality as behavior that results from the choices made by characters in the course of a story.

This paper presents a precise formulation of an idea previously described in [3]. We provide a formalization of the elements used to incorporate choice in planning-based narrative generation. Additionally, we give a detailed description of an algorithm that considerably improves upon, expands, and refines the approach previously proposed. Finally, we discuss a scenario that illustrates in detail the operation of the algorithm during story generation.

2   Background and Related Work

Previous approaches to introduce character personality in automatically-generated narrative have focused primarily on a character's immediate reaction to events [18, 26]. In contrast, this research focuses on the story as a whole and in particular the role of observable actions and their effect on the mental model that the audience forms when experiencing a story.

2.1   AI Planning

Planning is an artificial intelligence technique used to solve problems by finding a sequence of actions to achieve a goal state from a given initial state [36, 32]. A classical planning problem is represented by three inputs: an initial world state, a desired goal state, and a set of available actions that enable transitioning between world states. One of the applications of AI planning is the automatic generation of narrative.

An approach used in planning is to search through the space of plans. Search tree nodes represent partially ordered plans and edges represent plan refinements [36]. The planning algorithm maintains a partially ordered sequence of actions and a list of causal links [33] that indicate when the effect of an action is required to establish the precondition of another. The planning process starts with a null plan, where the *start* action has no preconditions and its effects are the literals in the initial state. The *goal* action has as its preconditions the literals in the goal state and does not have any effects. The algorithm non-deterministically selects an open precondition and chooses an existing or new action that establishes it. When a new action is added to the plan, ordering constraints and causal links are updated if necessary. The planning process succeeds when all preconditions have been satisfied and all the threats have been resolved and fails when the plan structure or bindings become inconsistent.

2.2   Planning-based Narrative Generation

A significant body of work has been dedicated to the development of AI systems for the automatic generation of narrative. The use of AI planners to automatically generate stories was first introduced in systems such as TALE-SPIN [20]. Considerable effort has been dedicated since then to the development and improvement of techniques, algorithms, and architectures to enable the application of the problem solving capabilities of AI planners to the automatic generation of narrative that is both interesting and coherent [31, 29]. It is

important to note that the work presented in this paper follows an approach that is distinct from that used by Riedl and Young in the IPOCL planning algorithm [31]. IPOCL focuses on character intentionality by identifying goals that explain a character's actions, which is done without considering the character's personality. In contrast, the approach described in this paper focuses on the use of specific actions to enable the portrayal of specific personality traits. It is envisioned that both approaches can be complementary as part of a solution aimed at producing more coherent and interesting narratives.

The work of Lebowitz on the UNIVERSE system [15] focused on the generation of stories using a plan-based approach combined with predefined character models. The system uses author goals to control the story-generation mechanism and character goals to ensure that their actions reflect their personality and backgrounds. The use of character goals is dependent on the definition of detailed character models based on the concept of stereotypes.

Some of the approaches for the automatic generation of narrative have focused on the implementation of systems that direct the interaction among characters. Work by Assanie on the extension of synthetic characters based on the Soar QuakeBot environment [13] dealt primarily with providing agents with the ability to adjust to changing goals provided by an external narrative manager [2]. One of the design objectives was the implementation of characters able to exhibit behavior that would resemble improvisational actors. Work by Riedl and Stern on drama managers focused on the implementation of semi-autonomous agents that have the ability to *fail believably* [28]. The system uses various techniques to ensure that agents avoid situations that are in conflict with the goals set by the drama manager and also to behave in a manner that justifies agent failure due to conflict with a goal set by the drama manager.

2.3 Emotion Expressed Through Facial and Physical Gestures

Research by Loyall [17] focused on the creation of believable agents, defined as autonomous agents with a rich personality and properties similar to those of characters in the traditional arts (e.g., film). The *Hap* agent architecture was developed to provide a language that enables authors to describe agent personality in terms of goals and behaviors.

André et al. worked on the development of lifelike characters as a means to improve interaction between humans and virtual characters [1]. Their approach centers on the use of the Five Factor Model model of personality and the OCC [23] theory of emotion to control the affective state of a virtual agent. The affective state determines the behavior, physical gestures, and dialog used by the agent when it communicates with humans.

The work of Doce et al. applied the OCC theory of emotion [23] and the Five Factor Model [11] to create distinguishable personalities in virtual agents [8]. The authors developed a model of personality that applies traits described in the Five Factor model (e.g., extroversion) to affect specific cognitive and behavioral processes, such as coping mechanisms and bodily expressions. The OCC theory of emotion is used to generate emotional states that influence the agent's planning mechanism and physical expressions [24].

Recent work in the area of virtual agents has focused on a specific subset of character actions: utterances in dialog [18, 26]. Of particular interest to this research is the work of Mairesse and Walker on PERSONAGE, a natural language generator that can be configured to generate dialog to meet a predefined set of personality requirements [18]. PERSONAGE is built on an architecture that uses the Five Factor Model [11] to create a mapping between personality traits and dialog utterances.

Table 1 Big-Five Factors Mapped to Observable Behavior.

Factor	Likely Behaviors	Unlikely Behaviors
Agreeableness	Honesty	Aggression
	Responsibility	Confrontation
	Ambition	Sabotage
	Empathy	Irritability
	Generosity	Selfishness
Conscientiousness	Organization	Impulsiveness
	Dutifulness	Lack of ambition
	Achievement	Mischief
	Reliability	Anti-social
	Risk avoidance	Criminal

3 Computational Model of Character Personality

The goal of this research is to facilitate the inclusion of compelling characters in narrative that is automatically generated. We do so by producing character behaviors that adjust: to authorial goals, to story events, and to user interaction. We aim to produce characters that can be distinguished by the visible manifestations of their personality.

We use a behavior model based on the Big Five Personality structure defined by Goldberg [11], which provides the following factors for the classification of personality types:

1. Extroversion
2. Agreeableness
3. Dependability
4. Emotional Stability
5. Culture (or Openness).

Each factor is linked to personality traits that can be mapped to behavioral manifestations, according to results obtained by social psychologists Mehl et al. [21] and Funder and Sneed [10, 21, 10, 35, 19]. A sample mapping is shown in table 1.

Our model uses a solution based on a declarative approach in which character properties are used to dynamically choose the actions they perform. Additionally, the effects of actions are evaluated to gauge whether these show behavior consistent with a character's personality traits. Finally, actions are considered in conjunction with a contrasting set of alternatives.

4 Definitions

In this section we provide a set of definitions necessary in the description of our algorithm for choice-based character personality. We work within the context of a planning-based narrative generation system, where a plan data structure is used to represent the events of a story and the temporal and causal relationships between them [37, 31].

4.1 Planning-based Narrative Generation

These definitions focus on the use of a partial order planner to generate stories, they are based on the work of Young and Riedl [37, 31].

▶ **Definition 1** (State). A state is a conjunction of literals used to describe what is true and what is false in a story world.

Any literals not explicitly described in the initial state of the story world are assumed to be false (closed-world assumption [27]).

▶ **Definition 2** (Character Name). A character name is a constant symbol that represents a story agent. Character names are unique.

▶ **Definition 3** (Action Schema). An action schema is a template for an action that is possible in the story world. An action schema is a tuple $\langle \mathsf{ActionType}, P, E, V, \mathsf{MainCharacter} \rangle$ where $\mathsf{ActionType}$ is a unique identifier for the action, P is a set of literals that must be true prior to the execution of the action (preconditions), E is a set of literals established by the execution of the action (effects), V is the list of free variables used in the template, and $\mathsf{MainCharacter}$ is a special variable used in the action schema to designate the story character primarily responsible for the execution of the instantiated action. The value of $\mathsf{MainCharacter}$ can be null ($\emptyset$).

We use a STRIPS-like representation [9] to describe the set of actions that are available in the story world. Preconditions are represented as a conjunction of all-positive literals, whereas effects may be a conjunct of both positive and negative literals.

For example, the plan step used to represent the story event "the knight kills the evil wizard" is an instance of the following action schema:

```
(action
    :action-type kill
    :variables ?main ?char
    :main-character ?main
    :preconditions (has sword ?main) (alive ?char)
    :effects (dead ?char))
```

▶ **Definition 4** (Planning Domain). A planning domain Λ is the set of all action schemas available in the story world.

▶ **Definition 5** (Planning Problem). A planning problem is a tuple $\langle \Lambda, S_0, S_G, C \rangle$ where Λ is a planning domain, S_0 is a set of literals that specify an initial state of the story world, S_G is a set of literals that specify a goal state, and C is a set of character names available in the story world. Each $c \in C$ is a unique identifier.

▶ **Definition 6** (Binding Constraint). A binding constraint on a pair of free variables or constants (u, v) indicates that u and v must unify in any well-formed formula. A negated pair $\neg(u, v)$ indicates that u and v cannot unify in any well-formed formula.

For example, the following binding constraints could be used to represent two characters that take part in a story event:

$\beta = \{(\texttt{?main-character, the-knight}), (\texttt{?char, the-evil-wizard})\})$

▶ **Definition 7** (Step). A step describes an instance of an action schema in a plan. A step is a tuple $\langle \mathsf{ActionType}, \mathsf{StepID}, \mathsf{Pre}, \mathsf{Eff}, \beta \rangle$, where $\mathsf{ActionType}$ is the unique identifier of an action schema, StepID is an identifier for a step that is unique within the plan in which the step occurs, Pre is the set of step preconditions, Eff is the set of step effects, and β is a set of binding constraints on free variables in Pre.

▶ **Definition 8** (Ordering Constraints). An ordering constraint over two steps s_i and s_j is denoted $s_i < s_j$, where s_i and s_j are steps in a plan and s_i must execute before s_j does.

In this representation, plan steps are partially ordered with respect to time [33].

▶ **Definition 9** (Causal Link). A causal link is denoted $s_i \xrightarrow{p} s_j$, where s_i is a plan step with an effect p and s_j is plan step with a precondition p. Causal links are used to keep track of dependencies between actions that exist when one action establishes a literal that is a precondition for another.

▶ **Definition 10** (Plan). A plan is a tuple $\langle S, B, O, L \rangle$ where S is a set of unique step identifiers, B is a set of binding constraints on free variables in S, O is a set of ordering constraints, and L is a set of causal links. A plan P is complete if and only if the following conditions are true:

- Every precondition of every step in $P = \langle S, B, O, L \rangle$ is satisfied,
  i.e., $\forall s_i \in S, s_i = \langle \mathsf{Pre}, \mathsf{Eff}, B \rangle, \forall p \in \mathsf{Pre}$ there is a causal link $s_i \xrightarrow{p} s_j \in L$.
- All threats have been resolved, i.e., for all causal links in $s_i \xrightarrow{p} s_j \in L$, there is no step $s_k \in S$ such that $s_i < s_k < s_j$ is a valid ordering under O and that has an effect $\neg q$, where q unifies with p.

4.2 Character Choice

These definitions are based on an analysis of creative writing principles and narrative structure, which was introduced in previous work [3]. The analysis focused on the essential role that characters play in the composition of a story and in particular their importance for the realization of key story elements such as events and dialog. We extend the plan-based representation of a story in a manner that enables the implementation of an intelligent action selection mechanism used to model the choices for action made by characters during the course of a story.

▶ **Definition 11** (Viable Alternative). A viable alternative over literal p is a tuple $\langle P, \mathsf{Prob}, s_{\mathrm{need}}, p, a_i \rangle$, where $P = \langle S, B, O, L \rangle$, $\mathsf{Prob} = \langle \Lambda, S_0, S_G, C \rangle$ is a planning problem, $s_{\mathrm{need}} \in S$, p is a precondition of s_{need}, and a_i is an action schema in Λ. We say that schema a_i is a viable alternative for p in plan $P = \langle S, B, O, L \rangle$ just when $\exists e_i \in \mathsf{Eff}_{a_i}$ such that p and e_i unify and there is no effect in a_i that is the negation of a literal in the preconditions of s_{need} in the context of B.

▶ **Definition 12** (Unintended Consequence). An unintended consequence with respect to a viable alternative a_i is a tuple $\langle P, \mathsf{Prob}, s_{\mathrm{need}}, p, a_i, x_i \rangle$, where $P = \langle S, B, O, L \rangle$, $\mathsf{Prob} = \langle \Lambda, S_0, S_G, C \rangle$ is a planning problem, $s_{\mathrm{need}} \in S$, p is a precondition of s_{need}, a_i is a viable alternative over p, and x_i is an effect of a_i. We say that x_i is an unintended consequence of a_i just when $x_i \neq p$, and $\neg \exists q \in \mathsf{Pre}_{s_{\mathrm{need}}}$ such that x_i and q unify in the context of B.

▶ **Definition 13** (Branching Point). A branching point is a tuple $\langle P, s_b, s_i, p \rangle$, where $P = \langle S, B, O, L \rangle$, $s_b, s_i \in S$, and $p \in \mathsf{Pre}_{s_j}$. We say that s_b is a branching point in plan P just when the following conditions are true:

- Step s_b can be ordered immediately before s_i, i.e., $s_b < s_i$ is a valid ordering under O and $\neg \exists s_k \in S$ such that $s_b < s_k < s_i$ is a valid ordering.
- Step s_i is the first step in a causal chain of events $s_i, s_{i+1}, \ldots, s_{i+n}$
- There exists an effect $e_b \in \mathsf{Eff}_{s_b}$ that unifies with p in the context of B and p is a precondition for s_i
- Step s_b can be instantiated from a specific action schema that has been selected from a set that contains two or more viable alternatives over p.

A branching point represents a narrative event that advances the story by requiring progress toward the story's goal state through one of at least two possible story paths. This definition is based on the concept of story kernels proposed by Barthes [4, 6].

▶ **Definition 14** (Choice). A choice is a tuple $\langle \mathsf{Prob}, \mathsf{BranchingPoint}, A, p \rangle$, where $\mathsf{Prob} = \langle \Lambda, S_0, S_G, C \rangle$ is a planning problem, $\mathsf{BranchingPoint} = \langle P, s_b, s_i, p \rangle$, and A is a set of viable alternatives over p, such that any action schema $a_i \in A$ may be instantiated to create step s_b.

A choice occurs in the story generation process when a branching point exists. Thus, a single step in the plan can be the source of zero or many choices, for example: $\{\{s_b, \emptyset\}\}$ represents a step without choices and $\{\{s_b, a_1\}, \{s_b, a_2\}, \{s_b, a_3\}\}$ represents a step with multiple choices (where a_1, a_2, a_3 are viable alternatives for an effect p established by step s_b).

5   The CB-POCL Algorithm

A preliminary version of an algorithm to incorporate choice in a planning-based story generation process was previously introduced in [3]. However, that version was limited because it did not take into account changes needed to support choices in the plan structure. In particular, it did not provide a detailed process for evaluating viable alternatives and establishing their preconditions. The previous version also did not present details on the method used to select specific action schemata for the instantiation of plan steps; details were preliminary both in terms of action selection and in the evaluation of their differences. Further, the algorithm's definition of the components of a choice within the context of a plan, as well as the elements necessary to support it, was not precise enough to disambiguate the non-deterministic selection of actions in a conventional POCL algorithm from the choice-based selection introduced in the paper. Finally, the previous version of the algorithm presented a method for ranking actions that did not include much detail about its integration with a behavior-based personality model.

This paper presents CB-POCL, an expanded and improved version that addresses in detail the changes to the planning process (see Algorithm 1). The current version is based on the extension of a partial-order causal link planning (POCL) algorithm, such as UCPOP [36, 25], to ensure that choice is treated as a first-class object during narrative generation, i.e., the story structure and contents promote the existence of choices and make their existence evident to the audience. This approach builds on prior research by Young and his colleagues [37, 29, 31, 12]. CB-POCL addresses the limitations of the previous work because it provides a detailed and formalized description of the mechanism used to incorporate choices in stories generated by a POCL process. Further, the definition of the elements and plan structure components needed to represent a choice in a planning context presented here precisely disambiguates the differences between the process for action selection used by conventional POCL algorithms and the algorithm we are developing.

5.1   Algorithm Description

In conventional POCL algorithms, actions are chosen nondeterministically to address open preconditions (flaws) until a complete plan is constructed or the process fails [36]. In our approach, specific actions may be added to the plan to facilitate the depiction of personality traits assigned to a character, e.g., honesty. Furthermore, the plan is built to support the ability to replace such actions with alternatives that depict contrasting personality traits.

The initial invocation of the algorithm is similar to that of UCPOP [36, 25]. The values in $\langle S, B, O, L \rangle$ are initialized to represent the null plan for the planning problem, and the value of agenda is set to the list of conjuncts in the goal state. Additionally, Choices is initialized to $\{\{s_0, \emptyset\}, \{s_G, \emptyset\}\}$, to indicate the lack of choices before the initial step or after the final step. On subsequent invocations of the algorithm, if the agenda is empty, the process succeeds. When flaws remain, one is selected, and the planner attempts to find an action schema or existing plan step whose effects establish the necessary precondition.

In the CB-POCL algorithm the selection is informed by an evaluation of the effects of all valid actions and in particular the unintended consequences that may result from the execution of an action. This is intended to represent a character's choice to perform a specific action from a set of possible alternatives. Two key factors are used when considering actions: (1) the action is viable and (2) the action can be used to portray a character's personality. Action evaluation is assisted by the use of an oracle function, focusing on the effects of the action upon other story participants. The function gauges whether the effects of a given action show behavior consistent with specific personality traits. This information is then used to determine if the action is part of a group of viable alternatives with enough contrast to clearly demonstrate a choice. For example, a set of alternatives for obtaining money that includes {Work, Borrow, Beg} does not provide the same contrast as a set that includes {SellFarm, Steal, Earn}. When only one alternative is viable or not enough contrast exists, the planner proceeds with a standard POCL process.

A modified process is used when multiple alternatives are viable, and a choice is added to the plan. The purpose of the modifications is twofold: (1) ensure that the viability of all the alternatives in the set is maintained and (2) preserve the contrast provided by the plan step during subsequent invocations of the planning algorithm. The viability of all the alternatives is ensured in order to create plans where choices exist, i.e., any of the viable alternatives in a set can be replaced with a member of the same set and the plan remains complete. To ensure viability, the union of the preconditions of all the alternatives is added to the agenda and these are treated as any other flaw. To preserve contrast, we flag the causal link drawn from the effect of the added step and the step whose flaw is being resolved (s_{need}). The flag is used to prevent the use of the same effect to solve other flaws, which would lessen the significance and clarity of the choice. The step is flagged to prevent its use in the plan unless it is in circumstances where it may enhance the portrayal of personality characteristics (e.g., honesty, aggression, or responsibility).

One of the potential problems that may result from the addition of choices to the plan is the existence of orphan steps, i.e., plan steps whose only purpose is to establish a precondition for one of the viable alternatives introduced by a choice. The following procedure is used when dealing with choice-related flaws in order to prevent the instantiation of orphan steps:

1. The use of steps that already exist in the plan is favored over the addition of new steps.
2. If using an existing step is not possible, the Initial State Revision (ISR) algorithm developed by Riedl and Young [30] is applied. ISR partitions the initial state of a planning problem into three sets: atomic ground sentences known to be true ($\mathcal{T}$), atomic ground sentences known to be false ($\mathcal{F}$), and atomic sentences whose truth value is undetermined ($\mathcal{U}$). The $\mathcal{U}$ set represents knowledge about the story world for which the author does not have a preference. ISR eliminates the flaw by committing to the value of a sentence in $\mathcal{U}$, enabling a causal link from the choice to the initial step (s_0).
3. If revision of the initial state is not applicable, a new step is added to the plan.

Other aspects of the POCL process work as expected, including: bindings update, causal link protection, failure detection, and the recursive invocation of the algorithm. Details of the functions that implement modified processes are provided in the following sections.

Algorithm 1 CB-POCL($\langle S, B, O, L \rangle$, Agenda, Λ, Choices)

1: **if** agenda is empty **then return** $\langle S, B, O, L \rangle$ ▷ Success, return the plan
2: **end if**
3: Select a flaw from the **agenda**. Each flaw in the agenda is represented by a pair $\langle p, s_{\text{need}} \rangle$, where $s_{\text{need}} \in S$ and p is a precondition of s_{need}.
4: $S_{\text{add}} = \text{GetViableAlternatives}(p, s_{\text{need}}, \langle S, B, O, L \rangle)$, where each item in S_{add} is either a brand new step instantiated from one of the schemas in Λ or an existing plan step that can be consistently ordered prior to s_{need}.
5: **if** $S_{\text{add}} = \emptyset$ **then**
        return failure   ▷ No viable alternatives exist, a plan cannot be constructed.
6: **end if**
7: **if** $\text{Count}(S_{\text{add}}) > 1$ **then**   ▷ **Update the plan structure**
8:     Choice = true   ▷ Annotate the step to indicate that it is part of a choice
9:     s_{need}:choice = true
10: **end if**
                        ▷ **Non-deterministically choose one of the viable alternatives**
11: s_c, Contrast-set, Consistent-bindings = $\text{ChooseAction}(p, s_{\text{need}}, S_{\text{add}}, \langle S, B, O, L \rangle, \text{Rank})$
                        ▷ **Update causal links, bindings, and ordering constraints**
12: $L' = L \cup \{s_c \xrightarrow{p} s_{\text{need}}\}$
13: **if** Choice = true **then**
14:     Annotate the causal link, $s_c \xrightarrow{p} s_{\text{need}}$, to indicate that it is part of a choice
15: **end if**
16: $B' = B \cup \{(u, v) | (u, v) \in MGU(q, p, B) \text{ and } q \text{ is an effect of } s_c\}$
17: $O' = O \cup \{s_c < s_{\text{need}}\}$
18: **if** Choice = true **then**   ▷ **Update the list of available choices, if applicable**
19:     Choices' = Choices.
20:     **for each** viable alternative s_i in S_{add} **do**
21:         Let Choices' = Choices' $\cup \{s_{\text{need}}, s_i\}$
22:     **end for**
23: **end if**
24: $S' = S$ and agenda' = agenda   ▷ **Update the plan steps and the agenda**
25: **for** each alternative s_i in $(s_c \cup \text{Contrast-set})$ **do**
26:     **if** $s_i \notin S'$ **then**
27:         Add $\langle \text{preconditions}(s_i) \backslash MGU(q, p, B), s_i \rangle$ to agenda'
28:         Add $\text{preconditions}(s_i)$ to the preconditions list of s_c
29:     **end if**
30: **end for**
31: $B' = B' \cup \text{Consistent-bindings}$
32: Add s_c to S'
33: **for** each causal link $l = s_i \xrightarrow{p} s_j$ in L **do**   ▷ **Causal link protection**
34:     **for** each plan step s_t that threatens l **do**   ▷ Nondeterministically choose one
35:         **Promotion:** If consistent, let $O' = O' \cup \{s_j < s_t\}$
36:         **Demotion:** If consistent, let $O' = O' \cup \{s_t < s_i\}$
37:         **if** neither Promotion nor Demotion can be chosen **then return** failure
38:         **end if**
39:     **end for**
40: **end for**
41: **if** B' is inconsistent **then**   ▷ **Recursive invocation**
        return failure
42: **else**
43:     Call CB-POCL($\langle S', B', O', L' \rangle$, agenda', Λ, Choices')
44: **end if**

5.2 The GetViableAlternatives Function

The GetViableAlternatives($p, s_{\text{need}}, \langle S, B, O, L \rangle$) function returns a set of actions newly instantiated from schemas in the planning domain library (Λ) or existing steps from the current plan that have p in their effects. An action (a_i) selected as a viable alternative must also be performed by a character whose personality the author wants to portray. Viable alternatives are evaluated to gauge whether their execution results in unintended consequences that show personality traits for a character. Choices are added to the plan only when the set of viable alternatives provides enough contrast to clearly portray a character's personality. When contrast does not exist, an action is chosen nondeterministically following the standard POCL process. For example, when all the viable alternatives result in unintended consequences consistent with honest behavior (a trait of agreeable personalities).

Algorithm 2 The GetViableAlternatives($p, s_{\text{need}}, \langle S, B, O, L \rangle$) function

---

 1: $A = \emptyset$
 2: **if** s_{need}:main-character = null **then**   ▷ **Only consider steps by a main character**
 3:     $A =$ Nondeterministically choose any existing step $s_i \in S$
 4:         OR a new action instantiated from Λ, such that the ordering constraint
 5:         $s_i < s_{\text{need}}$ is consistent with O, s_i has an effect q that unifies with p given B,
 6:         and s_i is not a step used in a Choice.
        return A
 7: **else**   ▷ **Prepare the initial list of viable alternatives**
 8:     $A =$ Select all existing steps $s_i \in S$ or new actions instantiated from Λ,
 9:         such that the ordering constraint $s_i < s_{\text{need}}$ is consistent with O,
10:         s_i has an effect q that unifies with p given B, $s_i : choice = false$, and
11:         s_{need}:main-character unifies with s_i:main-character given B.
                     ▷ **Evaluate the alternatives to check for contrast**
12:     $C =$ The set of story characters such that s_{need}:main-character
13:         unifies with s_i:main-character given B
14:     $A' = $RankActions($A, C, p, s_{\text{need}}, \langle S, B, O, L \rangle$)
15:     **if** Count(A') ≥ 2 and Rank(A_{first}) $\neq$ Rank(A_{last}) **then**
            return A'
16:     **else**   ▷ **Contrast does not exist**
            return (Nondeterministically choose an action from A')
17:     **end if**
18: **end if**

5.3 The RankActions Function:

This function evaluates a set of viable alternatives to gauge their compliance with the personality traits of a character. The output is a list of actions in descending order of compliance (see algorithm 3). An action's compliance with a personality trait is based on how closely the action is indicative of behavior that is typically associated with such trait, e.g., *avoidance of others* in the case of *introversion.*

5.4 The ChooseAction Function

The ChooseAction($p, s_{\text{need}}, S_{\text{add}}, \langle S, B, O, L \rangle$) function uses a selection mechanism informed by the personality traits of the character that performs the action and also by the context in

---

Algorithm 3 RankActions(Actions, Characters, p, s_{need}, $\langle S, B, O, L \rangle$)

1: **for** each action a_i in Actions **do**
2:     Rank$_i$ = 0
                          ▷ **Evaluate the unintended consequences of the action**
3:      **for** each character c_j in Characters **do**
4:          **for** each effect e_k established by a_i
5:              where $e_k \neq p$ and e_k does not unify with any other precondition of s_{need} **do**
6:              **for** each personality trait t_l assigned to c_j **do**
7:                  Rank$_i$ = Rank$_i$ + EvaluateEffect(character, e_k, t_l, $\langle S, B, O, L \rangle$)
8:              **end for**
9:          **end for**
10:     **end for**
11: **end for**
12: A' = Sort Actions by their corresponding rank in descending order
13: **return** (A', Rank)

---

which such action takes place. The output of the function is a triple containing an action in S_{add} that represents the personality traits of a given character, a corresponding set of contrasting alternatives, and a set of bindings that ensure the consistency between the elements in the contrast set and s_{need}. These bindings are the minimal set of bindings that ensure that there are no conflicts within the set of all the preconditions of s_{need} and all the preconditions of the viable alternatives, thus guaranteeing their consistency with the plan. In this version of the algorithm, we use a behavior model that is a hand-crafted approximation of the Big Five structure [11].

Algorithm 4 The ChooseAction(p, s_{need}, S_{add}, $\langle S, B, O, L \rangle$, Rank) function

---

1: C = The set of story characters such that s_{need}:main-character
2:        unifies with s_i:main-character given B
                          ▷ **Invoke the oracle function**
3: Choice = Nondeterministically choose an action from S_{add} that portrays the personality
            traits of s_{need}:main-character within the current story context.
        ▷ **For each alternative that is consistent with the character's personality,**
    determine its contrasting set
4: Contrast-set = Select the subset of actions from S_{add} that provides the most
                contrast with Choice.
5: Consistent-bindings = Consult B and then compute the set of bindings that ensure
                consistency between all the elements in Contrast-set and s_{need}.
6: **return** (Choice, Contrast-set, Consistent-bindings)

5.5   The *Oracle* Function for Trait Evaluation

An action a_i is said to demonstrate a personality trait t_k if its unintended consequences result in a state of the world that is consistent with behavior typical of people who score high on such trait. An action a_i is said NOT to demonstrate personality trait t_k if its unintended consequences result in a state of the world that is consistent with behavior that is typical of people who score low on the same trait. In a similar manner, the preconditions

Table 2 Rule for Generous Behavior – Effect.

Rule Name	GiveUpItems
Rule Type	Effect
Applies To	Agreeableness: generosity
Parameters	An action's unintended consequence $\Phi(x_1, x_2, \ldots, x_n)$, where $\Phi()$ is a predicate and $x_1, x_2, \ldots, x_n$ are variables. The current plan, $\langle S, B, O, L \rangle$. A character, c. OBJECT() is a predicate, such that OBJECT$(x_i) = True$ if x_i represents a physical object in the story world, e.g., car, plane, sword, a bag of gold.

Description	▷ **If any of the variables is bound to an entity that is of type OBJECT**
	▷ **check whether the action's unintended consequence is the negation of a predicate**
	▷ **that is currently true. Then check the character's personality.**
	affectsObjects = False
	for each variable x_i in the action's unintended consequence **do**
	    if OBJECT(x_i) = True
	        affectsObjects = True
	    end if
	end for
	if affectsObjects = True AND
	   One of the variables in $(x_1, x_2, \ldots, x_n)$ represents character c AND
	   None of the variables in $(x_1, x_2, \ldots, x_n)$ represents characters other than c AND
	   $\Phi(x_1, x_2, \ldots, x_n)$ is the negation of a predicate that is currently true **then**
	    if c is Highly-Agreeable **then**
	        return 1
	    else
	        return -1
	    end if
	else
	    return 0
	end if

that must exist prior to the execution of an action can demonstrate behavior consistent with specific personality traits. An action a_i is said to demonstrate a personality trait t_k if the preconditions needed for its execution indicate a state of the world that is the result of behavior typical of people who score high on such trait. An action a_i is said NOT to demonstrate personality trait t_k if the preconditions needed for its execution indicate a state of the world that is the result of behavior typical of people who score low on such trait.

Initially we focus on the basic case in which given one personality trait, for a specific character, the planner evaluates the viable alternatives and selects the one most consistent with the trait. Note that effects and preconditions common to all the actions in a set of viable alternatives are ignored because they do not provide information to help differentiate the alternatives from each other.

5.5.1 Effects Evaluation:

The purpose of the EvaluateEffect(Character, e_j, t_k, $\langle S, B, O, L \rangle$) function is to gauge whether an action's effect demonstrates behavior that is consistent with a personality trait. We do this by considering the resulting state of the story if such an effect is established. For example, an action that results in the loss of treasure would be consistent with the behavior

of a highly agreeable character, who is expected to act in a generous manner as indicated in table 1. In contrast, a non-agreeable character would be less willing to give up treasure and instead be more inclined to make choices that do not result in a state in which possessions are lost. This function applies all the rules of the appropriate type to the effect and character, in the context of the current partial plan, and returns an average of the score computed after applying all rules.

To enable the evaluation of effects, we use a declarative mapping between personality traits and behaviors. This mapping is implemented as an extensible set of behavioral rules that when applied can be used to evaluate individual effects. The rules are placed in a domain-independent knowledge base that can be maintained by the user. An example of a behavioral rule is shown in table 2.

6   Sample Scenario: A Hero's Quest

6.1   Story Argument

Princess Kayla has been kidnapped by the evil dragon Gomez. Joe, a local townsman, has been asked to rescue her. If Joe is to succeed in his quest, he must first obtain a sword that will enable him to slay Gomez and perform the rescue.

In this scenario, we consider Joe to be highly agreeable. Behavior that can be expected of him includes: honesty, responsibility, interaction, empathy, and generosity. Likewise, Joe is unlikely to exhibit behaviors such as aggression, sabotage, irritability, and selfishness. The story plan will be constructed to include choices that show behavior consistent with Joe's personality.

6.2   Scenario Alternatives

Consider a scenario where in order to obtain the sword Joe faces a choice from three alternatives: (1) he can buy the sword, (2) he can steal the sword through treachery, or (3) he can attack someone to take away the sword. Each one of these alternatives has a specific set of effects and preconditions that when added to the plan create a different version of the story. The contrast between the story produced by the chosen alternative and those that would result from the selection of the other alternatives is used to represent the act of making a choice – Joe's choice in this context. Furthermore, the selection of a specific alternative is determined by the personality traits that have been previously assigned to Joe. The plan is also constructed to support the alternate story paths that would result from the contrasting alternatives.

6.3   Action Schemata

The following action schemata are used to implement the alternatives in the scenario. Note that not all the actions used in the sample plan are included, just those with particular relevance to this discussion.

```
(action
        :action-type Buy
        :variables ?item ?buyer ?seller
        :main-character ?buyer
        :preconditions (has ?item ?seller)
                       (has gold ?buyer)
                       (knows ?buyer ?seller)
        :effects (has gold ?seller)
                 (not (has gold ?buyer))
                 (has ?item ?buyer))

(action
        :action-type Steal
        :variables ?item ?char ?victim
        :main-character ?char
        :preconditions (has ?item ?victim)
                       (confused ?victim)
        :effects (upset ?victim)
                 (has ?item ?char)
                 (not (has ?item ?victim)))


(action
        :action-type Fight
        :variables ?attacker ?victim
        :main-character ?attacker
        :preconditions (knows ?victim ?attacker)
                       (has sword ?victim)
        :effects (hurt ?victim)
                 (has sword ?attacker)
                 (not (has sword ?victim)))

(action
        :action-type Polish
        :variables ?char ?item ?location
        :main-character ?char
        :preconditions (at village ?char )
        :effects (has gold ?char))
```

6.4 Story Generation Walkthrough

Consider the partial story plan shown in figure 1, which shows the *Start* and *End* steps of the plan. At this point only the initial invocation of the algorithm has taken place and no steps have been added to address any open preconditions. After a few recursive invocations of CB-POCL, we may have the partial plan shown in figure 2. Note that the preconditions of the *End* step have been addressed by the addition of two new plan steps: *Slay* and *Rescue*. However, the addition of the new steps results in new preconditions that must also be addressed, in particular (`has sword Joe`) for the *Slay* step.

Figure 1 Sample Scenario – Initial Story Plan. In this figure, the rectangular boxes indicate steps in the plan. Time flows roughly horizontally downward. Initially, the start and end steps are the only steps in the plan.
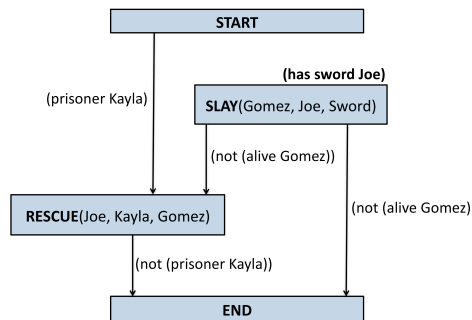
Figure 2 Sample Scenario – Partial Story Plan. In this figure, the rectangular boxes indicate steps in the plan. Time flows roughly horizontally downward. Solid arcs from one step to another indicate a causal link between the source step's effects and the destination step's preconditions.

Figure 3 Partial Story Plan with Viable Alternatives Under Consideration. In this figure, steps drawn with dashed lines indicate viable alternatives that do not actually appear in the final plan used to generate the story. These are flagged as being part of a choice, which makes it possible to use them during discourse generation to help convey the existence of such choice. Similarly, dashed arcs indicate causal links that could potentially provide the needed condition from a viable alternative to a specific step; however, these links do not appear in the final plan. A box with a double border indicates a choice.

On a later invocation of the algorithm, we could have the situation shown in figure 3. The graph includes a set of viable alternatives being considered by CB-POCL to address the (`has Sword Joe`) precondition of the *Slay* step. Note that the addition of any of the alternatives – *Buy, Steal, or Fight* – establishes the necessary precondition to enable the execution of the *Slay* step. However, each of the alternatives has different effects not needed for the execution of the *Slay* step (unintended consequences). Evaluation of the unintended consequences by the oracle function would determine that the *Buy* alternative is more representative of a character with a highly agreeable personality and that *Steal* and *Fight* are appropriate contrasting options. Additionally, note that the oracle function also considers the (`has gold Joe`) precondition of the *Buy* option, which can be established by the *Polish* action further depicting behavior in accordance with that expected of a highly agreeable character. A highly agreeable character will choose the option of buying the sword, even though the precondition of having gold requires the extra effort of polishing armor and the buy action represents a loss of personal treasure. In contrast, a non-agreeable character will choose one of the remaining options because neither performing extra work to obtain gold nor giving it up in exchange for the sword is consistent with a non-agreeable personality.

Figure 4 shows the plan after the (`has Sword Joe`) precondition of the *Slay* step has been addressed by the addition of the *Buy* step. Steps to enable the execution of the *Buy* step have also been added. Alternate story paths based on the other two viable alternatives (*Steal, Fight*) are shown with a dashed border to illustrate that they could be part of the plan, yet they are not. Steps shown with a dashed border are not included in the final plan used to generate the story; however, since these are flagged as being part of a choice it may be possible to use them during discourse generation to help convey the existence of such choice. Also, the ordering of the *Travel* actions is important in this scenario because the main character must have visited all three locations (village, tavern, forest), which cannot be done simultaneously hence proper sequencing of these steps is required. Note that CB-POCL will ensure that the necessary supporting steps can be added to the completed plan by also addressing the corresponding preconditions: (`confused Blacksmith`) and (`knows Blacksmith Joe`). ISR is applied to set the precondition that indicates who has the sword, which is needed for any of the alternatives being considered.

7   Discussion and Future Work

7.1   Algorithm Limitations

The CB-POCL algorithm currently does not guarantee that all the steps added to the plan are consistent with a character's personality traits. For example, in the scenario previously described the *Use Potion* step needed to enable the *Steal* viable alternative is not consistent with the agreeable personality of the character (Joe). A more robust version of the process used to construct the plan should take into account the entire chain of events that result from the addition of a viable alternative and not just the immediate steps. Additionally, plans produced by the algorithm may contain truncated action sequences that have been added only to establish preconditions for a viable alternative. Even though CB-POCL is designed to reduce the addition of isolated actions by favoring the use existing plan steps or applying Initial State Revision, there is no guarantee that truncated action sequences will never be part of a plan. Future work on this research will include exploring methods to further mitigate and ideally eliminate the occurrence of such sequences. Finally, this work in its current form does not address the limitations that stem from the inability to determine the completeness of a plan until the planning process has stopped. Working with a complete plan would be the ideal environment to identify the optimal structure and contents

■ **Figure 4** Partial Plan Showing Alternate Story Paths Based on Choice. In this figure, dotted arcs indicate causal links that are established from the initial state via initial state revision. In this figure, steps drawn with dashed lines indicate viable alternatives that do not actually appear in the final plan used to generate the story. These are flagged as being part of a choice, which makes it possible to use them during discourse generation to help convey the existence of such choice. Similarly, dashed arcs indicate causal links that could potentially provide the needed condition from a viable alternative to a specific step; however, these links do not appear in the final plan. A box with a double border indicates a choice.

to portray specific personality traits. However, since this is not feasible, future versions of CB-POCL should incorporate techniques to obtain an approximation of the knowledge needed to improve the plan structure while plan refinement is still in progress.

7.2   Planned Algorithm Enhancements

The current version of the algorithm does not consider action preconditions during the initial evaluation of viable alternatives. Preconditions are necessary for a personality-based action selection process when we reflect on their relevance to the state of the story world prior to the execution of an action. The need to establish a particular precondition may also result in the addition of other actions to the plan, thus providing more opportunities to portray a character's behavior. In addition to the elements considered for individual choices, it is necessary to consider changes to the process used to construct the plan structure. The algorithm should enable operations such as: changing the ordering of actions currently in the plan, increasing or reducing action decomposition, changing or introducing a causal chain of events, and dynamically introducing behavior-related constraints. These modifications would facilitate the construction of plans that treat choice as a first-class object. Finally, the action selection process presented in this paper uses a step-wise approach when reasoning about choices during story construction. An improved version of the algorithm should extend the model to enable reasoning about sub-plans and action decomposition.

7.3   Behavioral Model

The *oracle* function mentioned in this document must be replaced by a robust behavioral model that is less dependent on knowledge engineering and thus more flexible in its use.

Such a model will implement an operationalization of principles identified by behavioral psychology research, such as that of Mehl et al. and Funder and Sneed [21, 10].

One of the possibilities being considered for this component is the implementation of a model based on the goal-based hierarchical taxonomy proposed by Chulef, Reed, and Walsh [7]. The taxonomy includes 135 achievement goals identified through experimentation. Additionally, the authors propose conceptual relationships between personality traits and specific goals in the hierarchy.

Applying the concept of goal hierarchies, personality could be implemented as a set of goals of varying importance. For example, being likable, following social norm, and safety. In this model, the importance of specific goals can be determined by the personality traits they are associated with. Furthermore, characters could have goals of two types: communal and individual [34]. Another factor to consider is that goals require stimuli that activates them and resources that enable their achievement [34]. These affect the likelihood of enacting behavior. Goals can also be seen as incentives or deterrents for behavior. Finally, conflict may also have a role in the expression of personality. For example, when a character's goals are frustrated this could result in aggressive behavior. Highly social-goal oriented individuals are less likely to behave in an uncivil manner when they experience frustration [16].

8   Conclusion

This paper presents the next step toward the development of an intelligent mechanism that enables the automatic generation of narrative that elicits the perception of distinct character personalities without the need of a labor-intensive process. The solution is based on a declarative approach, in which character properties and the story context are used to model the choices that determine the set of actions that they perform in the course of the story.

The approach described extends a conventional POCL algorithm to ensure that choice is treated as a first-class object. The planner's data structures and supporting processes are modified to this effect. Additionally, the nondeterministic selection of actions is replaced by a mechanism that incorporates choice as a key component of the story generation process. Furthermore, personality is operationalized in terms of the behavior that results from the choices made by characters as they perform their role in the narrative.

Finally, an essential part of our future work will be to validate whether narratives generated using the CB-POCL algorithm result in characters whose personality traits are distinguishable to human audiences. We plan to conduct user studies designed to measure whether character behavior generated by the algorithm elicits in the audience the perception of corresponding personality traits.

References

1   Elisabeth André, Martin Klesen, Patrick Gebhard, Steve Allen, and Thomas Rist. Integrating models of personality and emotions into lifelike characters. In Ana Paiva, editor, *Affective Interactions: Toward a New Generation of Computer Interfaces?*, number 1814 in Lecture Notes in Computer Science, pages 150–165. Springer, 2000.

2   Mazin Assanie. Directable synthetic characters. In Ken Forbus and Magy Seif El-Nasr, editors, *Proceedings of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, number SS-02-01 in AAAI Technical Reports, pages 1–7, 2002.

3   Julio César Bahamón and R. Michael Young. A choice-based model of character personality in narrative. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 164–168, İstanbul, 2012.

4   Roland Barthes and Lionel Duisit. An introduction to the structural analysis of narrative. *New Literary History*, 6(2):237–272, 1975.

5   Colin Bulman. *Creative Writing: a Guide and Glossary to Fiction Writing.* Polity, Cambridge, UK, 2007.

6   Seymour Benjamin Chatman. *Story and Discourse: Narrative Structure in Fiction and Film.* Cornell University Press, Ithaca, NY, 1978.

7   Ada S. Chulef, Stephen J. Read, and David A. Walsh. A hierarchical taxonomy of human goals. *Motivation and Emotion*, 25(3):191–232, 2001.

8   Tiago Doce, João Dias, Rui Prada, and Ana Paiva. Creating individual agents through personality traits. In Jan Allbeck, Norman Badler, Timothy Bickmore, Catherine Pelachaud, and Alla Safonova, editors, *Intelligent Virtual Agents, 10th International Conference, IVA 2010, Philadelphia, PA, USA, September 20-22, 2010, Proceedings*, number 6356 in Lecture Notes in Computer Science, pages 257–264. Springer, 2010.

9   Richard E. Fikes and Nils J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3-4):189–208, 1971.

10   David C. Funder and Carl D. Sneed. Behavioral manifestations of personality: An ecological approach to judgmental accuracy. *Journal of Personality and Social Psychology*, 64(3):479–490, 1993.

11   Lewis R. Goldberg. An alternative "description of personality": The big-five factor structure. *Journal of Personality and Social Psychology*, 59(6):1216, 1990.

12   Justin Harris and R. Michael Young. Proactive mediation in plan-based narrative environments. In Themis Panayiotopoulos, Jonathan Gratch, Ruth Aylett, Daniel Ballin, Patrick Olivier, and Thomas Rist, editors, *International Conference on Intelligent Virtual Agents*, number 3661 in Lecture Notes in Computer Science, pages 292–304. Springer, 2005.

13   John E. Laird and Randolph M. Jones. Building advanced autonomous AI systems for large scale real time simulations. Presented at the Computer Games Development Conference, Long Beach, California, 1998.

14   Alice LaPlante. *The Making of Story: a Norton Guide to Creative Writing.* W. W. Norton, New York, 2007.

15   Michael Lebowitz. Creating characters in a story-telling universe. *Poetics*, 13(3):171–194, 1984.

16   Wu Liu, Shu-Cheng Steve Chi, Ray Friedman, and Ming-Hong Tsai. Explaining incivility in the workplace: The effects of personality and culture. *Negotiation and Conflict Management Research*, 2(2):164–184, 2009.

17   A. Bryan Loyall. *Believable agents: building interactive personalities.* PhD thesis, Carnegie Mellon University, 1997.

18   François Mairesse and Marilyn Walker. PERSONAGE: Personality generation for dialogue. In Annie Zaenen and Antal van den Bosch, editors, *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 496–503, Prague, Czech Republic, 2007. Association for Computational Linguistics.

19   Robert R. McCrae and Oliver P. John. An introduction to the five-factor model and its applications. *Journal of Personality*, 60(2):175–215, 1992.

20   James R. Meehan. TALE-SPIN, an interactive program that writes stories. In Dabbala Rajagopal Reddy, editor, *Proceedings of the 5th International Joint Conference on Artificial Intelligence. Cambridge, MA, August 1977*, pages 91–98. William Kaufmann, 1977.

21   Matthias R. Mehl, Samuel D. Gosling, and James W. Pennebaker. Personality in its natural habitat: Manifestations and implicit folk theories of personality in daily life. *Journal of Personality and Social Psychology*, 90(5):862–877, 2006.

22   Matt Morrison. *Key concepts in creative writing.* Palgrave Macmillan, Houndmills, Basingstoke, Hampshire; New York, 2010.

23   Andrew Ortony, Gerald L. Clore, and Allan Collins. *The cognitive structure of emotions.* Cambridge University Press, 1990.

24   Catherine Pelachaud. Multimodal expressive embodied conversational agents. In Hongjiang Zhang, Tat-Seng Chua, Ralf Steinmetz, Mohan S. Kankanhalli, and Lynn Wilcox, editors, *Proceedings of the 13th Annual ACM International Conference on Multimedia*, pages 683–689. Association for Computing Machinery, 2005.

25   J. Scott Penberthy and Daniel S. Weld. UCPOP: A sound, complete, partial order planner for ADL. In Bernhard Nebel, Charles Rich, and William R. Swartout, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, pages 103–114. Morgan Kaufmann, 1992.

26   Aaron A. Reed, Ben Samuel, Anne Sullivan, Ricky Grant, April Grow, Justin Lazaro, Jennifer Mahal, Sri Kurniawan, Marilyn Walker, and Noah Wardrip-Fruin. A step towards the future of role-playing games: The SpyFeet mobile RPG project. In Vadim Bulitko and Mark O. Riedl, editors, *Proceedings of the Seventh AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment.* AAAI Press, 2011.

27   R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1–2):81–132, 1980.

28   Mark O. Riedl and A. Stern. Failing believably: Toward drama management with autonomous actors in interactive narratives. In Stefan Göbel, Rainer Malkewitz, and Ido Iurgel, editors, *Technologies for Interactive Digital Storytelling and Entertainment, Third International Conference, TIDSE 2006, Darmstadt, Germany, December 4–6, 2006*, number 4326 in Lecture Notes in Computer Science, pages 195–206. Springer, 2006.

29   Mark O. Riedl and R. Michael Young. Character-focused narrative generation for execution in virtual worlds. In Olivier Balet, Gérard Subsol, and Patrice Torguet, editors, *Virtual Storytelling: Using Virtual Reality Technologies for Storytelling, International Conference, ICVS 2001, Avignon, France, September 27–28, 2001, Proceedings*, number 2897 in Lecture Notes in Computer Science, pages 47–56. Springer, 2003.

30   Mark O. Riedl and R Michael Young. Open-world planning for story generation. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30-August 5, 2005*, pages 1719–1720. Professional Book Center, 2005.

31   Mark O. Riedl and R. Michael Young. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:164–167, 2010.

32   Stuart J. Russell and Peter Norvig. *Artificial intelligence: a modern approach.* Prentice Hall/Pearson Education, Upper Saddle River, N. J., 3rd edition, 2010.

33   Earl D. Sacerdoti. The nonlinear nature of plans. In *Advance Papers of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, Georgia, USSR, 3–8 September 1975*, pages 206–214, 1975.

34   Mei Si, Stacy C. Marsella, and David V. Pynadath. Thespian: Using multi-agent fitting to craft interactive drama. In Michal Pechoucek, Donald Steiner, and Simon Thompson, editors, *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 21–28. Association for Computing Machinery, 2005.

35   Jason Vandenberghe. The five domains of play: Mapping the five-factor model in psychology to game design. *Game Developer Magazine*, 19(5):44–46, 2012.

36   Daniel S. Weld. An introduction to least commitment planning. *AI Magazine*, 15(4):27–61, 1994.

37   R. Michael Young. Notes on the use of plan structures in the creation of interactive plot. In M. Mateas and P. Sengers, editors, *Narrative Intelligence: Papers from the AAAI Fall Symposium*, number FS-99-01 in AAAI Technical Reports, pages 164–167, 1999.

Cognitive Interpretation of Everyday Activities — Toward Perceptual Narrative Based Visuo-Spatial Scene Interpretation

Mehul Bhatt[1,2], Jakob Suchan[1,2], and Carl Schultz[1,2]

1   **Cognitive Systems**
    University of Bremen
    Bremen, Germany
    `{bhatt,jsuchan,cschultz}@informatik.uni-bremen.de`
2   **Sonderforschungsbereich Transregional Collaborative Research Center 8:**
    Spatial Cognition
    University of Bremen
    Bremen, Germany

---- **Abstract** ----

We position a narrative-centred computational model for high-level knowledge representation and reasoning in the context of a range of assistive technologies concerned with *visuo-spatial perception and cognition* tasks. Our proposed narrative model encompasses aspects such as *space, events, actions, change, and interaction* from the viewpoint of commonsense reasoning and learning in large-scale cognitive systems. The broad focus of this paper is on the domain of *human-activity interpretation* in smart environments, ambient intelligence etc. In the backdrop of a *smart meeting cinematography* domain, we position the proposed narrative model, preliminary work on perceptual narrativisation, and the immediate outlook on constructing general-purpose open-source tools for perceptual narrativisation.

1   Introduction: Cognitive Interpretation by Narrativisation

Narratives have been a focus on study from several perspectives, most prominently from the viewpoint of language, literature, and computational linguistics; see for instance, discourse analysis and computational narratology [1, 13, 14, 12]. From the viewpoint of commonsense reasoning, and closely related to the computational models of narrative perspective, is the position of researchers in logics of *action and change*; here, narratives are interpreted as "*a sequence of events about which we may have incomplete, conflicting or incorrect information*" [16, 18]. As per McCarthy [15], "*a narrative tells what happened, but any narrative can only tell a certain amount. A narrative will usually give facts about the future of a situation that are not just consequences of projection from an initial situation*". The interpretation of narrative knowledge in this paper is based on these characterisations, especially in regard to the commonsense representation and reasoning tasks that accrue whilst modelling and reasoning about the perceptually grounded, narrativised epistemic state of an autonomous
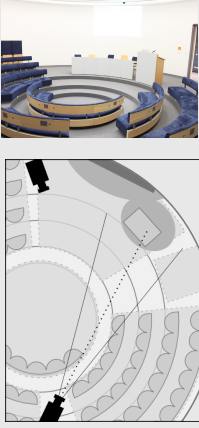
Listing 1. Smart Meeting Cinematography

The smart meeting cinematography domain focusses on professional situations such as meetings and seminars. A basic task is to automatically produce dynamic recordings of interactive discussions, debates, presentations involving interacting people who use more than one communication modality such as hand-gestures (e.g., raising one's hand for a question, applause), voice and interruption, electronic apparatus (e.g., pressing of a button), movement (e.g., standing-up) and so forth. The scenario consists of people-tracking, gesture identification closed under a context-specific taxonomy, and also involves real-time dynamic collaborative co-ordination and self-control of pan-tilt-zoom (PTZ) cameras in a *sensing-planning-acting* loop. The long-term vision is to benchmark with respect to the capabilities of human-cinematographers, real-time video editors, surveillance personnel to record and semantically annotate individual and group activity (e.g., for summarisation, story-book format digital media and promo generation). [5]

agent pertaining to *space, actions, events, and change* [2]. In particular, this encompasses a range of inference patterns such as: (a) spatio-temporal abduction for scenario and narrative completion [3]; (b) integrated inductive-abductive reasoning with narrative knowledge [7]; (c) narrative-based postdiction for abnormality detection and planning [8].

Perceptual Narratives.   These are declarative models of visual, auditory, haptic and other observations in the real world that are obtained via artificial sensors and / or human input. Declarative models of perceptual narratives can be used for interpretation and control tasks in the course of assistive technologies in everyday life and work scenarios, e.g., behaviour interpretation, robotic plan generation, semantic model generation from video, ambient intelligence and smart environments (e.g., see narrative based models in [10, 11, 17, 3, 7, 8]).

High-Level Cognitive Interpretation and Control.   Our research is especially concerned with large-scale cognitive interaction systems where high-level perceptual sense-making, planning, and control constitutes one of many AI sub-components guiding other low-level control and attention tasks. As an example, consider the *smart meeting cinematography* domain in Listing 1. In this domain, *perceptual narratives* as in figure 1 are generated based on perceived spatial change interpreted as interactions of humans in the environment. Such narratives explaining the ongoing activities are needed to anticipate changes in the environment, as well as to appropriately influence the real-time control of the camera system. To convey the meaning of the presentation and the speakers interactions with a projection, the camera has to capture the scene including the speakers gestures, slides, and the audience. E.g., in figure 1, when the speaker explains the slides, the camera has to capture the speaker and the corresponding information on the slides. To this end, the camera records an overview shot capturing the speaker and the projection, and zooms on the particular element when the speaker explains it in detail, to allow the viewer to follow the presentation and to get the necessary information. When the speaker continues the talk, the camera focuses on the speaker to omit unnecessary and distracting information. To capture reactions of the audience, e.g., comments, questions or applause, the camera records an overview of the attending people or close-up shots of the commenting or asking person.
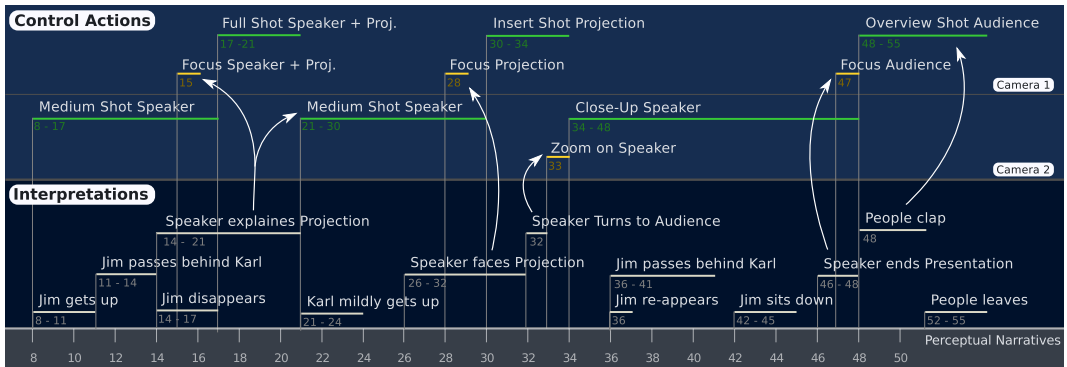
Figure 1 Cognitive Interpretation and Control by Perceptual Narrativisation.
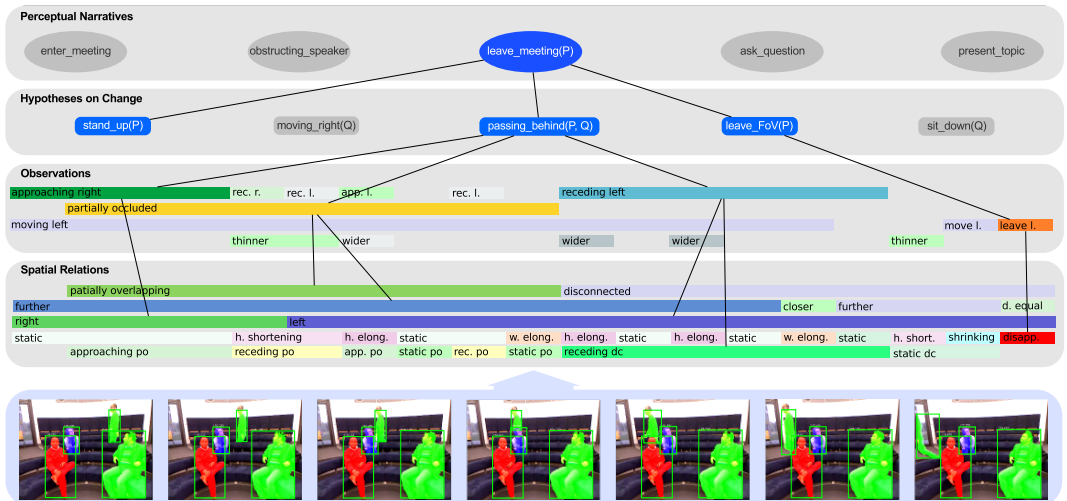
Figure 2 Perceptual Narratives of Depth, Space, and Motion.

2   Perceptual Narrative Generation for Activity Interpretation

Systems that monitor and interact with an environment populated by humans and other artefacts require a formal means for representing and reasoning about spatio-temporal, event and action based phenomena that are grounded to real public and private scenarios (e.g., logistical processes, activities of everyday living) of the environment being modelled. A fundamental requirement within such application domains is the representation of *dynamic* knowledge pertaining to the spatial aspects of the environment within which an agent, system, or robot is functional. This translates to the need to explicitly represent and reason about dynamic spatial configurations or scenes and, for real world problems, integrated reasoning about perceptual narratives of *space, actions, and change* [2]. With these modelling primitives, the ability to perform *predictive* and *explanatory* analyses on the basis of sensory data is crucial for creating a useful intelligent function within such environments [7].

Perceptual Narratives of Space and Motion

To understand the nature of perceptual narratives (of space, and motion), consider the aforediscussed work-in-progress domain of *smart meeting cinematography* (Listing 1). The particular infrastructural setup for the example presented herein consists of Pan-Tilt-Zoom (PTZ) capable cameras, depth sensors (Kinect), and a low-level vision module for people tracking (whole body, hand gesture, movement) customised on the basis of open-source algorithms and software. With respect to this setup, declaratively grounded perceptual narratives capturing the information in figure 1 is developed on the basis of a commonsense theory of *qualitative space* (Listing 2), and interpretation of motion as qualitative spatial change [9]. In particular, the overall model as depicted in figure 2 consists of:

Space and Motion: A theory to declaratively reason about qualitative spatial relations (e.g., topology, orientation), and qualitative motion perceived in the environment and interpret changes as domain dependent observations in the context of everyday activities involving humans and artefacts.

Explanation of (Spatial) Change: Hypothesising real-world (inter)actions of individuals explaining the observations by integrating the qualitative theory with a learning method (e.g., Bayesian and Markov based (logic) learning) to incorporate uncertainty in the interpretation of observation sequences.

Semantic characterisation: as a result of the aforementioned, real-time generation of declarative narratives of perceptual data (e.g., RGB-D) obtained directly from people/object tracking algorithms.

Hypothesised object relations can be seen as building blocks to form complex interactions that are semantically interpreted as activities in the context of the domain. As an example consider the sequence of observations in the meeting environment depicted in figure 2.

> Region P elongates vertically, region P approaches region Q from the right, region P partially overlaps with region Q while P being further away from the observer than Q, region P moves left, region P recedes from region Q at the left, region P gets disconnected from region Q, region P disappears at the left border of the field of view

To explain these observations in the 'context' of the meeting situation we make hypothesis about possible interactions in the real world.

> Person P stands up, passes behind person Q while moving towards the exit and leaves the room.

Listing 2.   Qualitative Abstractions of Space and Motion

To represent space and spatial change we consider spatio-temporal relations [6] holding between individuals in the environment, i.e., *topology, orientation, size, movement.* Combinations of spatial and temporal relations serve as observations describing perceived phenomena in the real world.

The theory is implemented using **CLP(QS)** [4], which is a *declarative spatial reasoning framework* that can be used for representing and reasoning about high-level, qualitative spatial knowledge about the world. CLP(QS) implements the semantics of qualitative spatial calculi within a constraint logic programming framework (amongst other things, this makes it possible to use spatial entities and relations between them as native entities). Furthermore it provides a declarative interface to qualitative and geometric spatial representation and reasoning capabilities such that these may be integrated with general knowledge representation and reasoning (KR) frameworks in artificial intelligence.

The **semantic interpretation of activities** from video, depth (e.g., time-of-flight devices such as Kinect), and other forms of sensory input requires the representational and inferential mediation of qualitative abstractions of space, action, and change. Such relational abstractions serve as a bridge between high-level domain-specific conceptual or activity theoretic knowledge, and low-level statistically acquired features and sensory grammar learning techniques. Generation of perceptual narratives, and their access via the declarative interface of logic programming facilitates the integration of the overall framework in bigger projects concerned with cognitive vision, robotics, hybrid-intelligent systems etc. In the smart meeting cinematography domain the generated narratives are used to explain and understand the observations in the environment and anticipate interactions in it to allow for intelligent coordination and control of the involved PTZ-cameras.

3   Immediate Outlook

The smart meeting cinematography scenario presented in this paper serves as a challenging benchmark to investigate narrative based high-level cognitive interpretation of everyday interactions. Work is in progress to release certain aspects (pertaining to space, motion, real-time high-level control) emanating from the narrative model via the interface of constraint logic programming (e.g., as a Prolog based library of depth–space–motion). We also plan to release general tools to perform management and visualisation of activity interpretation data.

References

1   Roland Barthes and Lionel Duisit. An Introduction to the Structural Analysis of Narrative. *New Literary History*, 6(2):237–272, 1975.

2   Mehul Bhatt. Reasoning about space, actions and change: A paradigm for applications of spatial reasoning. In S. Hazarika, editor, *Qualitative Spatial Representation and Reasoning: Trends and Future Directions*, pages 284–320. IGI Global, USA, 2012.

3   Mehul Bhatt and Gregory Flanagan. Spatio-Temporal Abduction for Scenario and Narrative Completion: a preliminary statement). In Mehul Bhatt, Hans Guesgen, and Shyamanta Hazarika, editors, *ECAI 2010, Workshop Proceedings, Spatio-Temporal Dynamics*, pages 31–36, 2010.

4   Mehul Bhatt, Jae Hee Lee, and Carl Schultz. CLP(QS): A Declarative Spatial Reasoning Framework. In Max J. Egenhofer, Nicholas A. Giudice, Reinhard Moratz, and Michael F. Worboys, editors, *Spatial Information Theory, 10th International Conference, COSIT 2011,*

Belfast, ME, USA, September 12-16, 2011. Proceedings, number 6899 in Lecture Notes in Computer Science, pages 210–230. Springer, 2011.

5 Mehul Bhatt, Jakob Suchan, and Christian Freksa. ROTUNDE – A Smart Meeting Cinematography Initiative. In M. Bhatt, H. Guesgen, and D. Cook, editors, *Proceedings of the AAAI-2013 Workshop on Space, Time, and Ambient Intelligence (STAMI).*, Washington, US, 2013. AAAI Press. (to appear).

6 Anthony G. Cohn and Jochen Renz. Qualitative spatial representation and reasoning. In Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors, *Handbook of Knowledge Representation*, pages 551–596. Elsevier, 2007.

7 Krishna Dubba, Mehul Bhatt, Frank Dylla, Anthony Cohn, and David Hogg. Interleaved Inductive-Abductive Reasoning for Learning Event-Based Activity Models. In Stephen Muggleton, Alireza Tamaddoni-Nezhad, and Francesca A. Lisi, editors, *Inductive Logic Programming, 21st International Conference, ILP 2011, Windsor Great Park, UK, July 31 – August 3, 2011, Revised Selected Papers*, number 7207 in Lecture Notes in Computer Science, pages 113–129, 2011.

8 Manfred Eppe and Mehul Bhatt. Narrative based postdictive reasoning for cognitive robotics. In *Proceedings of the 11th International Symposium on Logical Formalizations of Commonsense Reasoning*, 2013.

9 Antony Galton. *Qualitative Spatial Change.* Oxford University Press, Oxford, 2000.

10 Hannaneh Hajishirzi, Julia Hockenmaier, Erik T. Mueller, and Eyal Amir. Reasoning about robocup soccer narratives. *CoRR*, abs/1202.3728, 2012.

11 Hannaneh Hajishirzi and Erik T. Mueller. Symbolic probabilistic reasoning for narratives. In Ernest Davis, Patrick Doherty, and Esra Erdem, editors, *Logical Formalizations of Commonsense Reasoning, Papers from the 2011 AAAI Spring Symposium, Stanford, California, USA, March 21-23, 2011*, number SS-11-06 in AAAI Technical Reports, pages 123–126, 2011.

12 George Lakoff and Srini Narayanan. Toward a computational model of narrative. In Mark Alan Finlayson, editor, *Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium*, number FS-10-04 in AAAI Technical Reports, pages 21–28. AAAI Press, 2010.

13 Inderjeet Mani. *Computational Modeling of Narrative.* Number 5 in Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2012.

14 Inderjeet Mani. Computational Narratology. In Peter Hühn, Jan Christoph Meister, John Pier, and Wolf Schmid, editors, *The Living Handbook of Narratology.* Hamburg University Press, 2013.

15 John McCarthy. Concept of logical AI. In Jack Minker, editor, *Logic-based artificial intelligence*, pages 37–56, Norwell, MA, USA, 2000. Kluwer Academic Publishers.

16 Rob Miller and Murray Shanahan. Narratives in the Situation Calculus. *Journal of Logic and Computation*, 4(5):513–530, 1994.

17 Erik T. Mueller. Modelling space and time in narratives about restaurants. *Literary and Linguistic Computing*, 22(1):67–84, 2007.

18 Javier Pinto. Occurrences and narratives as constraints in the branching structure of the Situation Calculus. *Journal of Logic and Computation*, 8(6):777–808, 1998.

Exploring the Betrothed Lovers*

Andrea Bolioli[1], Matteo Casu[2], Maurizio Lana[3], and Renato Roda[4]

1   **Cross Library Services srl**
     Trento, Italy
     `bolioli@cross-library.com`
2   **CELI srl**
     Torino, Italy
     `casu@celi.it`
3   **Dipartimento di Studi Umanistici**
     Università degli Studi del Piemonte Orientale, Vercelli, Italy
     `m.lana@lett.unipmn.it`
4   **Dipartimento di Culture, Politiche, Società**
     Università degli Studi di Torino, Torino, Italy
     `renato.roda@unito.it`

Abstract

We present the ongoing activities and the first results achieved in a research project concerning the understanding of narrative in the high school. Students and teachers experimented with new ways to learn linguistic and digital skills, by using a collaborative learning environment built around the novel *I Promessi Sposi*. We analyzed the literary text, extracting social networks of characters and other fundamental narrative elements (sequences, locations, etc.), in order to provide the students with appropriate tools and resources to conduct their own inquiries on the novel.

1   Introduction

In this paper we describe the ongoing activities and the first results achieved in a research project concerning the understanding of narrative in the high school. The project partners are four Italian high schools (in the province of Trento), the Human Language Technology research unit at FBK-irst, two small enterprises, and two public educational organizations (IPRASE and Centro Formazione Insegnanti). The research project *Sharing Educational Content* (Sèduco) aims to study, develop and test digital tools for the creation, management and sharing of educational resources in the Italian upper secondary school (age 14 to 18).

Two ongoing activities in the project are concerned with the topic of narrative understanding:

- the linguistic and narratological analysis, with computational tools, of the historical novel *I promessi sposi* (*The Betrothed*) by Alessandro Manzoni,
- the collaborative creation of an anthology of Italian literature, which is named Antologia 2.0.

This paper focuses on the first issue. Students and teachers experimented with new ways to learn linguistic and digital skills, by using a collaborative learning environment[1] built around the novel *I Promessi Sposi*. This educational application can be considered an example of "semantic mashup" for narrative, a discovering method which is gaining attention and it is considered attractive by general users (see for example `http://lotrproject.com`) as well as a good training camp by researchers. In agreement with Inderjeet Mani, we believe "that computational narratology has the potential to revolutionize the way we create and study literature" [6].

The remainder of the paper is organized as follows. Section 2 briefly describes the *Promessi Sposi* learning environment. Section 3 presents the basic ontology underlying the system. Section 4 describes the social networks of narrative interactions created and used in the environment.

2   The learning environment as a semantic mashup

I promessi sposi is a very important Italian novel written by Alessandro Manzoni and published in its final illustrated version in 1840. For many and various reasons it was recognized as a work representing the cultural and linguistic unity of Italy and rapidly became a book read and studied in the secondary high school. Today it constitutes a difficult matter of study for students because of the growing distance from its themes, matter, and linguistic flavor (nineteenth century vocabulary and syntax).

The work can be – in a typical postmodern way – rebuilt while keeping it intact. This has been done by making it the center of a universe of relation, tools and activities.

Relations: around the work relations are built first of all among schools whose students work on the text, among schools and SMEs, and among schools and territorial institutions.

Tools: we built a set of tools allowing the students to browse the text following characters in theirs wanderings, to pay attention to a specific place in order to look at passing by characters, to explore the text by means of social network graphs, and so on.

Activities: in addition to "exploration" activities, students can link self-made reports, photographs, images (such as the drawings by Francesco Gonin), other texts, bibliographic references to the appropriate passages.

The basis for this rebuilding of the work has been the annotation of the novel for entity mentions and narrative sequences. We annotated a version of Promessi sposi [7] for Entity mentions of type Person (characters) and Location using two NER tools (TextPro by FBK-HLT and Sophia Semantic Engine by CELI). The automatic annotation were manually corrected. The novel has been split into sequences, that is the basic narrative units studied in high schools, (chunks of text) concerning a specific action or place or character. Dialogs, being the object of a specific interest, constitute always distinct sequences. Given this criteria, this operation has been done by humans.[2]

What actually happens is that students can browse the novel following characters or searching for places, can find the sequences pertaining to their interest, can read them, can publish content related to a specific sequence or to a block of sequences linking it to them, can study graphs showing relations between characters based on their co-presence in the same place or on their being part of a same dialogue.

---

[1] `http://seduco.cross-library.com/promessisposi/`
[2] A portion of the resulting dataset will be available to the research community.

Even technical institute students, with a few literary interests, have enjoyed the novel and they succeeded in performing linguistic and narrative analysis that would not have been possible in a traditional frontal lesson based solely on the paper book.

3   Ontological modelling of *I Promessi Sposi*

An interactive application with such a rich set of contents has faced the task of finding a good balance between exploiting all the information (acquired via automated or manual annotation) and avoid its explosion in terms of conceptual complexity.

Our treatment of *I Promessi Sposi* contains different sets of entities, which are linked by means of relations: the novel itself, discourse elements (Chapters, Paragraphs, Textual chunks), narrative elements (Sequences, Agents and Locations).

As for the novel, we considered the FRBR (Functional Requirements for Bibliographic Records) model [5], which distinguishes between the abstract Work from its Expressions (e.g., different revisions by the author) and its Manifestations.

Discourse elements, such as chapters and paragraphs, are defined by the author, and – given a certain edition – they are fixed. Ontologically, in order to treat this subset of entities, we reused classes from DoCO, from the SPAR (Semantic Publishing And Referencing) set of ontologies [9].

Narrative elements are the most interesting in this learning environment. We defined our own subclasses for a sequence: signally, Narrative (including Dialogues and Monologues), Descriptive and Reflexive kinds of sequences. Sequences are in many-to-many relation with paragraphs. A narrative sequences is linked to a location (the place of the narrative unit), to characters which act in it, and possibly to speakers, which are not necessarily the same.

Characters are, for the most part, persons or group of persons. A relevant exception is played by the *Divine Providence*, which can be considered as an agent in Manzoni's work. For these reasons we adopt the Friend Of A Friend (FOAF) vocabulary [3], which provides the class `foaf:Agent` with subclasses `foaf:Person` and `foaf:Group`, with the property `foaf:member` to link a group (such as the *bravi*, in our case) to its members.

The notion of character is not a class in our model, because we interpret it as a role played by an agent in a work. For this reason we modelled *character* as an OWL object property between a `foaf:Agent` and a `frbr:Work`.

Places and locations play a particular role in the model, for they can be regions, municipalities (such as Milan, Lecco) as well as relative locations (e.g., Lucia's house). Another geographic information in the novel is given by journeys between locations. In order to structure places – e.g., declaring that Lucia's house is in Lecco – we reuse the `parentFeature` property from the GeoNames ontology [10]. Another aspect about the places deals with the changing of places in time: some historical places do not exist anymore, have changed name or have changed status (e.g., the home of Lucia was in Acquate, formerly a municipality, now a neighborhood in Lecco). Again, we use the GeoNames ontology when possible and custom properties when necessary. In the context of the application these aspects are related to two maps depicting the places of the novel: namely, a XVIIth century map of northern Italy, and the current Google map.

4   SNA in *I Promessi Sposi*

In narratology, SNA (Social Network Analysis) has been used mostly as a new instrument for the study of plot evolution. By the extraction of the interactional networks of characters

from narrative works and the subsequent synthesis of the obtained data in network graphs it is possible to open a whole new perspectives to better comprehend the dynamics and the structure of a narrative plot. Even excluding the numerous quantitative analysis options available, the mere rearrangement of the narration from the written context to a fairly clear and understandable display – where the whole plot is summarized in a single eidetic structure – represents a powerful explanatory enhancement [8]. This SNA analytical approach has already been employed with a fairly large selection of different literary text, ranging from Shakespeare's tragedies [8] to Lewis Carroll's *Alice in Wonderland*[1], including the whole Marvel comics universe [2].

However, the vast heuristic potential of quantitative network analysis methods has rarely been employed for studying a text as structured and as complex as *I Promessi Sposi* is. In addition, it is the first time it is used in an educational context (high schools), as far as we know. The composite narration of Alessandro Manzoni's most famous work comprises plentiful flashbacks, historical and ethical excursuses, changes in the main narration perspective.

Such narrative complexity challenges the heuristic potential of SNA techniques: the available methodologies must be carefully selected and polished to obtain the most effective analytical tool. While the identification of the network's "nodes" with the characters of the studied texts is a choice common to the absolute majorities of literary researches that employed network analysis methodologies, the definition of "edges" appears less immediate and more difficult. One option, in such sense, focuses on "conversational edges" [4], where an edge is formed between two characters/nodes every time the studied text features an explicit dialogue between aforementioned characters.

The graph in Figure 1 represents the complete conversational network of *I Promessi Sposi*, elaborated with the SNA visualization software Gephi. The dimension of each node is due to the number of conversational interactions ("degree" in SNA terminology) in which that node's character is involved: in fact the largest node, in the very middle of the graph, embodies Renzo Tramaglino, the me lead of the novel. The thickness of each "edge" is directly correlated to the number of interactions between the relative couple of nodes: if a dyad of characters shares many dialogues, the edge that tie their node will be more thick. Different colors distinguish each community of characters (that is a set of highly interconnected nodes).

Remarkably, the community optimization algorithm perfectly recognizes the various "narrative blocks" of characters interactions: for example, the group of green nodes at the left of Renzo, quite isolated from the rest of the network, identifies one of the novel's most famous flashbacks, centered on the origin of Gertrude's character.

Alongside the "conversational network", we have also extracted a second network from *I Promessi Sposi*: while the nodes still represent the main characters of the novel, this time an edge between two characters identify their co-occurrence in the same narrative sequence. The graph in Figure 2 illustrates this second action-oriented network. As it is possible to notice even at a first glance, while the general structure of the two networks is similar, there are few interesting differences: for example, in the second graph the node Lucia Mondella, the main female character, has a slightly larger degree than the one of Renzo, and achieve the status of larger hub of the graph. This is particularly interesting since the main plot of the book splits very soon into 2 separate narrative lines, one for each of the titular betrothed lover. According to our graphs, role of Lucia is more action-oriented while the one of Renzo more involved in conversational interactions.

The twofold SNA approach used, however, revealed not to be only a valuable visualization/summarizing education tool, but also a powerful instrument per a deeper exploration of the text's narrative structure and characters dynamics. With results that sometimes go

beyond the mere teaching/learning scope that is the main focus of our project. In fact, thanks to the quantitative metrics obtained from the two characters network, it was even possible to open new perspective of narrative analysis and to discover aspects unknown (and unexpected) even to our team's experts in the work of Alessandro Manzoni. For example, while investigating on the relational "importance" of each character in the novel through the evaluation of betweenness centrality (a measure of brokerage centrality, that is the value of a node as bridge between other nodes in a given network), we were surprised to find out that some characters considered moderately secondary have an higher value than some of the main actors of the plot. Figures like "il Griso" – the main antagonist's sidekick – and even more "Agnese" – Lucia's mother – play a way stronger bridging role in the relational balance than stronger narrative characters like the charismatic villain (then reformed helper) "Innominato" or the iconic "Don Abbondio". This is just one example of the many analytic exploration perspectives opened by the SNA approach, perspectives than can be investigated by both students and the researchers, taking advantage of the nature of versatile research tool – and not only passive visualization solution – of the network graph.

──── **References** ────

1 Apoorv Agarwal, Augusto Corvalan, Jacob Jensen, and Owen Rambow. Social network analysis of Alice in Wonderland. In *Proceedings of the NAACL-HLT 2012 Workshop on Computational Linguistics for Literature*, pages 88–96. Association for Computational Linguistics, 2012.

2 Ricardo Alberich, Joe Miro-Julia, and Francesc Rosselló. Marvel universe looks almost like a real social network. Preprint, 2002.

3 Dan Brickley and Libby Miller. The Friend Of A Friend (FOAF) vocabulary specification, 2007. `http://xmlns.com/foaf/spec/`.

4 David K. Elson, Nichoas Dames, and Kathleen McKeown. Extracting social networks from literary fiction. In Jan Hajič, Sandra Carberry, Stephen Clark, and Joakim Nivre, editors, *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 138–147, Stroudsburg, PA, 2010. Association for Computational Linguistics.

5 IFLA Study Group on the Functional Requirements for Bibliographic Records. Functional requirements for bibliographic records: final report, 1997. As amended and corrected through February 2009.

6 Inderjeet Mani. *Computational Modeling of Narrative.* Number 5 in Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2012.

7 Angelo Marchese, editor. *Alessandro Manzoni: I Promessi Sposi.* Arnoldo Mondadori Editore, 1985.

8 Franco Moretti. Network theory, plot analysis. *New Left Review*, 68:80–102, 2011.

9 David Shotton and Silvio Peroni. DoCO, the document components ontology, 2011. `http://purl.org/spar/doco`.

10 Bernard Vatant. GeoNames Ontology – Semantic Web. `http://www.geonames.org/ontology/`.

A   Character Networks
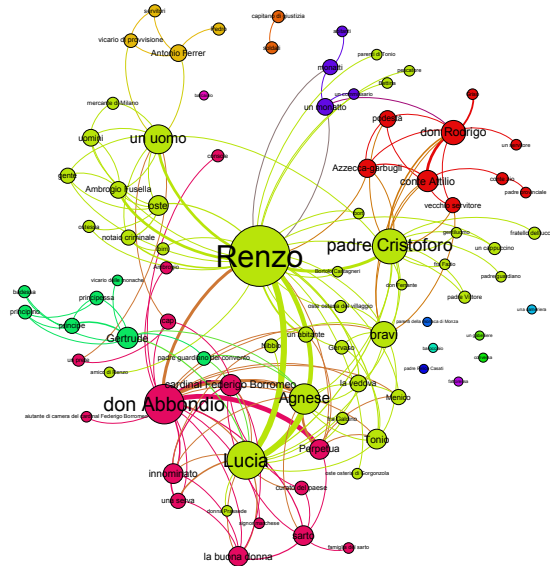
Figure 1 Conversational network.
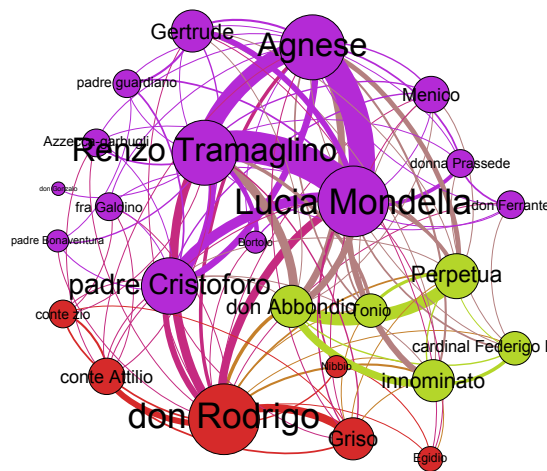
Figure 2 Co-occurrence in narrative sequences network.

The Disappearance of Moral Choice in Serially Reproduced Narratives*

Fritz Breithaupt[1], Kevin M. Gardner[2], John K. Kruschke[3], Torrin M. Liddell[3], and Samuel Zorowitz[4]

1 Department of Germanic Studies
   Indiana University
   Bloomington, IN, USA
   `fbreitha@indiana.edu`
2 Hutton Honors College
   Indiana University
   Bloomington, IN, USA
   `kmgardne@indiana.edu`
3 Department of Psychological and Brain Sciences
   Indiana University
   Bloomington, IN, USA
   `kruschke@indiana.edu,tliddell@indiana.edu`
4 Department of Psychological and Brain Sciences
   Johns Hopkins University
   Baltimore, MD, USA
   `szorowi1@gmail.com`

── **Abstract** ──────────────────────────────

How do narratives influence moral decision-making? Our ongoing studies use serial reproduction of narratives, that is multiple retellings as in the telephone game, of morally ambiguous situations. In particular, we tested stories that include a minor misdemeanor, but leave open whether the wrongdoer will be punished by a bystander. It turns out that serial reproduction (retelling) of stories tends to eliminate the possibility of intervention by the bystander under certain conditions. We reason that this effect can be explained either by preferences of the readers or by the reader's discomfort to get involved. A second finding is that retellings of third-person narratives of moral situations lead to a higher degree of change and invention of the outcome than first-person narratives.

1 Introduction

The large question informing our research is how narratives shape the way people think about moral issues. One aspect of this question concerns moral choice in ambiguous situations. We discuss the conditions under which narratives support moral ambiguity or moral clarity.

The approach we use is serial reproduction of narratives [1] as in the telephone game or, in German, "Stille Post".

Most types of narrative deal with morally relevant situations, ranging from stories involving good or bad characters to ethical dilemmas. This is true for everyday reports of events, gossip (tracking who did what and why), and also for narrative fiction, which rarely does not deal with morally charged events. In general, we assume that narratives are not neutral with regard to morality and that moral issues are at the core of story-telling [5]. One of the questions following from our assumption is whether narratives support moral ambiguity or tend toward polarized solution of moral dilemmas.

One could speculate, for example, that moral reasoning often is the product of narrative apprehension, that is induced by particular narrative structures and not necessarily by moral content per se. In this case, one could suggest that it makes for better stories if there are many sides to an event which elicit and engage moral cognition. This is in line with narratologists' theories of the event that stress that each narrative event induces at least two different interpretations or competing contexts [2]. Such theories reach back to Goethe's articulation of the "unheard of event" [4]. Complex moral issues without a clear solution could be seen as a form of the multifaceted nature of events. If so, effective story-tellers might overemphasize moral ambiguity for narrative suspense (to engage attention, memory, and intensity). An opposite opinion could hold that narratives provide a system to deal with and resolve moral dilemmas. This could mean that narratives display morally charged events in such a way to expose all relevant features and thereby lead toward dissipating moral tensions. In the first case, we would expect to see narratives holding on to or overemphasizing moral problems and ambiguity. In the latter case, we would expect narratives leading toward a solution and thereby resolving moral conflicts.

To test this question, we have employed serial reproductions of narratives in a set of experiments, each containing a morally charged situation. In serial reproduction, a participant reads the narrative and then attempts to reproduce it, in writing, from memory alone. The reproduced story is then given to another participant, who attempts to reproduce it for the next participant. Here we ask what happens if a story with a morally ambiguous event or character gets serially repeated. Is the moral ambiguity increased, maintained, or reduced?

Prior research dealing with serial reproduction of narratives suggests that a "minimally counter-intuitive" event in a story might be most optimal for stories and should be maintained more accurately over a series of reproductions (or periods of time in memory studies) than narratives with no counter-intuitive events or "maximally counter-intuitive" events [9]. The studies by Norenzayan et al. did not focus explicitly on moral choices and moral ambiguity. Instead they dealt with events in fairy-tales. Still, we wonder to which degree the "minimally counter-intuitive" events also present the most morally relevant situations since they involve a moral problem that begs for a plausible solution.

However, we speculate that the moral pressure to achieve a resolution can in specific situations outweigh the narrative demand of retaining complexity (or suspense). For this pressure-to-a-resolution, it is not relevant whether moral reasoning leads to a resolution or simple quick side-taking or rush decision-making. In fact, quick side-taking may be key since humans are exceptionally rapid side-takers or judgment-makers [10, 12]. Indeed, DeScioli and Kurzban [3] argued that rapid side-taking, serving to reduce the costs of conflict, is the key evolutionary force behind human morality. Hence, we reason that the tendency of preserving a "minimally counter-intuitive" and morally charged situation can be overshadowed by a different dynamic, namely quick judgment, side-taking, or preference.

We speculate that when stories get retold, the implicit side-taking, judgment, or preference by the re-teller manifest itself in various ways. This in turn will influence the next reader to reinforce the side-taking and lead to a polarization of future story generations. Hence, we predict that serial reproductions of morally ambiguous situations will lead to more clearly polarized and resolved narratives, and thereby also offer an exception to the findings of Norenzayan et al. 2006 [9].

At the same time, we also speculate that there are certain conditions in which the side-taking or preference will not be allowed to manifest itself strongly. One of these conditions could be first-person narratives in which readers are "locked into" an ambivalent event that is driven by other characters. Here the reader is bound to one perspective that remains subject to the ambivalent event. Examples will be provided below.

2   Study 1

In these studies, we asked participants to read a short story and then retell it "in your own words". The retold story was then retold again up to a sixth version. The general design of the study of serial reproduction of narratives follows Kashima 2000, Lyons & Kashima 2006, and Kashima 2010 [6, 7, 8]. Our stories involve a morally ambiguous situation in which a moral choice is to be made by either a character in the story or by an observer. However, this choice is not expressed explicitly. In this pilot study, the same core story is presented in three forms to different participants: one time as a third-person narration with two named characters (Susan and Jessica), and twice as a first-person narrative in which either "Susan" or "Jessica" are replaced by a first-person narrator throughout the short story. The story puts one character as the agent of a minor misdemeanor, the other into the position to potentially intervene. Here is the prototype:

> Jessica is working at the cash register of her university food court. She sees her friend Susan approaching the food line. Jessica knows that Susan does not have much money on her account and can only afford a very simple meal. And indeed, Susan opts for a soup and takes water to drink, just enough for her diminished account. However, close to the checkout Jessica watches as Susan conceals a small apple in her bag. Of course, this is illegal, but Jessica knows that the apple Susan had bought yesterday had a worm inside it. In fact, Jessica had teased her about the apple and had said: "Oh, you want to kiss the worm? Maybe it is a prince." After the teasing, Susan did not eat the apple. Jessica knows that the university overprices healthy food, while junk food like French fries are held artificially cheap. Jessica knows that Susan does not have rich parents and has to work through summer and during the terms to make ends meet. Now Susan comes to the cash register and does not reveal the apple. Jessica starts to ring up the food.

We assumed that most participants would clearly notice that the actual decision (intervention or no intervention) is not made at the end of the story. Pretesting confirmed that this narrative was morally ambiguous. The response to "Do you blame one of the individuals for failing to respect other people or follow social rules?" yielded a response of nearly the perfect middle of 3.88 on a 1–7 scale ($n = 74$). Another question lead to a similar response near the mathematical middle of 4. It asked, "If there was a main wrong-doer, how important is it that he or she gets punished or at least confronted?" Responses were indicated on a 1–7 scale, with 1 labeled "Not Important" and 7 labeled "Very Important". The mean rating was 3.67 ($n = 74$).

In the following, we analyze and report only on the first retelling condition (iteration 1) since the results are already informative after a single retelling. All first iteration versions ($n = 45$) included the basic situation at the food court, the intentional concealing of the apple or fruit, and at least one of the excusing circumstances (relative poverty, the "worm", or the overpricing). The original story ended where the potential intervener could act, but there was no indicator at that point whether or not she intended to do so. We coded the retelling for clear markers that would establish that indeed a decision was made, either to confront Susan or not to intervene. Examples of not intervening include: "Jessica lets it go because...", "Jessica is OK with this" or "but I do not report it". In the absence of these markers, we coded that there was no decision made (as in the original story). Typically, retellings of this kind ended in something like: "the clerk starts to ring stuff up."

The results divided the different story conditions in a surprisingly clear fashion. Both stories with a first-person narrator tended to retain the ambiguous end as it was given in the source story. More than 75% of the retellings (23 of 30 in ongoing data collection) were coded as having an ambiguous end and only 17% (6/30) lacked ambiguity either because the situation was passed (the ringing up has already happened with no action taken by the teller) or because of a clearly marked intent of non-intervention. However, in the third-person condition, participants leaned toward not making an intervention and thus making the decision to let the misdemeanor pass (above 50%, and only 40% maintaining the ambiguous end, with 10% uncertain coding). Here is a typical result of the third-person condition, illustrating a decision to let the misdemeanor pass:

> Jessica works as a cashier at the cafeteria. She saw her friend Susan coming to buy food. She knows that Susan is poor. Jessica sees Susan buying soup and water and then steal an apple by putting it in her pocket. Jessica remembers how Susan had an apple with worm in it and jokingly wanted to kiss the apple to see if it was a prince. Susan comes and Susan pays for the food. Jessica doesn't say anything.

Pretesting of the different versions according to a standard questionnaire with eight numeric questions revealed that participants also would more strongly lean toward non-intervention in the third-person condition. Among the questions with a clear leaning is: "If there was someone facing a moral choice [at the end of the story], how important is it that he or she takes action (e.g., intervenes or calls for help)?" On a scale from 1–7 (1 = not important to act, 7 = very important to act), participants in the third-person condition (Jessica and Susan) had a lower response than in the first-person as wrongdoer scenario ($n = 74$; 25 or 24 per version; participants in the pretesting were different from the 45 retellers of the study). In several other questions, however, the average was similar for all three conditions, for example, "did someone in this story display bad character?"

2.1   Discussion of Study 1

To be sure, the sample is quite small at the time of this writing. Testing is ongoing and will be reported at the CMN meeting. However, assuming that the results hold, we could reason that the first-person narratives "lock" participants into a position that makes them look toward the outcome (intervention or no intervention). This is especially true when the first-person is the one who steals the apple. Intuitively, the unresolved anticipation of possibly getting caught is quite salient when understood from the first person. From the other first-person perspective, the unresolved anticipation of having to decide whether to intervene is also quite salient. Hence, marking the end as open is an essential part of the story for either first-person perspective.

This changes in the third-person condition. Judgment of the situation by third parties is not bound by what actually happens at the end (outcome), but by what should happen. Hence, side-taking seems to be a logical possibility. In this sense, the invented statement of what Jessica at the cash register will do (namely, not intervene) seems to express the preference of the reteller [11]. Notice that this change of the narrative leads to a disappearance of the situation of moral choice, because, in the retelling, the choice is already made. This is remarkably unfaithful to the source story.

There is another puzzling aspect to this. The invented statement by Jessica always favors non-intervention. This might suggest that participants consistently lean towards non-punishment. However, this seems not be the case. Recall that the prestudies established that participants in all three story conditions deem the offender to be blameworthy near the 50%-line (on a 1–7 scale). Perhaps there is more to the disappearance of moral choice than mere preference? In the following, we will consider a second possible reason for the invented addition of non-intervention.

3   Study 2

In the pilot study to another experiment, we again asked participants to read a short story and then retell it "in your own words." Again, it is the story of a bystander who can intervene. And once again, the story features a misdemeanor, but now the bystander does not necessarily face the choice to intervene. At the time of this writing, we only have results for third person accounts.

> Original Version (1)
> Max is walking to a job interview. He is a senior in college, and now it is time for him to enter the working world. He goes over all interview questions in his head. He took the train into Chicago this morning, so he feels important. It is his first time wearing his new suit. On the way to the job interview, he comes through a park. Suddenly he happens to see two kids steal the umbrella of an old man who is sitting on a park bench. The two kids run away, while the old man starts to shout at them. It is obviously a little prank since the kids throw the umbrella into a bush. The kids call something back to the old man, and do not see that Max is right behind them.

At the end of this story, Max implicitly has the option to intervene (for example scold or punish the kids) or to ignore the prank, but unlike the source story in Study 1, he does not have to engage in any action with the other characters (such as ringing the food up) and does not bear responsibility (as a teller at the food court has). We predicted that the serial reproduction would quickly lead to a polarization of the choice in the one or the other direction. We measured polarization in three ways. First, we asked control groups about their moral preferences ("How should Max act: Should he intervene or not?"). Second, we tallied a number of indicators in retellings, including omission of information between two versions, changes of word choice ("kids" become "guys", the "umbrella" becomes a more valuable object, etc.), negatively or positively marked words and ideas, and perspective-taking. Third, we noted changes of the plot involving a punishment or acquittal of the wrongdoer within the frame of the iteration. In this example, we expected some story versions after 2 or 3 iterations to demonize the kids more and have Max intervene. And we expected some other versions after 2 or 3 iterations to make the action of the kids more harmless. Results showed that polarization did indeed occur, and detailed results will be reported at a later time.

However, after five or fewer reproductions something else occurred. Here is a typical example:

Iteration 5

Max has newly arrived in Chicago. He is walking by a park when he sees two young boys harassing the old man sitting on a bench. The boys take the old man's backpack and throw it into the bushes. Max keeps walking, but having seen this, he can't help but wonder about what kind of place he's entered.

In this retelling, the entire moral decision for Max has disappeared. Whereas Max in the original version faced a choice, iteration 5 has him as a mere distanced observer without the pressure or opportunity to engage with what he observes.

Our ongoing studies will establish the frequency of this disappearance of moral choice in serially reproduced narratives. At this point, the disappearance of the moral choice occurred in 10 of 26 cases and it occurred instead of radical polarization that also could have eliminated the pressure of choice.

There are also related ongoing experiments, which must be omitted here for lack of space. The full presentation will include these additional studies.

It should be noted that a slightly different story did not lead to the disappearance of moral choice with a high frequency. The only change to the source story was that the two kids kept the umbrella rather than throwing it in the bushes. In this condition, most retellings (12 of 14; iteration 1 only) maintained the possibility of intervention at the end, namely the fact that the kids come close to Max.

4   Interpretation

We originally predicted that the serially reproduced stories would quickly lose their ambiguity and show a more polarized description. The preliminary results partly confirm this prediction. However, another form of retelling was to drop the moral choice in its entirety, often because an invented choice was simply made within the retelling. Both Study 1, in third-person condition, and Study 2 led to a disappearance of moral choice. Still, Study 1 and Study 2 seem to suggest different tendencies. Retellers in Study 1 omitted the choice by presenting a story with an invented choice already made according to what seems to be the preferred outcome. Retellings in Study 2 suggest a discomfort to make a moral choice by displacing the observer from the situation. The interpretations are not contradictory. Perhaps some "discomfort" is a driver to make the moral choice in Study 2. And Study 1 also may express the preference of the participant since a non-intervention is desired.

A possible frame for the interpretation of these serially reproduced narratives is provided by DeScioli & Kurzban (2012) [3]. They hypothesized that the purpose of moral judgment is the coordination of bystanders of a conflict in such a way that bystanders will choose to support the same side. The purpose of bystander coordination is group coherence and the avoidance of conflict. From this perspective, it is plausible to predict that people tend to choose sides (esp. when observed, see Kurzban et al. 2007) and will aim to influence others to choose the same side. In the condition of serial reproduction, this could lead to the described radicalization ("black-and-white painting").

However, the results only partially fit this idea. Rather than painting clear black-and-white pictures, or by involving a punishment or acquittal of the wrongdoer, the serially reproduced versions tended to justify the non-intervention by the bystander in these cases of minor misdemeanors. Max, in the quoted example of iteration 5, does not have to position himself vis-à-vis some crime or harmless prank by the kids any longer. The tension of a moral choice has disappeared. In the food court story, the woman at the food court already makes the decision, especially in the third-person condition.

Whereas these findings do not contradict DeScioli & Kurzban (2012), they suggest an additional urge to bypass the moral choice in its entirety. This is itself a way to avoid conflict, which is the very motive for side taking presumed by DeScioli & Kurzban. To be sure, morality does not disappear in its entirety – Max still observes the events, the woman at the cash register makes a decision – but the pressure to make a choice or get involved disappears from the story. Hence, we speculate that narratives seek an optimal position or perspective on moral choice. This position is one of being an observer, but not a decision-maker. Morality, within narratives, is a spectator sport.

References

1   Frederic C. Bartlett. *Remembering: An Experimental and Social Study.* Cambridge University, Cambridge, 1932.

2   Fritz Breithaupt. *Kultur der Ausrede.* Number 2001 in suhrkamp taschenbuch wissenschaft. Suhrkamp, Berlin, 2012.

3   Peter DeScioli and Robert Kurzban. A solution to the mysteries of morality. *Psychological Bulletin*, 139(2):477–496, 2013.

4   Johann Peter Eckermann. *Conversations of Goethe with Johann Peter Eckermann, 1836.* Da Capo Press, Cambridge, 1998. transl. J. Oxenford.

5   Jonathan Gottschall. *The Storytelling Animal: How Stories Make Us Human.* Houghton Mifflin Harcourt, 2012.

6   Yoshihisa Kashima. Maintaining cultural stereotypes in the serial reproduction of narratives. *Personality and Social Psychology Bulletin*, 26(5):594–604, 2000.

7   Yoshihisa Kashima and Victoria Wai-Lan Yeung. Serial reproduction: An experimental simulation of cultural dynamics. *Acta Psychologica Sinica*, 42(1):56–71, 2010.

8   Anthony Lyons and Yoshihisa Kashima. Maintaining stereotypes in communication: Investigating memory biases and coherence-seeking in storytelling. *Asian Journal of Social Psychology*, 9(1):59–71, 2000.

9   Ara Norenzayan, Scott Atran, Jason Faulkner, and Mark Schaller. Memory and mystery: The cultural selection of minimally counterintuitive narratives. *Cognitive Science*, 30(3):531–553, 2006.

10   Stephen Porter, Leanne ten Brinke, and Chantal Gustaw. Dangerous decisions: The impact of first impressions of trustworthiness on the evaluation of legal evidence and defendant culpability. *Psychology, Crime & Law*, 16(6):477–491, 2010.

11   David N. Rapp and Richard J. Gerrig. Predilections for narrative outcomes: The impact of story contexts and reader preferences. *Journal of Memory and Language*, 54(1):54–67, 2006.

12   Alexander Todorov, Manish Pakrashi, and Nikolaas N. Oosterhof. Evaluating faces on trustworthiness after minimal time exposure. *Social Cognition*, 27(6):813–833, 2009.

Gist and Verbatim in Narrative Memory*

David A. Broniatowski[1] and Valerie F. Reyna[2,3]

1   Center for Advanced Modeling in the Social, Behavioral, and Health Sciences
    Johns Hopkins University
    Baltimore, MD, USA
    broniatowski@jhu.edu
2   Center for Behavioral Economics and Decision Research
    Cornell University
    Ithaca, NY, USA
    vr53@cornell.edu
3   Cornell Magnetic Resonance Imaging Facility
    Cornell University
    Ithaca, NY, USA

Abstract

A major concern regarding the study of narratives regards how they are indexed and retrieved. This is a question which touches on the structure of human memory in general. Indeed, if narratives capture the substance of human thought, then data that we have already collected regarding human memory is of central importance to the computational study of narrative. Fuzzy Trace Theory assumes that memory for narrative is simultaneously stored at multiple levels of abstraction and, whenever possible, decision-makers interpret a stimulus qualitatively and therefore operate on a simple – typically categorical – "gist" representation. Here, we present a computational model of Fuzzy Trace Theory applied to explain the impact of changes in a narrative upon risky-choice framing effects. Overall, our theory predicts the outcome of 20 experimental effects using only three basic assumptions: 1) preference for lowest level of gist, that is, categorical processing; 2) decision options that fall within the same categorical description are then interpreted using finer-grained (ordinal or verbatim) distinctions; and 3) once the options are mentally represented, decision preferences are generated on the basis of simple positive vs. negative valences stored in long-term memory (e.g., positive value for human lives). A fourth assumption – that negatively-valenced decision options are preferentially converted to positive decision options – is used when categories are not otherwise comparable.

1   How Are Narratives Stored in Memory?

Narratives entail sequences of events that must be both recalled from the narrator's memory and stored in the memory of the audience. It has long been known that narratives are not only recalled, nor stored, in a verbatim manner – rather, memory for narratives tends to favor events that are central to the story's causal structure [17, 18] – i.e., events that communicate the narrative's meaning to its audience. For example, Dehghani et al. [4] have

---

shown that changes in a narrative's surface form did not change how subjects responded to a questionnaire about how the narrative's main character was expected to respond, whereas changes in the narrative's underlying structure did lead to a change in the expected behavior of the main character. Dehghani et al. argue that changes in the narrative's surface form lead subjects to make inferences based on analogy with the original narrative, whereas changes in the underlying structure of the narrative preclude such an analogy from forming. Such observations motivate our research question: How are narratives stored in, and recalled from, memory?

2   Is Memory Schematic?

Modern psychological theories of memory storage and retrieval may be broadly classified into two categories: schema theories and association theories, both of which are widely found throughout the artificial intelligence (AI) and psychology literatures [7, 15]. Schema theories posit the existence of a schema or frame that is used to structure memory and experience such that what is recalled is coherent and confirms existing expectations. Such schemata are typically generated through repeated experience and may be part and parcel of membership within a given culture, domain of expertise, or any other situation in which frequent exposure to a set of regularities drives future expectations in a structured way. Schema, once formed, can drive perception and explanation, and may therefore be said to have an existence that is independent of the world features with which they interact. In contrast, association theories, assume that meaning emerges from co-occurrence patterns among world feature. Whereas schema models posit a higher-level data structure which guides surface feature recognition and interpretation, association models posit a bottom-up approach whereby surface feature co-occurrence drives meaning. In their classic paper, Alba and Hasher [1] synthesized evidence demonstrating that memory, although possessing several features of schematic representation, nevertheless also was strongly affected by surface form. In fact, that literature review documented robust but contradictory findings from these alternative theoretical perspectives. Empirical research suggests that elements of the schema and associationist models are both correct [13]. For example, phenomena such as false recognition, where "false" memories that are consistent with a given schema are recalled, support schema theory. On the other hand, phenomena such as false recognition reversal (higher levels of rejection for schema-consistent recognition items), have also been observed, supporting the recall of surface form. These contradictory findings have led to the development of "Fuzzy Trace Theory" (FTT) – a theory of memory for narrative that posits the simultaneous existence, and encoding, of two types of memory. Verbatim memory is memory for surface form and is typically detailed, yet brittle (i.e., quick to fade). Verbatim memory might include the specific words used in the telling of a narrative. On the other hand, FTT also posits the existence of gist memory, which captures the essential (often schema-consistent) meaning of a stimulus.

3   The Asian Disease Problem as a Testbed for Narrative Memory

The difference between gist and verbatim may be explained using a classic decision problem due to Tversky and Kahneman, known as the Asian Disease Problem (ADP) [19, 20], which is accompanied by the following pair of narratives: Assume that 600 people are expected to die from an outbreak of disease. You have a choice between two programs to combat the disease: (a) 200 people will be saved versus (b) a 1/3 probability that 600 people will be

saved and a 2/3 probability that no one will be saved. The "people saved" version of the problem is described in terms of gains relative to a reference point of 0. The "people die" version of the same problem is: Assume that 600 people are expected to die from an outbreak of disease. You have a choice between two programs to combat the disease: (a) 400 people will die versus (b) a 2/3 probability that 600 people will die and a 1/3 probability that no one will die. This problem has been used to highlight a framing effect, which refers to the typical result that most people prefer the certain option in the gain frame, but they prefer the risky option (gamble) in the loss frame. FTT holds that the narrative described above is simultaneously encoded on two levels: gist and verbatim. The verbatim level behaves in a manner similar to traditional economic theory, whereby a comparison is made between two options of equal expected value, as follows:

(a) 200 people will be saved = 200 saved vs.
(b) 1/3 * (600) people will be saved = 200 saved

In contrast, the gist level represents the core meaning of each of the decision options. Thus, a representation of the above problem at the gist level might be:

(a) Some are saved vs.
(b) Maybe some are saved or maybe none are saved

This would lead a decision-maker to choose option a, consistent with empirical findings. Similarly, the gist of the negatively-framed alternative might be:

(d) Some die vs.
(e) Maybe some die or maybe none die

This framing would lead a decision-maker to choose option b, also consistent with empirical findings. Importantly, gist and verbatim interpretations are recorded in parallel – i.e., gist is not derived from a verbatim representation [13]. Furthermore, several levels of gist can exist, such as at the level of the word, the sentence, the paragraph, or the entire narrative. We are able to use the concepts of gist and verbatim to replicate the findings of 20 separate experiments from the psychology literature. These effects are the consequences of the following four assumptions:

1. preference for categorical representations that are simultaneously the most meaningful and least detailed
2. decision options that fall within the same categorical description are then interpreted using finer-grained (ordinal or verbatim) distinctions
3. once the options are mentally represented, decision preferences are generated on the basis of simple positive vs. negative valences stored in long-term memory (e.g., positive value for human lives).
4. negatively-valenced decision options are preferentially converted to positive decision options –when categories are not otherwise comparable. Following is a computational implementation of a risky decision-making problem of the sort outlined above, meant to further clarify the distinction between gist and verbatim.

Our computational model has the following elements:

1. Configuration space: A narrative must occur within a given context – what is known in literary theory as a small world [5]. In the case of our decision-making problem outlined above, we posit the existence of a configuration space [6], which serves as the contextual

```
create_configuration_space(num_dimensions, special_points)
for i=1:num_decision_alternatives
        category(i) = find_category(decision_alternative(i))
end

for i=1:total_decision_alternative_pairs
        representation = categorical_representation(i)
        if~(categories_comparable)
                representation = convert_negative_to_positive(i)
        end

        while(representation~=verbatim_representation(i))
                if(dominant_category_exists)
                        return dominant_option
                else
                        representation =
                add_precision(representation)
                end
        end
end
```

■ **Listing 1** FTT Pseudocode.

grounding for our decision. In our context, a configuration space is a mathematical formalization of the universe of possible decisions that might be made. For example, for the ADP, there are two types of numbers that a decision-maker is required to understand. The first represents the number of people who are saved, and the second represents the probability with which the first number occurs. This may be represented as a 2-dimensional space, where one axis (ranging from 0 to 600) represents the number of people who might be saved under the different treatment conditions and the other axis (ranging from 0 to 1) represents the probability of success in the gamble option. Each decision option complement represents a point in this space. For example, the certain option is located at (200, 1) because there is a certain probability that 200 people will be saved. The non-zero complement of the gamble option is located at (600, 1/3) since there is a 1/3 probability that 600 people will be saved. Finally, the zero complement of the gamble option is located at (0, 2/3) since there is a 2/3 probability that 0 people will be saved under the gamble option.

2. Special points: We may subject points in the configuration space to certain constraints. Restricting points to a subset of this space reflects one such constraint. A concrete example of such a restriction might be that the total number of people saved equals zero which restricts all points to be on the vertical axis (i.e., the line $x = 0$). This corresponds to the statement "none are saved". Another example might be the horizontal line defined by the equation $y = 1$, corresponding to a certain outcome. Further, Feldman observes that constraints of this type are highly improbable by chance alone – i.e., a point chosen at random from the configuration space is highly unlikely to satisfy the constraint by accident. These values are associated with qualitatively different categories: an event happens or not, is certain versus uncertain, or is impossible versus possible. Thus, none saved ($x = 0$) is likely to be salient. All ($x = 600$) saved may be salient if it is made clear in the problem definition that 600 represents the entire population (therefore all are saved;

Figure 1 An example of a 2-dimensional configuration space for the Asian Disease Problem. Each point represents a decision complement.

Figure 2 A 2-dimensional configuration space with 1-dimensional restrictions defined by the lines $x = 0$ and $y = 0$ (heavy dashed lines). The likelihood that a randomly chosen point in this 2-dimensional space will fall on either of these lines approaches 0.

see below). As noted above, the constraints that are salient in the configuration space are determined by biological, psychological, or socio-cultural factors. Given a configuration space, certain points within this space are psychologically special (or "non-accidental" see, e.g., [8]), and are therefore categorically different than other points in the space. In the ADP, "none saved" or "none die" are special points because they are categorically different from "some saved", or "some die". In this way, these special points define a set of categories in the configuration space. Although, formally, "0 die" is still part of the closed interval containing the "die" axis, and is therefore contained within the "some die category", it is preferentially interpreted by the most restrictive category available, consistent with Feldman's genericity principle. These categories therefore define a partial

order over categories that may be thought of as a perception lattice [6]. For example, Kasturirangan [10] has proposed a narrative interpretation of non-accidental points whereby narrative contexts change whenever characters in a story encounter conflicts, which are non-accidental features in his framework.

3. Values: Given a pair of categories, one is either strictly preferred to the other, or they are unrelated. Such values allow a decision-maker to choose between these categories. For example, "some live" is preferred to "none live", whereas "none die" is preferred to "some die".

According to FTT, gist-level processing is preferred, explaining the framing effects observed by Tversky and Kahneman [19, 20]. This model has been implemented in GNU/OCTAVE and faithfully replicates 20 different experimental effects, including truncated and expanded versions of the ADP, and frequently-observed outcomes in the Allais paradox [2]. In each of these cases, minor changes in the narrative describing a decision-situation can lead to large changes in the decision outcome.

4   The 20 Effects Predicted

Our model has been tested against, and successfully predicts, the outcomes of the 20 experiments listed in the following table:

Effect	Stimulus	Gist Representation	Decision Outcome	Data Source
1	200 saved vs. 1/3 chance 600 saved OR 2/3 chance none saved	some saved vs. some chance some saved OR some chance none saved	A	[19]
2	400 die vs. 1/3 chance none die OR 2/3 chance 600 die	some die vs. some chance none die OR some chance some die	B	[19]
3	200 saved AND 400 not saved vs. 1/3 chance 600 saved OR 2/3 chance none saved	some saved AND some not saved vs. some chance some saved OR some chance none saved	Indifference	[11, 16, 3]
4	400 die AND 200 do not die vs. 1/3 chance none die OR 2/3 chance 600 die	some die AND some do not die vs. some chance none die OR some chance some die	Indifference	[11, 16, 3]
5	400 not saved vs. 1/3 chance 600 saved OR 2/3 chance none saved	some not saved vs. some chance some saved OR some chance none saved	B	[11]
6	200 do not die vs. 1/3 chance none die OR 2/3 chance 600 die	some do not die vs. some chance none die OR some chance some die	A	[11]
7	200 saved vs. 1/3 chance 600 saved	some saved vs. some chance some saved	Indifference	[14, 16, 12]
8	400 die vs. 2/3 chance 600 die	some die vs. some chance some die	Indifference	[14, 16, 12]

9	200 saved vs. 2/3 chance none saved	some saved vs. some chance none saved	A	[14, 12]
10	400 die vs. 1/3 chance none die	some die vs. some chance none die	B	[14, 12]
11	200 saved vs. 2/3 chance 600 die	some saved vs. some chance some die	A	[14]
12	400 die vs. 1/3 chance 600 saved	some die vs. some chance some saved	B	[14]
13	200 saved vs. 1/3 chance none die	some saved vs. some chance some saved	Indifference	[14]
14	400 die vs. 2/3 chance none saved	some saved vs. some chance none saved	A	[14]
15	200 saved vs. 1/3 chance none die OR 2/3 chance 600 die	some saved vs. some chance some saved OR some chance none saved	A	[14]
16	400 die vs. 1/3 chance 600 saved OR 2/3 chance none saved	some saved vs. some chance some saved OR some chance none saved	A	[14]
17	1m dollars with certainty 1m dollars with 0.89 chance OR Nothing with 0.01 chance OR 5m dollars with 0.10 chance	some money with certainty some money with some chance OR no money with some chance OR some money with some chance	A	[2]
18	Nothing with 0.89 chance OR 1m dollars with 0.11 chance Nothing with 0.90 chance OR 5m dollars with 0.10 chance	no money with some chance OR less money with some chance no money with some chance OR more money with maybe	B	[2]
19	200 saved vs. 1/3 chance all 600 saved OR 2/3 chance none saved	some saved vs. maybe all saved OR maybe none saved	Indifference	[22, 21, 9]
20	400 die vs. 1/3 chance none die OR 2/3 chance all 600 die	some die vs. maybe none die OR maybe all die	Indifference	[22, 21, 9]

Each of these effects is representative of a class of decision problems that share a similar structure.

5 Discussion

We began this paper by asking how narratives are stored in, and retrieved from, human memory. In order to answer this question computationally, we take steps towards formalizing Fuzzy Trace Theory – a theory of human memory. Although we focus on decision problems for tractability, the more general concepts of gist and verbatim representation are applicable

to the study of narrative. Furthermore, decision problems are particularly relevant to the computational modeling of narrative because they entail short narratives that set up the decision situation. Minor changes in the way these narratives are framed can have major impacts on the ultimate decision made. We argue that this is central to the computational modeling of narrative because understanding may depend significantly on how a narrative is framed. Future work will focus on extending our formalization beyond the domain of the simple decision problems to more complex accounts of false memory and narrative structure.

References

1   Joseph W. Alba and Lynn Hasher. Is memory schematic? *Psychological Bulletin*, 93(2):203, 1983.

2   Maurice Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, 21(4):503–546, 1953.

3   Tilmann Betsch and Maria Kraus. Die Auswirkungen von Ergebnis-Framing und dem Wechsel der Problemdomäne auf monetäre Entscheidungen. *Zeitschrift für Experimentelle Psychologie*, 46:296–304, 1999.

4   Morteza Dehghani, Sonya Sachdeva, Hamed Ekhtiari, Dedre Gentner, and Ken Forbus. The role of cultural narratives in moral decision making. In Niels Taatgen and Hedderik van Rijn, editors, *Proceedings of the 31th Annual Conference of the Cognitive Science Society*, pages 1912–1917, 2009.

5   Umberto Eco. *The limits of interpretation*. Indiana University Press, 1994.

6   Jacob Feldman. The structure of perceptual categories. *Journal of Mathematical Psychology*, 41(2):145–170, 1997.

7   David A. Gallo. *Associative illusions of memory: False memory research in DRM and related tasks*. Psychology Press, 2006.

8   Allan Jepson and Whitman Richards. What makes a good feature. In Laurence Harris and Michael Jenkin, editors, *Spatial Vision in Humans and Robots: The Proceedings of the 1991 York Conference on Spatial Vision in Humans and Robots*, pages 89–126. Cambridge University Press, 1994.

9   Jerwen Jou, James Shanteau, and Richard Jackson Harris. An information processing view of framing effects: The role of causal schemas in decision making. *Memory & Cognition*, 24(1):1–15, 1996.

10 Rajesh Kasturirangan. Thinking is believing. *Progress in Brain Research*, 168:105–114, 2007.

11 Anton Kühberger. The framing of decisions: A new look at old problems. *Organizational Behavior and Human Decision Processes*, 62(2):230–240, 1995.

12 Anton Kühberger and Carmen Tanner. Risky choice framing: Task versions and a comparison of prospect theory and fuzzy-trace theory. *Journal of Behavioral Decision Making*, 23(3):314–329, 2010.

13 Valerie F. Reyna. A new intuitionism: Meaning, memory, and development in fuzzy-trace theory. *Judgment and Decision Making*, 7(3):332–359, 2012.

14 Valerie F. Reyna and Charles J. Brainerd. Fuzzy-trace theory and framing effects in choice: Gist extraction, truncation, and conversion. *Journal of Behavioral Decision Making*, 4(4):249–262, 1991.

15 Murray Singer and Gilbert Remillard. Veridical and false memory for text: A multiprocess analysis. *Journal of Memory and Language*, 59(1):18–35, 2008.

16 Volker Stocké. Framing oder Informationsknappheit? Zur Erklärung der Formulierungseffekte beim Asian-Disease-Problem. In Ulrich Druwe and Volker Kunz, editors, *Anomalien in Handlungs-und Entscheidungstheorien*. Leske + Budrich, 1998.

17   Tom Trabasso and Linda L. Sperry. Causal relatedness and importance of story events. *Journal of Memory and Language*, 24(5):595–611, 1985.

18   Tom Trabasso and Paul Van Den Broek. Causal thinking and the representation of narrative events. *Journal of Memory and Language*, 24(5):612–630, 1985.

19   Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1981.

20   Amos Tversky and Daniel Kahneman. Rational choice and the framing of decisions. *Journal of Business*, 59(4:2):251–278, 1986.

21   Xiao Tian Wang. Framing effects: Dynamics and task domains. *Organizational Behavior and Human Decision Processes*, 68(2):145–157, 1996.

22   Xiao-Tian Wang and Victor S. Johnston. Perceived social context and risk preference: A re-examination of framing effects in a life-death decision problem. *Journal of Behavioral Decision Making*, 8(4):279–293, 1995.

Assessing Two-Mode Semantic Network Story Representations Using a False Memory Paradigm*

Steven R. Corman[1], B. Hunter Ball[2], Kimberly M. Talboom[2], and Gene A. Brewer[2]

1   Center for Strategic Communication
     Arizona State University
     Tempe, Arizona, USA
     `steve.corman@asu.edu`
2   Department of Psychology
     Arizona State University
     Tempe, Arizona, USA
     `{hunter.ball,kimberly.talboom,gene.brewer}@asu.edu`

──── **Abstract** ────

This paper describes a novel method of representing semantic networks of stories (and other text) as a two-mode graph. This method has some advantages over traditional one-mode semantic networks, but has the potential drawback (shared with n-gram text networks) that it contains paths that are not present in the text. An empirical study was devised using a false memory paradigm to determine whether these induced paths are remembered as being true of a set of stories. Results indicate that participants report false memories consistent with the induced paths. Implications for further research and two-mode semantic representations are discussed.

1   Introduction

In recent years interest has grown in the use of networks to represent relationships revealed in text, including the text of stories. The classic work in this vein is by Schank and Abelson who sought to "emulate the human conceptual mechanisms that deal with language" [20, p. 1] by specifying story scripts that people use to organize knowledge about the world. These works provide the basis for plans that can deal with novel situations (i.e., ones not covered by existing scripts). The scripts take the form of a network of connected acts and states. Classical work in artificial intelligence also sought to represent knowledge in natural language systems as directed networks of relations (e.g., class inclusion, actions) between objects (see [17, 23]).

Recently the idea of semantic networks has broadened to include any representation of connections between words in texts. Some researchers derive these networks using simple n-grams of proximate words [8, 9, 3, 21], on the assumption that semantic meaning flows between adjacent words. Some introduce elements of syntactic analysis to such networks to reduce noise they can otherwise contain [4].

──────────

Recently we have sought to represent the text of stories as a network of relations between entities and actions using semantic role labeling [5]. We initially considered using one-mode networks of entities connected with directed edges labeled with actions, in the manner of classical semantic networks. While this method avoids a problem to be discussed below, it does so at the cost of excluding possibly important information from the representation of the story. To illustrate this, consider the following simple story:

> John had an appointment with Bill. Bill went shopping then waited for John at Starbucks. John drove to Starbucks, but he had an accident on the freeway. He panicked. He could miss Bill. He waited for the police, then he walked to Starbucks and he just caught Bill.

Some simple subject-verb-object relationships exist in this story that can be straightforwardly represented as labeled directed edges.

bill $\xrightarrow{go}$ shopping

john $\xrightarrow{have}$ accident

john $\xrightarrow{catch}$ bill

For indirect objects that modify verbs, it is possible to combine the verb and preposition into an edge label, as is conventionally done in semantic networks.

bill $\xrightarrow{wait\_for}$ john

john $\xrightarrow{wait\_for}$ police

But it is unclear what to do with the phrase *on the freeway*, since it modifies the noun *accident*, rather than the verb *have*. One sentence *he [John] panicked* cannot be represented at all because it has no object. Another disadvantage is that since all the verb edges in the one mode representation are distinct, there is no way to represent that *have* occurs more frequently in this story (as one might be able to do with a weighted edge, for example).

A solution to these shortcomings is to represent the story as a two-mode graph [5], where mode 1 represents nouns, and mode 2 represents verbs. Directed edges in this graph are labeled S (subject_of), O (has_object), or [preposition]. The complete two-mode representation of the above story is shown in Figure 1.

Figure 1 A two-mode representation of the example story.

In addition to better representing some of the grammatical constructions, the two-mode scheme also allows for extraction of the individual modes, showing verbs that are related because of their connections to nouns, and nouns that are related because of their connections to verbs, as shown in Figure 2.

Figure 2 Extracted modes of the graph in Figure 1.

1.1 Induced Connections

However the two mode scheme has a potential problem. Where nouns are connected through a common verb, as is the case with *wait* in the example above, an ambiguity is created. Was it John or Bill who waited for the police? Which one waited at Starbucks? Both possibilities are readable from the two-mode graph, and this is reflected in the noun-mode graph as well, where both John and Bill are connected to police even though Bill never came in contact with them in the story.[1]

Our question is: Is this a "fatal flaw" of the two-mode representations provided their proposed relation to human cognition? Or could these induced paths generate semantic activation in the minds of subjects when they are remembering elements of stories? Some existing research suggests the latter outcome as a possibility. Research in human memory has examined the role of semantic activation in prompting false memories for events (for an excellent review see [10]). Generally, this class of theories proposes that false memories emerge from infelicitous semantic activation occurring when a recognition probe mismatches information stored in memory. This spurious activation is misinterpreted by individuals leading them to claim that they remember the probe. If two-mode networks structurally resemble the semantic network of participants' memories for stories then induced paths from these networks should provide fertile grounds for falsely remembering story elements that never occurred in a story. To address the possibility we examined whether these "induced paths" from two-mode networks derived from stories would reflect sources of semantic activation that lead to false memories.

1.2 False Memories

Examination of the factors involved in the creation of false memories provides a fruitful method of investigating the underlying mechanisms involved in the organization of human

---

[1] This problem is even worse for n-gram representations, which do not restrict connections to grammatical relations. A 3-gram graph of the story (with stop-words eliminated) has connections between *john* and *shopping*, *police* and *walk*, *accident* and *starbucks*, etc., none of which reflect events in the story.

memory. Memory errors have been studied using a variety of methods and have been found to occur when new words or sentences are similar in meaning to studied items (e.g., [1]), for new items that are visually or phonologically similar to old items (e.g., [14]), during free recall tests [22], and during eyewitness testimony [16]. A method that has been widely used to investigate false memories is the Deese-Roediger-McDermott (DRM) paradigm wherein after studying a list of related words such as *bed*, *rest*, *tired*, and *dream*, people often erroneously claim that a non-presented critical lure (*sleep*) was originally included in the list [6, 19]. Thus, the critical lure serves as an overarching theme of the related study list but the theme itself is never actually studied.

One theory that has been proposed to explain the DRM illusion is the activation/monitoring theory. This theory posits that activation of critical lures occurs during processing of list items via spreading activation of conceptual representations within a semantic network [11]. During encoding, a summation of multiple implicit associative responses produced by the studied associates may internally activate the conceptual representation of the critical lure thus making it available in memory [13]. During retrieval, test probes may serve to reactivate the associative network that subsequently makes the critical lure susceptible to false remembering due to a high degree of overlap between the lure and its activated representation within the associative network [15]. Thus, the relationships formed between items or sentences may allow extraction of the overall theme of the associated list or cause the critical lure to become activated and available in memory, depending on the subjective organization imposed by the participant during encoding [10].

To date, only a handful of DRM studies have placed associates in the context of sentences or stories during encoding. Placing related items in the context of sentences or text (e.g., "after work she was very tired", "after dinner she laid down in bed") influences false memories [7, 18], and new sentences that are related to a previously heard story are often falsely identified as having been previously studied [2]. These findings suggest that people make inferences and associations that are consistent with the overall meaning of a passage or text. For example, Dewhurst, Pursglove, and Lewis [7] found that 5-year-olds were more likely to falsely recognize critical lures after reading stories with DRM associates placed within the sentences of the text as compared to the standard list encoding (although there were no differences for 8- and 11-year-olds). They argued that the story context made it easier for younger children to identify its overall theme, whereas older children may have a greater ability to identify thematic associations during list processing and therefore had equally high false alarms in both encoding conditions.

In this experiment we employed a similar false memory paradigm in which participants read two stories at encoding and later made recognition decisions for sentences that either occurred (i.e., targets) or did not occur (i.e., lures) in stories that we developed. More specifically, we chose the two lure types that differed based on their occurrence in the two-mode network generated from the stories. For example, one lure type occurred in the two-mode network but did not occur in either story (i.e., an *induced path*; "Bill waits for the police" from the story above). Another lure type chosen from the stories reflects a path that neither occurred in the stories nor in the two-mode network but is composed of nouns and verbs in the graph (i.e., *synthesized path*; "Bill drove to Starbucks"). If induced paths in two-mode networks represent underlying semantic structure of human memories for stories then participants should incorrectly claim that induced path statements occurred in the stories more than synthesized path statements. Therefore, we hypothesized that statements from induced paths would falsely remembered more often on average than statements from synthesized paths.

2 Methods

2.1 Participants

A total of 34 undergraduate students from the Arizona State University were recruited from the psychology subject pool. Each participant was tested individual or in groups of two to eight in sessions that lasted approximately 20 minutes.

2.2 Materials

We created two stories of about 500 words each as stimuli for the experiment (see Appendix). Both were set in the context of the current Syrian civil war in order to maximize possibilities for finding induced paths in the resulting two-mode semantic networks. One was a "Fighters" story in which a group of men cross the border from Turkey into Syria in hopes of joining the fight against government forces. The other was a "Refugees" story about a family living in a Syrian refugee camp who cross the border from Syria into Turkey in order to escape the fighting.

Both stories were represented as two-mode, directed semantic graphs of the kind described above, and these were joined to produce an aggregate graph representing both stories. A researcher then examined the paths in comparison to the original stories to produce three kinds of stimulus sentences. *Old sentences* were simple sentences in the stories corresponding to paths present in the graph. Two types of new sentences were composed from elements in the graphs. *Induced sentences* were based on paths present in the graph but not present in the stories. *Synthesized sentences* were plausible sentences composed from nodes in the graph but were no represented by paths in the graph.

2.3 Design and Procedure

Participants were informed that they were going to read a set of stories and that their memory would later be tested on the content of the stories. The two stories used were counterbalanced such that the "Fighters" story was read first by half the participants, whereas the "Refugees" story was read first by the other half of the participants. Encoding was self-guided, such that when participants finished with a sentence or paragraph they pressed the spacebar to continue to the next screen. At the end of the first story, participants were informed that they had finished the first story and could begin reading the second story when ready. Separate stories were demarcated simply with the heading "Story 1" and "Story 2". Upon conclusion of the study phase, a 5-minute distractor phase consisting of solving Sudoku puzzles was administered. Following this, instructions for the recognition test were given. Participants were told they were going to be shown a series of sentences. Upon presentation, they were to think back to the stories earlier and if they remembered seeing the presented sentence in one of the stories they were to press the "old" key. If the sentence was new, they were to press the key labeled "new" to indicate that they did not read the sentence in the previous stories. The test phase consisted of 20 old sentences and 40 new sentences. Of the old items, 10 were taken from each of the two stories. The new items consisted of 20 synthesized sentences and 20 induced sentences. Upon conclusion of the test phase participants were debriefed and thanked for their participation.

3 Results

The proportion of sentences correctly recognized as old ($M = .60$) was greater than the proportion of new sentences incorrectly called old ($M = .20$), $t(33) = 13.97, p < .001, d = 2.69$. Critically, however, there was a significant difference in the proportion of new sentences incorrectly called old across the different classes of new items. The proportion of false alarms to synthesized sentences ($M = .23$) was greater than false alarms to induced sentences ($M = .18$), $t(33) = 2.43, p < .05, d = .239$.

4 Discussion

Overall, we found that our participants" abilities to discriminate target from lure statements drawn from the stories was fairly accurate. However, participants also made quite a few false alarms, which is reasonable considering that all of the lures consisted of nouns, verbs, and objects that occurred in the stories. Contrary to our hypothesis, participants made more false alarms, on average, to synthesized than induced lures. These results were unexpected in the sense that they were not predicted by our motivating hypothesis that induced paths in two-mode networks represent semantic structure that leads to ambiguous activation at test. These results have experimentally and theoretically interesting implications regarding false memories for textual passages. In the current study we have developed a novel extension of the basic false memory paradigm for investigating false memories from narrative. The combination of this paradigm with semantic network modeling suggests new and interesting methodological innovations for studies of false remembering linked to nodes in semantic networks.

Theoretically, these results indicate that semantic network modeling can provide useful information about the underlying semantic structure in which inappropriate activation occurs. However, it is important to note that additional methodological and theoretical work is needed to clarify the link between semantic space derived in these models and semantic space in the human mind. For example, longer delays between the reading the stories and taking a recognition memory task may lead to a reversal of the false alarm effects. Based on this line of reasoning, induced paths may reflect stronger thematic connections in semantic memory that are invariant to forgetting. In the current study we focused on only semantic origins of false memories. In future work additional structural and discourse related overlap could be manipulated across stories to promote false memories.

Implications of these results for the two-mode semantic networks are also interesting. The fact that we observed false-alarms in both the induced and synthesized paths suggest that when people read stories they do indeed form memories about connections between entities and events that are not present in the story itself. This means that if the objective of a representation of the text is to reflect the way it is remembered by readers (as opposed to providing a "factual" representation of the text per se), then the induced paths in the two-mode networks are not necessarily a fatal flaw in the approach. Further research is needed on the reasons behind the higher false alarm rate for synthesized vs. induced paths; our intuition is that introducing a delay between reading and remembering may reverse this effect. If so, then it strengthens the case for a two-mode over a one-mode semantic graph representation, as this would provide evidence that memories are structured in a way similar to the two-mode graphs.

The findings are also interesting outside the domain of computerized representations of stories. If false memories induce connections between story elements where no explicit links exist, then this provides fertile ground for interpretations that connect different stories into a

system—what we [12] call a narrative—even when the particulars of those stories may not strictly support this integration. This could explain why stories tend to change from their original forms over time as their accounts circulate in discourse and are passed from person to person. If these changes can be considered random, they may even hint at a mechanism for narrative evolution.

References

1   J. D. Bransford and J. J. Franks. The abstraction of linguistic ideas. *Cognitive Psychology*, 2(4):331–350, 1971.

2   A. L. Brown, S. S. Smiley, J. D. Day, M. A. R. Townsend, and S. C. Lawton. Intrusion of a thematic idea in children's comprehension and retention of stories. *Child Development*, 48:1454–1466, 1977.

3   K. M. Carley. Extracting team mental models through textual analysis. *Journal of Organizational Behavior*, 18:533–538, 1997.

4   S. R. Corman, T. Kuhn, R. McPhee, and K. Dooley. Studying complex discursive systems: Centering resonance analysis of communication. *Human Communication Research*, 28(2):157–206, 2002.

5   S. R. Corman, S. W. Ruston, and M. Fisk. A pragmatic framework for studying extremists' use of cultural narrative. In Gavriel Salvendy and Waldemar Karwowski, editors, *Advances in Human Factors and Ergonomics*. CRC Press, 2012.

6   J. Deese. On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, 58(1):17–22, 1959.

7   S. A. Dewhurst, R. C. Pursglove, and C. Lewis. Story contexts increase susceptibility to the DRM illusion in 5-year-olds. *Developmental Science*, 10(3):374–378, 2007.

8   J. Diesner, T. Frantz, and K. M. Carley. Communication networks from the Enron email corpus: It's always about the people, Enron is no different. *Computational and Mathematical Organization Theory*, 11(3):201–228, 2005.

9   M. A. Doerfel and G. A. Barnett. A semantic network analysis of the International Communication Association. *Human Communication Research*, 25(4):589–602, 1999.

10   D. A. Gallo. *Associative illusions of memory: False memory research in DRM and related tasks*. Psychology Press, New York, NY US, 2006.

11   D. A. Gallo. False memories and fantastic beliefs: 15 years of the DRM illusion. *Memory & Cognition*, 38(7):833–848, 2010.

12   J. R. Halverson, H. L. Goodall, and S. R. Corman. *Master Narratives of Islamist Extremism*. Palgrave-Macmillan, New York, 2011.

13   T. W. Hancock, J. L. Hicks, R. L. Marsh, and L. Ritschel. Measuring the activation level of critical lures in the Deese-Roediger-McDermott paradigm. *American Journal of Psychology*, 116(1):1–14, 2003.

14   T. C. Jones and L. L. Jacoby. Feature and conjunction errors in recognition memory: Evidence for dual-process theory. *Journal of Memory and Language*, 45(1):82–102, 2001.

15   D. R. Kimball, W. J. Muntean, and T. A. Smith. Dynamics of thematic activation in recognition testing. *Psychonomic Bulletin & Review*, 17(3):355–361, 2010.

16   E. F. Loftus. Leading questions and the eyewitness report. *Cognitive Psychology*, 7(4):560–572, 1975.

17   A. B. Markham. *Knowledge Representation*. Psychology Press, New York, 1998.

18   G. Plancher, S. Nicolas, and P. Piolino. Influence of suggestion in the DRM paradigm: What state of consciousness is associated with false memory? *Consciousness and Cognition: An International Journal*, 17(4):1114–1122, 2008.

19   H. L. Roediger and K. B. McDermott. Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 21:803–814, 1995.

20   R. C. Schank and R. Abelson. *Scripts, Plans, Goals, and Understanding.* Erlbaum, Hillsdale, NJ, 1977.

21   A. E. Smith. Automatic extraction of semantic networks from text using Leximancer. In *HLT-NAACL 2003, Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, May 27–June 1, Edmonton, Canada, Companion Volume: Demonstrations*, pages 23–24, 2003.

22   N. Unsworth and G. A. Brewer. Individual differences in false recall: A latent variable analysis. *Journal of Memory and Language*, 62(1):19–34, 2010.

23   W. Van Atteveldt. *Semantic Network Analysis: Techniques for Extracting, Representing, and Querying Media Content.* Booksurge, New York, 2008.

A   Appendix: Stories Used in the Study

A.1   Fighters Story

Three young men drive along a dirt road in a small silver car. The road winds along the Turkey side of the Turkey/Syrian border. The car stops on the side of the road. One of the men exits the car. His name is Asan. Asan grabs a pair of binoculars and looks at the valley. The other men exit the car too.

"That is the border between Turkey and Syria there," Asan points.

Army troops swarm the area.

The other young men, Raji and Abdullah, make frustrated noises. They hoped this part of the border would be clear. Raji and Abdullah are part of a larger group of young men. These men want to sneak into Syria to join the Free Syrian forces.

"We need to cross soon," says Raji. "I want to fight."

"We have one more chance," replies Asan. "I know someone who might still be able to help us make the crossing."

The three men return to the car. Asan drives in a different direction. Soon the car arrives in a small village. Asan parks in front of a small house. The three men exit the car. An old man sits at a home-made table in front of the house. The old man smiles at Asad. He offers tea to the young men. The cups are mismatched.

They drink the tea. Asan asks if the old man still knows the secret way into Syria. The old man says yes. He tells them about a narrow, winding path through the hills. It will take the fighters around the border patrols. The Army forces have not found it.

"It is dangerous," the old man warns the men. "You must walk and go at night or they might see you."

Asan wants to wait for one more night but Raji is impatient. Raji, Asan and Abdullah drive back to their apartment and collect the rest of the group. Asan tells them to pack food and water in backpacks that they can carry. Once they cross the border they will be met by members of the Free Syrian forces.

The group crams into three cars and they drive back to the beginning of the trail. Night falls as the men get out of the cars. The drivers will stay behind and wait for the next group who wants to cross into Syria. Asan wishes luck to Raji and Abdullah. Raji has the only flashlight for the group. He takes the lead and finds the small trail. The men walk in single file. No one talks. They walk for hours. They are not caught.

Morning comes. The group arrives in Syria. They are met by Free Syrian Forces. They have crossed the border.

A.2 Refugees Story

Ahmed stands on a dusty road. The Syrian/Turkey border lies in front of him. A city of tents and shacks sprawls behind him. Refugees live in the city. Ahmed was a teacher in Aleppo. His students have all left. Assad's Army turned the city into a war zone. They shelled houses indiscriminately. The Free Syrian forces destroyed the rest of the city trying to fight back. Ahmed and his family now try to escape the conflict. Ahmed hopes to take his family to Turkey where they will be safe. There are people who will smuggle refugees across the border. The cost is expensive.

Ahmed and his wife sold the last of their possessions several days ago. They live in the refugee camp now.

"This country was my home," Ahmed says. He points in the direction of Aleppo. "I lived in Aleppo my entire life. I met my wife there. My daughter and son were born there. Now the city is destroyed. I must leave. I must keep my family safe."

Ahmed walks back through the refugee camp. People look sad and scared. Ahmed talks to no one.

He steps inside a shelter made from a blue tarp. A pretty woman dishes beans into four small, chipped bowls. There is no table. Two children sit on the floor. They eat quickly. Ahmed thanks his wife. Another man walks up to the tent.

Ahmed greets Niam. Ahmed invites the younger man inside. Niam declines. "I have checked the road. It is clear. We leave tonight. Be ready at sundown. Make sure you bring plenty of water," Niam says.

Ahmed nods. He looks at Fatima, his wife. She looks scared. It will be a long walk into Turkey. The children are young. They will get tired. "We will be ready," Ahmed says.

At sundown Niam returns. Ahmed and Fatima carry food and water in their backpacks. Fatima puts two blankets in each backpack for the children. The children carry nothing. Niam leads them out of the refugee camp. Night falls. Niam turns on a flashlight. He hands another flashlight to Fatima. Fatima shines it in front of the children.

The group walks for hours. Ahmed and Niam talk softly. Fatima tells the children stories. The moon rises in the sky. The children grow tired. Ahmed and Fatima pick them up and carry them.

The moon begins to set. Ahmed and Fatima grow weary. Suddenly Niam points his flashlight.

"There," Niam says. They see a lake. The moonlight shines on the water. "You are safe. Welcome to Turkey," says Niam.

Processing Narrative Coherence: Towards a Top-Down Model of Discourse

Erica Cosentino¹, Ines Adornetti², and Francesco Ferretti³

1   Department of Philosophy
    University of Calabria
    Cosenza, Italy
    `ericacosentino@libero.it`
2   Department of Philosophical Researches
    University of Roma "Tor Vergata"
    Rome, Italy
    `inesado@yahoo.it`
3   Department of Philosophy, Communication and Visual Arts
    University of Roma Tre
    Rome, Italy
    `francesco.ferretti@uniroma3.it`

──────── **Abstract** ────────

Models of discourse and narration elaborated within the classical compositional framework have been characterized as bottom-up models, according to which discourse analysis proceeds incrementally, from phrase and sentence local meaning to discourse global meaning. In this paper we will argue against these models. Assuming as a case study the issue of discourse coherence, we suggest that the assessment of coherence is a top-down process, in which the construction of a situational interpretation at the global meaning level guides local meaning analysis. In support of our hypothesis, we explore the role of executive functions (brain functions involved in planning and organization of goal-oriented behaviors) in coherence's establishment, discussing the results of several studies on narrative abilities of patients with brain injuries. We suggest that, compared to other models of discourse processing focused on comprehension, our model is a viable candidate for an integrated account of discourse comprehension and production.

1   Introduction

In spite of the criticisms to generative grammar in recent years, Chomsky's concept of language continues to be the standard model in cognitive science. At the base of this model are the assumptions that the structure of the internal constituents of the sentence represents the core of language and that, consequently, the general device at the basis of language is a module specialized in the analysis of the syntactic structure of the sentence. But what does this model tell us about human narrative abilities? Even if Chomsky does not address the issue explicitly, scholars who are inspired by universal grammar and who are interested in the study of narrative processing maintain that the building of the coherent flow of discourse (the basis of any narrative abilities) must be interpreted in terms of syntactic parser functioning. In their view, in fact, the principles that govern the sequence of sentences in the construction of the discourse are specified in some constituents of the sentence's structure. According

to this perspective, information processing that underlies narrative abilities has a strong "bottom-up" character: the global level of discourse is gained starting from the analysis of the single sentences actually said, through a sequential process of accumulation of information (e.g., [36]).

In this paper we propose an alternative model of language and narrative abilities. Against the Chomskian idea of the primacy of the sentence, we propose that the ability to process discourse takes priority over the ability to process sentences. Such a proposal, which we characterize as a "top-down" hypothesis, implies the adhesion to two aspects particularly relevant to the analysis of narrative abilities. First, sentence comprehension is driven by a prior comprehension of the coherence of the flow of discourse: the understanding of the narrative flow has logical and temporal priority on the comprehension of sentences. Second, discourse comprehension implies the involvement of cognitive devices other than those implicated in the analysis of the constituents structures of sentences. Contrary to the devices specialized in the analysis of the syntactic structure, indeed, the processing systems involved in the analysis of discourse have a specific character of projection. While the bottom-up model is firmly anchored to the processing of what the speaker effectively says in a given moment, the top-down model we are proposing is largely fueled by information on what she has already said (projections into the past) and on what she is going to say (projections into the future). What kind of devices can analyze this kind of information?

From a general point of view, our idea is that the projection devices responsible for the construction of discourse are similar to those involved in the processes of navigation. The relationship between narrative abilities and spatial navigation is a good metaphor empirically grounded. Experimental data corroborate, indeed, the idea that the devices involved in the analysis of discourse (ones that allow us to "keep the route" to reach the goal we have in mind) are closely related to systems that allow us to navigate through space and time [14, 18, 19, 20]. In this paper we will analyze only a limited aspect of the relationship between narrative abilities and navigation: the role played by a specific device of projection (the executive functions of action planning) in the building of the "global coherence" of discourse. In general, the projection devices can ensure a strong attachment of the expressions uttered in the flow of speech to the context. For this reason, the model of the navigation can explain the link between linguistic expressions and the extralinguistic context [14, 18, 20]. In this paper we focus the attention on the theme of global coherence. This theme leads us to consider the link between sentences in the intralinguistic context. The hypothesis that underlies our work is that human narrative abilities cannot be explained only in terms of local coherence (cohesion). In spite of the key role assigned to the syntactic constituents by the standard model, indeed, some properties of discourse may be explained only by referring to the global coherence. As we'll see, in this regard a key role is played by executive functions.

This paper is divided into three parts: in the first (pars destruens) we discuss the nature and limits of the bottom-up model of language and discourse according to which the phrase has a logic and temporal priority on the discourse. In the second part (pars construens) we present a top-down model of language and narrative abilities in which the discourse has a logical and temporal priority on the sentence. In the third part, referring to the literature on narrative pathologies, we present empirical evidence for our model, analyzing the processing systems that regulate the coherence of the flow of discourse. We have to specify that, although discourse and narrative are two phenomena not completely equivalent, in this paper we discuss them together as we examine coherence that is a fundamental property that is common to both.

2   The primacy of the sentence and the bottom-up model of discourse

The idea of language that emerges from the perspective of classical modularity is reflected in the dedication that Fodor [23] addressed to Garrett:

> One day—it must have been five years or so ago—my friend, colleague, and sometime coauthor Merrill Garrett made what seems to me to be the deepest remark that I have yet heard about the psychological mechanisms that mediate the perception of language. "What you have to remember about parsing," Merrill said, "is that basically it's a reflex." This work is, in effect, a sustained meditation on Merrill's insight, and it is gratefully dedicated to him.

This idea that language is a reflex, actually means it is a way of treating language as a module. The similarity between modules ("stupid" systems that operate in an automatic and mandatory way) and reflexes depends on an evolutionarily significant reason: the processing speed [53, 46]. The stupidity of the modules is an adaptive strategy through which the cognitive system can process large amounts of information quickly. The amount of information that language can transmit is amazing. Without a fast processing system, human communication probably would be doomed to failure. But the speed comes at a cost. The devices specialized for the rapid processing of language must have an automatic and mandatory character; they have to be able to focus exclusively on certain aspects of the stimulus. Devices of this kind must have, in other words, specific bottom-up properties. The reference to bottom-up devices has important consequences for the understanding of the nature of human communication.

In the standard version of cognitive science, communication is interpreted in terms of the code model [48], a model that Fodor [22, p. 106] considers "not just natural but inevitable". According to this model, "we have communicated when you have told me what you have in mind and I have understood what you have told me" [22, p. 109]. The fact that verbal communication relates to the sharing of thoughts between speaker and listener raises a number of interesting questions. The most important one for our current purposes concerns what makes possible the transformation of thought into language. How is it possible that a mental state (the nature of which is abstract and conceptual) can be coded in a physical structure (a succession of sounds, in the case of verbal language) and, through this step, can be communicated to the receiver? How is it possible, in other words, that the structure of sounds is able to respect the structure of mental states?

The move taken by the proponents of the standard model to answer this problem is to question the structural isomorphism between language and thought. The argument of Fodor [22, 24] is clear: language can express thought because it reflects the basic structure of thought, the "logical form." The point is particularly relevant for our purposes. The primacy assigned to the phrase by the advocates of standard conception depends heavily on the propositional nature of thoughts. The thesis of the isomorphism between the constituents structure of thoughts and the syntactic structure of sentences that express them has a direct impact on the understanding of human communication. According to Fodor (as well as Chomsky and all the authors who adopt the standard model), the analysis of language is entirely governed by a device specifically used in the processing of the logical form (syntax, basically) of the utterances. From such a perspective, contextual information it is not only irrelevant but also harmful: everything needed to understand what the speaker said is encoded in the utterance.

The thesis of the isomorphism between language and thought fits perfectly with the idea that the processing devices have a bottom-up character. At the base of this character is the

mechanistic nature of the language module that is considered a self-sufficient processing system. Each time the linguistic device detects the appropriate stimulus in the environment (utterances that exhibit a logical form), the comprehension follows as a automatic and mandatory consequence of the processing of the stimulus. This mechanistic idea of comprehension is highlighted by Fodor, who argues that "one cannot avoid hearing a phrase that has been said (in a known language) as a phrase that has been said" [23, p. 91]. In this paper, we do not care to discuss this issue. What interests us here is to note that, in the standard version, the analysis of the logical form of the sentences is a necessary and sufficient condition for the comprehension of the content expressed by those sentences and that such analysis implies a specific processing device wired on the syntactic properties of the sentences.

Bottom-up models in cognitive science have always been considered, because of the emphasis on the mechanical and material aspects of the information processes, the trump card in the debate on the naturalization of the mind [23][1]. In spite of these considerations, our idea is that the bottom-up models based on the primacy of the phrase are founded on a highly abstract concept of language. Indeed, the processing device, focalized on the syntactic aspects of the phrases, has to operate independently from any background noise: so to speak, it must analyze the shape of the proposition in its "purity." The models that explain human narrative capabilities using the theoretical paradigm of the primacy of the phrase have enormous explanatory difficulties, as we shall see in the next section. These difficulties will lead us to change the interpretative model.

3   Processing coherence: how local and global processes are intertwined

According to standard compositional theories, sentences encode meaning by the means of a context-free rule-based combination of lexical–semantic features of the words within a sentence. This step of the comprehension process is considered necessary, and it corresponds to the level of sentences' truth conditions. Such a thesis is a "literal meaning thesis" and is the basis of all traditional semantic theories. According to the literalist thesis, the contribution of world knowledge to the truth conditions is limited to cases of indexicality and ambiguity; this means that the role of the context of utterance should be traceable to syntactic elements in the logical form of sentences. Currently, many scholars interested in language functioning recognize that processes of the type just described are not enough to account fully for language comprehension and production. They seem to agree that, at a certain point, context is taken into account. The disagreement is about when, exactly, this happens.

Here it might be useful to draw a distinction between a two-step model of linguistic comprehension and a one-step model. Classical theories of meaning are two-step models, according to which contextual information is considered only after establishing phrase or sentence local meaning. On the basis of a one-step model, contextual information may be used in a more top-down fashion, such that the local contribution of individual words or sentences is a function of the construction of a situational interpretation at the global meaning level. In this article we argue against two-step interpretations of language comprehension,

---

[1] It must be said that there have been multiple attempts to devise Chomsky-style narrative grammars (e.g., [39, 38]) and that these operate top-down. Nevertheless, although they represent a step forward compared with Chomskian perspective, our idea is that such attempts are ineffective to account for the global coherence of the discourse since it is a pragmatic property that cannot be explained in terms of universal grammar.

and we propose a one-step model according to which the wider discourse context has an immediate effect on the unfolding linguistic information.

From this perspective, our criticism extends also to some models of discourse elaborated in the field of discourse analysis. Even if from the 1980s onward the explicit goal of people working in that field has been going beyond the sentence, our specific charge here is that models of discourse elaborated within the classical compositional framework (e.g., [36]) are still characterized by the idea that sentence analysis has priority over discourse analysis. The model proposed by Kintsch and van Dijk [36], for example, is a bottom-up model, which is centered on the role of mental propositions expressed by predicate-argument structure. According to this model, the structure of a text or a discourse ("macrostructure") can be formally derived from the structure of the relations between sentences (which form the "microstructure") by the means of the application of some general rules. This model, then, fits into a two-step perspective of the interpretation process as discourse meaning is inferred only once sentences' analysis is completed. Such a model has been very popular among linguists during the last thirty years, and it still represents the dominant view on discourse processes. In spite of this, we argue that this model, as well as other two-step models, can be seriously undermined by challenging two key psycholinguistic assumptions that lie behind it.

First, the model assumes the incremental nature of the interpretative process, which means that the processing of coherence is based on a word-by-word analysis and integration. This is the very essence of two-step models because incrementality is consistent with the idea that local meanings are built up from the meanings of individual words, which in turn is consistent with compositionality principle. Second, the orthodox view of language comprehension is that the processes involved are fully completed, namely, that semantic information for each word is fully retrieved during the incremental process. For example, Just and Carpenter [35] state: "readers interpret a word while they are fixating it, and they continue to fixate it until they have processed it as far as they can" (p. 30).

From these two assumptions follows the classical compositional view that local coherence is established prior to global coherence and that a thorough check at the local level is part of the normal process of coherence establishment. In the current paper, these two assumptions are called into question, presenting arguments in favor of alternative claims that (1) discourse processing is driven by global processes oriented to topic maintenance as opposed to the maintenance between utterances, (2) the extent to which an item is analyzed at a local semantic level is a function of the general fit of that item to the discourse context. The higher the global fit of the item, the lower the accuracy of the semantic processing of that item is. In order to test these hypotheses, we focus on anomaly detection. In particular, the case study is the *survivor effect* observed in the way certain discourses are interpreted naturally by human subject, and described below.

4 Towards a top-down model of discourse

The idea that discourse processing is based on the construction of a coherent mental representation of what is in the discourse is a widely accepted view. In particular, the aim of comprehension and production processes is the construction of an integrated representation that reaches a "coherence threshold." The question is how that threshold is reached. As mentioned, according to the standard compositional view, discourse processing proceeds in a bottom-up fashion, constructing a complete representation based on a thorough check of each sentence. From this perspective local meaning is established prior to global meaning [36, 44].

Accordingly, McKoon and Ratcliff [43] have proposed the minimalist hypothesis that claims people only try to establish global coherence when there is a break in meaning at the local level. Alternatively, discourse-level information may be used in a more top-down fashion, such that the local contribution of individual words or sentences is a function of the construction of situational interpretation at the global meaning level. In this latter view, partial or incomplete semantic analysis at the sentence level would often be sufficient to fulfill the comprehender's coherence need. Looking for discourse coherence, listeners would activate global instead of local processes, namely processes that are oriented to the maintenance of global coherence (topic maintenance), as opposed to local coherence (maintenance between utterances). A number of studies have provided evidence supporting this latter view.

For the purposes of this paper, studies concerning a well-documented effect during language processing, the "survivor effect," are particularly relevant. In the following example, subjects were asked to write solutions to a version of the subsequent problem (adapted from [6]:

> A tourist flight crashes in the Pyrenees, and wreckage is strewn equally in France and Spain. Where should the survivors be buried?

Results show that only 66% of the subjects noted that survivors are not the sorts of things that should be buried. Even more striking, when the term "survivors" was replaced by the phrase "surviving dead," only 23% of the subjects noted any anomaly. The extremely low detection rate suggests that local semantics of the phrase "surviving dead" is not computed prior to its incorporation into the more global representation of the text. If it were, the anomaly should be noted at that initial stage. It would seem that subjects understand the story by developing a global, situational interpretation of the discourse. To the extent that the coherence at the global level can be maintained (for example, "dead" is consistent with the global plane-crash situation), local problems are ignored and perhaps not even computed.

To support this conclusion further, it should be noted that if the critical expression ("surviving dead") is embedded within an incongruent context, for example a "bicycle accident," then subjects' reaction to the anomalous phrase is completely different. This time they will more easily notice the anomaly, and changing the scenario can manipulate the relevance of critical phrases. This is extremely significant as it shows that the extent of semantic analysis the critical item receives is a function of the general fit of the item to expectations based on context. If the global fit of phrases in the context is high, then more detailed, effortful, time-consuming analysis may not take place. In contrast to strict bottom-up, incremental interpretation, these findings are consistent with the idea that coherence's establishment is a top-down process guided by listener's expectations. Once the system has a satisfactory level of information supporting coherence, further analysis might not take place. That said, the question we need to address now is how exactly the process of coherence's evaluation works at cognitive architectural level. Which are the principles and the actual cognitive components that guide the top-down assessment of coherence?

We mentioned that cognitive processes of projection in space and time (i.e., navigational abilities) may have a crucial role in the processing of coherence. In the current paper we will not go into detail about the general navigational framework of communication (for an extended presentation of the model (see [14, 21, 20]). Here we will rather focus on a very specific aspect of navigational abilities, i.e., the contribution of executive functions. Before turning to this, we should note that our approach fits into a cognitive pragmatic conceptual framework and thus shares some general aspects of other pragmatic accounts of language, in particular Relevance Theory (RT;[50]. However, as we discuss elsewhere [14], several

characteristics distinguish our model from RT's model. For our current purposes, it will be sufficient to contrast a "theory of mind" account of pragmatics with a "navigation" framework on the basis of two related claims. First, we claim that a navigational model provides a richer notion of context compared to RT's model. According to a navigational model, the interpretative process is guided not only by the attribution of mental states to interlocutors, but also by the exploration of spatial and temporal perspectives that, even if activated by the current circumstances, represent alterative states of the actual situation. In this view, context is defined by the concurrent functions of grounding and projection [21], which allows the individual to take into account the extra-linguistic world by projecting himself toward spatial and temporal alternative scenarios.

The second relevant aspect concerns more closely the very notion of coherence. We would like to point out that most of the models of discourse processing discussed so far emphasize the comprehension side of the interpretative process. Classical compositional models are focused on mental processes of understanding what an interlocutor or a text expresses, providing a bottom-up analysis of such a process. However, RT's model acknowledges the role of top-down processes, but it also seems to be limited to the analysis of linguistic comprehension. Indeed, the notion of relevance, which is the key notion of RT's model, is much more concerned with comprehension than production. In the last part of this paper we will argue that this aspect may be considered a "side effect" of RT's model and that a more powerful unifying model, in which comprehension and production are placed side by side, can be reached by elaborating on the notion of coherence. As we will see, the one-step model we are proposing, according to which coherence is processed in a top-down fashion and guided by the role of executive functions, presents itself as a viable candidate for a unifying model of discourse processing.

5 It is not only a matter of relevance: coherence intuitions

Our idea is that, in order to propose a unified model of discourse processing, it is necessary to analyze not only the processes of interpretation, but also the processes of production. We discuss this idea through the analysis of the cognitive devices involved in the establishment of coherence. To clarify the issue, we begin by highlighting once again the aspects that distinguish our model of pragmatics from that of RT's.

According to the model proposed by Sperber and Wilson, verbal communication's burden falls mainly on the listener, who engages to reconstruct, through inferential chains, the speaker's intention. In fact, the main reason why RT's is primarily a model of comprehension lies in the adoption of Grice's assumption according to which, in the communication processes, the starting point is the intention of the speaker [30]. From this point of view, the speaker's intention is a phenomenon already given and each verbal cue introduced by the speaker it is necessary for the listener in order to reconstruct that intention. Now, while Grice [31] through the formulation of conversational maxims, tried to give an account of the processes involved in language production, in RT's model an explanation of this kind lacks.

Sperber and Wilson [49] argue that the main purpose of RT's model, and more generally of pragmatics, is to clarify the nature of the processes and the skills that allow the listener to reconstruct inferentially the intention of the speaker on the basis of the sentence's meaning. In fact, they characterize pragmatics as inferential comprehension oriented to relevance detection. As Wilson [55] points out, "the main aim of relevance theory in the domain of verbal communication is to explain how utterances are understood" (p. 58). At a general level, the inferential comprehension is made possible by a specific cognitive system, the theory

of mind (ToM) module, that underlies the ability to attribute mental states such as beliefs, intentions, and feelings to others and to explain and to predict the actions that derive from them (e.g., [5, 37]). What is important to note is that relevance theorists see pragmatics as a specific component, a relevance-based comprehension module, of the ToM module with its own proprietary concepts and procedures distinct from general ToM module [11, 49]. This means that the relevance principle characterizes, from a pragmatic point of view, the essence of pragmatics.

RT's model (heavily focused on the aspects of language comprehension and on principle of relevance) evidently represents an overly limited view of pragmatics and, consequently, a limited view of human communication [1]. Relevance, indeed, is not the only principle that governs communication. As highlighted, for example by Giora [27, 28], "speakers and hearers are not constrained only by the search for relevance. In addition, coherence considerations constrain communication and play a major role in discourse structuring and understanding" [27, p. 31]. To see how this is possible, one must analyze some verbal expressions and discuss them in reference to the notions of RT's model. Though such an analysis may appear extremely technical, it is important for the purposes of our argument to show that 1) relevance is not the only property of communication and 2) discourse coherence has a key role in pragmatic processes.

According to Sperber and Wilson [50, 56] an input (e.g., an utterance or a memory) is relevant to an individual when it connects with background information she has available to yield conclusions that matter to her. More in detail, an input is relevant to an individual when its processing in a context of available assumptions yields a positive cognitive effect, that is to say, a worthwhile difference to the individual's representation of the world (e.g., a true conclusion). The most important type of cognitive effect achieved by processing an input in a context is a contextual implication, a conclusion deducible from the input and the context together. Besides the cognitive effect, relevance of an input relies also on processing effort. Other things being equal, the greater the processing effort, the lower the relevance of an input to an individual in a given time. Thus, the relevance of an input for an individual at a given time is a positive function of the cognitive benefits that he would gain from processing it and a negative function of the processing effort needed to achieve these benefits. Giora points out that there are cases where the verbal productions are inappropriate from a pragmatic point of view because they lack coherence, but they are relevant (in Sperber and Wilson's terms) to an individual. To illustrate the point, Giora starts from a central notion of RT's: the choice or selection of the context. Sperber and Wilson argue that in communication the context is not given beforehand, but is open to choices and revisions during the process of comprehension. There are several ways through which it is possible choose or expand a context. For example, the listener, in order to understand a specific statement uttered during a conversation, can include in the context the interpretation of preceding utterances and/or the interpretation of her responses during the course of the dialogue. The relevant point is that, according to Sperber and Wilson, relevance determines the selection of the context; the set of assumptions that allows to get the best balance between processing effort and cognitive effect is chosen as the appropriate context.

According to Giora [27] an idea of this type is problematic because it may lead to situations in which the information may be relevant (in Sperber and Wilson's terms), although pragmatically inappropriate because it lacks coherence. To clarify the point, Giora discusses an example proposed by Sperber and Wilson [50, p. 125]. Consider the context composed by a), b) c) and the utterances 1), 2), 3), 4) and 5) below:

- a) People who are getting married should consult a doctor about possible hereditary risks to their children.
- b) Two people with thalassemia should be warned against having children.
- c) Susan has thalassemia.

1. Susan, who has thalassemia, is getting married to Bill.
2. Bill, who has thalassemia, is getting married to Susan.
3. Bill, who has thalassemia, is getting married to Susan, and 1976 was a great year for French wines.
4. Susan and Bill should consult a doctor about possible hereditary risks to their children.
5. Susan and Bill should be warned against having children.

The point is to establish what utterance is the most relevant, that is to say, what has the greater cognitive effect at the minimum processing effort. Sperber and Wilson state that (1) and (2) are equally difficult to process because they are similar in length and require the same context (a-c). However, (2) has greater cognitive effects (contextual implications) than (1), while (1) has only one contextual implication. For example (4), (2) has an additional contextual implications (5). So, Sperber and Wilson state that (2) is more relevant that (1).

Giora [27]; ([28], however, on the basis of the assumption that contexts are searched for, states that the context needed to render (1) relevant is smaller than that needed to render (2) relevant. In fact, in order to render (1) relevant, only two assumptions (a, c) should be activated. Instead, to process (2), it is necessary to add (b) to the context. So, (2) is not really more relevant than (1); (2) has more cognitive effects than (1), but it also needs more effort processing, necessitating the expansion of the context.

The same procedure (of extending the context) may apply to render a discourse such as (3) relevant, albeit inappropriate. Sperber and Wilson affirm that (2) is more relevant than (3); they have the same amount of contextual implications, but (3) requires more effort because the extra information in (3) is completely unrelated to the given context and, consequently, has no contextual effect. However, since contexts are searched for, it is possible to extend the context so as to render (3) relevant. For example, the speaker and hearer of (3) should have heard that a neighbor bought them a 1976 bottle of French wine. In the initial context (a, b, c) now there is a new assumption:

- d) Our neighbor bought us a 1976 bottle of French wine.

This extended context (a, b, c, d) renders (3) as equally relevant to the context as (2) is. While (3) requires more processing effort than (2), it also has more contextual effects. Thus, the utterance (3) "Bill, who has thalassemia, is getting married to Susan, and 1976 was a great year for French wines" is the more relevant in the context (a, b, c, d).

However, Giora outlines that, in spite of its relevance, (3) is an incoherent text that the people evaluate as inappropriate: the hearer must be left puzzled as to how the two propositions in (3) are related to each other (rather than to a context). Then, because (3) is more relevant than (2), the information that listener evaluates as the most relevant is (2) (even after the extension of the context). Why? Giora's hypothesis is that such a evaluation depends on the fact that the speaker and the listener are driven in communication processes not only by intuitions of relevance, but also by intuitions of narrative coherence. If relevance's detection were the only basic principle of human communication, in fact, the hearer (given the expanded context) should automatically consider (3) as the most pertinent information. But, although relevant, (3) is pragmatically inappropriate. Now, important for the scope of our argument is that the inappropriateness of (3) depends on the fact that

it violates the listener's intuitive expectations of coherence. Thus, the hearer's reluctance to consider (3) more relevant than (2) is because (3) is incoherent. The existence of such reluctance shows that relevance is not the only principle that regulates the communicative exchanges. An important role in this regard it also played by the principle of narrative coherence. Now, although the examples just described are related to the intuitions of the listener (i.e., to the interpretive processes) as previously mentioned, our idea is that in order to give an account of the cognitive devices involved in the processing of discourse it must analyze the processes of both comprehension and production. In the next section we will see how the reference to executive functions as the processing systems of coherence allows us to present a unifying one-step model of discourse processing that takes into account both interpretative and productive processes.

6 The key role of executive functions in the building of narrative coherence

As we have seen, coherence refers to conceptual organizational aspects of narration at the suprasentential level. Thus, the coherence of a narrative discourse depends, at least in part, on the speaker's ability to maintain thematic unity [2]. When is a discourse coherent? A dominant idea, especially among linguists, is that the coherence of a narration depends on the linear relations between adjacent sentences, that is to say on cohesion between pairs of consecutive sentence (e.g., [7, 10, 32, 52]). For example, consider the following text:

> After the forming of the sun and the solar system, our star began its long existence as a so-called dwarf star. In the dwarf phase of its life, the energy that the sun gives off is generated in its core through the fusion of hydrogen into helium [8, p. 2].

In this text the sentences are connected through lexical cohesion; the lexical cohesive relations hold among the lexical items sun, solar system, star, dwarf star and dwarf phase in the text. Now, although the cohesive relations (the local meaning) have an important role in the expression and recognition of coherence relations, the cohesion between consecutive sentences seems an unnecessary and insufficient condition for the narrative coherence (see also [26]). With reference to this a crucial distinction is that between global and local coherence. Global coherence is the manner in which discourse is organized with respect to an overall goal, plan, theme, or topic; it refers to the relationship between the content of a verbalization with that of the general topic of narration. Local coherence concerns the conceptual links between individual sentences or propositions that maintain meaning in a text or discourse [29]. Now, while the local coherence is made possible by cohesion relationships, the same is not true for global coherence. Consider for example the following text:

> I bought a Ford. The car in which President Wilson rode down the Champs Élysées was black. Black English has been widely discussed. The discussions between the presidents ended last week. A week has seven days. Every day I feed my cat. Cats have four legs. The cat is on the mat. Mat has three letters. [16, pp. 110–111].

In this text the sentences are connected through the cohesive mechanism of repetition. However, the set of sentences, despite the abundance of cohesive ties, is not perceived as a coherent whole. In this text the sentences do not hang together in a reasonable way; the text lacks of global coherence. So we can argue that global coherence of a narration is independent from cohesion, that is to say, the macrostructure of a narrative discourse cannot be formally derived by the microstructure of the sentence.

The idea that global coherence of a narration is independent from cohesion has received much evidence in recently from neurolinguistics research (e.g., [15, 40, 42]). These studies have highlighted the dissociation between the abilities that underlie sentence processing (microstructure or microanalysis) and those that underlie narrative processing (macrostructure or macroanalysis). Particularly relevant for our purpose are the data that come from studies of patients with traumatic brain injury (TBI). These subjects generally have impairments of specific prefrontal areas. Such impairments cause deficits of executive functions. Executive functions (EFs) is an umbrella term for a wide range of cognitive and behavioral skills whose main neural substrate is constituted by the prefrontal cortex. EFs have a key role in regulating the equilibrium between the organism and the environment. In fact, they are implicated in the temporal organization of goal-oriented behaviors (e.g., [4, 25, 51, 54]). EFs allow formulating a plan, starting its execution, and maintaining attention (perseveration) on that plane until its realization. Moreover, EFs allow the rapid shift of attention for the adaptation to novel contexts, while they inhibit inappropriate behavioral responses to the current situation.

Although there no precise taxonomy of executive functions (see [3, 34]), it is possible to delineate some aspects of convergence among the neuropsychological models. There is a general agreement that EFs are implied in processes such as planning, working memory, inhibition, and mental flexibility, as well as in the initiation and monitoring of action [12]. Numerous studies have shown that TBI subjects generally have deficits in EFs of action planning and monitoring. Because of such deficits, the behaviors of TBIs appear confused and disordered; they cannot organize and complete goal-oriented behaviors because they are not able to conceptually formulate and execute a sequence of actions [17, 47, 57]. Our hypothesis is that the deficits of planning and monitoring in these patients are the principle causes of their problems in narrative discourse. Our idea, in fact, is that EFs of planning and monitoring play an important role also in building the global coherence of a narrative. As global coherence is the manner in which discourse is organized with respect to an overall goal, plan, theme, or topic [29], the building of narrative coherence should be conceived as a specific case of goal- oriented behavior. From our perspective, it is possible to conceive coherence as the way to achieve the general goal (the general topic) toward the narrative discourse tends. Thus, the establishment of coherence implies a form of goal planning (a conceptual formulation of the general topic) and a form of organization of the single steps necessary to achieve that goal (organization of the single verbal expressions tied to the general topic). Moreover, during the execution of a plan—that is, during the stage of narrative production—it is necessary to continue estimation of the task in order to make sure that the elements introduced are in accordance with the general topic of conversation [19].

Analyses of the narrative production of TBI patients confirm this idea. Many neurolinguistics studies have shown that these patients connect sentences correctly by using cohesion ties (grammatical devices), but they are unable to construct and maintain the global coherence of their verbal productions because they cannot relate the individual sentences to a plan or to a more general purpose, and often introduce material that is irrelevant to the current context in their verbal productions (e.g., [9, 29, 33, 41, 13]). Coherence appears to be controlled by a higher-order conceptual process, whereas lexical cohesion may be driven by more automated linguistic processes that are not disrupted after TBI. As an example of this fact, consider the following narrative produced by a TBI subject:

> I have got faults and. my biggest fault is. I do enjoy sport. it's something that I've always done. I've done it all my life. I've nothing but respect for my mother and father and. my sister. and basically sir. I've only come to this conclusion this last

two months. and. as far as I'm concerned. my sister doesn't exist. [45, p. 305].

In this text the sentences are well formed from a strictly syntactic a point of view; the single local sentences are not problematic. However, taken as a whole, this fragment of speech is pragmatically inappropriate because it lacks global coherence. In fact, it is characterized by sudden and irrelevant changes of topic. What is important to note is that, as Biddle and colleagues [9, p. 463] pointed out, "the narrative impairment of adults and children with TBI [. . .] appeared to be the result of problems with planning, production and monitoring discourse".

Our idea is that examples of this kind, which show a dissociation of microlinguistic and macrolinguistic cognitive functions, provide support for the distinction between microstructural and macrostructural discourse component. More specifically, confirming the idea that coherence is processed in a top-down fashion, they support a unifying top-down model of discourse processing according to which the global meaning of a narration constraints in a substantial way the local meaning of the sentence.

7   Conclusions

Classical models of language functioning in cognitive science have been characterized by bottom-up models, which are centered on sentences' analyses. In this paper we have argued that the priority given to sentences' analyses undermines classical models' capacity to explain narrative processing because it undermines their capacity to explain a crucial property of narration: coherence. In particular, we have suggested that coherence processing is a top-down process in which the construction of an interpretation at the global meaning level takes priority over local meaning analysis. Analyzing the processing systems that underlie narrative coherence, we have shown that such a property has to be explained by focusing on macro-analysis rather than on microanalysis. Evidence regarding narrative abilities of TBI supports the distinction between microstructural and macrostructural discourse component and suggests that the processing of discourse and the processing of sentence are based on different cognitive devices. Specifically, these data show that discourse processing does not rely on devices involved in the structural analysis of the internal constituents of the sentence. Moreover, they support our hypothesis that coherence is processed in a top-down fashion by cognitive systems oriented to the future (anticipation of the general theme of narration). This general top-down account of coherence processing, according to which discourse global meaning constrains local meaning analysis in a substantial way, provides a unifying framework for discourse comprehension and production processes.

──── **References** ────

1   I. Adornetti. Why philosophical pragmatics needs clinical pragmatics. *Humana.Mente*, 23:159–174, 2012.

2   M. Agar and J. Hobbs. Interpreting discourse: Coherence and the analysis of ethnographic interviews. *Discourse Processes*, 5(1):1–32, 1982.

3   J. A. Alvarez and E. Emory. Executive function and the frontal lobes: A meta-analytic review. *Neuropsychology Review*, 16:17–42, 2006.

4   R. Barkley. *Executive Functions. What They Are, How They Work and How They Evolved.* The Guildford Press, New York, 2012.

5   S. Baron-Cohen. *Mindblindness: An Essay on Autism and Theory of Mind.* The MIT Press, Cambridge, 1995.

6   S. B. Barton and A. J. Sanford. A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory and Cognition*, 21(4):447–487, 1993.

7   I. Bellert. On a condition of the coherence of texts. *Semiotica*, 2:335–363, 1970.

8   I. Berzlánovich. Lexical cohesion and the organization of discourse. First year report, University of Groningen. http://www.rug.nl/let/onderzoek/onderzoekinstituten/clcg/berzlanovich.pdf.

9   K. Biddle, A. McCabe, and L. Bliss. Narrative skills following traumatic brain injury in children and adults. *Journal of Communication Disorders*, 29:447–469, 1996.

10   W. Bublitz. Cohesion and coherence. In J. Zienkowski, editor, *Discursive Pragmatics*, pages 37–49, Amsterdam/Philadelphia, 2011. John Benjamins Publishing Company.

11   R. Carston, S. Guttenplan, and D. Wilson. Introduction: special issue on pragmatics and cognitive science. *Mind and Language*, 17:1–2, 2002.

12   R. C. K. Chan, D. Shum, T. Toulopoulou, and E. Y. H. Chen. Assessment of executive functions: Review of instruments and identification of critical issues. *Archives of Clinical Neuropsychology*, 23(2):201–216, 2008.

13   C. Coelho, K. Lê, J. Mozeiko, F. Krueger, and J. Grafman. Discourse production following injury to the dorsolateral prefrontal cortex. *Neuropsychologia*, 50(14):3564–3572, 2012.

14   E. Cosentino and F. Ferretti. Communication as navigation: A new role for consciousness in language. *Topoi*, to appear.

15   G. A. Davis, T. M. O'Neil-Pirozzi, and M. Coon. Referential cohesion and logical coherence of narration after right hemisphere stroke. *Brain and Language*, 56(2):183–210, 1997.

16   N. E. Enqvist. Coherence, pseudo-coherence, and non-coherence. In J.-O. Ostman, editor, *Cohesion and Semantics*, number 41 in Meddelanden från Stiftelsens for Åbo Akademi Forskningsinstitut, pages 109–128, 1978.

17   P. Eslinger, G. Zappalà, F. Chakara, and Barrett A. M. Cognitive impairments after TBI. In N. D. Zasler, D. I. Katz, and R. D. Zafonte, editors, *Brain Injury Medicine. Second Edition*, pages 60–61, New York, 2011. Demos Medical Pub.

18   F. Ferretti. Navigation, discourse and the origin of language. In A. Marini and P. Brambilla, editors, *Brain evolution, language, and psychopathology in schizophrenia*. Psychology Press, (in press).

19   F. Ferretti and I. Adornetti. Discourse processing and spatial navigation. In B. Kokinov, A. Karmiloff-Smith, and N. J. Nersessian, editors, *European Perspectives on Cognitive Science, Proceedings of the European Conference on Cognitive Science, EuroCogSci 2011*, NBU Series in Cognitive Science, Sofia, 2011. New Bulgarian University Press.

20   F. Ferretti, I. Adornetti, E. Cosentino, and A. Marini. Keeping the route and speaking coherently: the hidden link between spatial navigation and discourse processing. *Journal of Neurolinguistics*, 26(2):327–334, 2013.

21   F. Ferretti and E. Cosentino. Time, language and flexibility of the mind: The role of mental time travel in linguistic comprehension and production. *Philosophical Psychology*, 26(1):24–46, 2013.

22   J. Fodor. *The Language of Thought.* Thomas Y. Crowell, New York, 1975.

23   J. Fodor. *The Modularity of Mind.* The MIT Press, Cambridge, 1983.

24   J. Fodor. *Psychosemantics: The problem of Meaning in the Philosophy of Mind.* MIT Press/Branford Books, Cambridge, 1987.

25   J. Fuster. *The Prefrontal Cortex.* Academic Press, London, fourth edition, 2008.

26   R. Giora. Notes towards a theory of text coherence. *Poetics Today*, 6(4):699–715, 1985.

27   R. Giora. Discourse coherence and theory of relevance: Stumbling blocks in search of a unified theory. *Journal of Pragmatics*, 27:17–34, 1997.

28   R. Giora. Discourse coherence is an independent notion: A reply to Deidre Wilson. *Journal of Pragmatics*, 29:75–86, 1998.

29   G. Glosser and T. Deser. Patterns of discourse production among neurological patients with fluent language disorders. *Brain and Language*, 40:67–88, 1990.

30   P. Grice. Meaning. *Philosophical Review*, 66:377–388, 1957.

31   P. Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Speech Acts*, number 3 in Syntax and Semantics, New York, 1975 (1967). Academic Press.

32   M. A. K. Halliday and R. Hasan. *Cohesion in English*. London, Longman, 1976.

33   M. S. Hough and I. Barrow. Descriptive discourse abilities of traumatic braininjured adults. *Aphasiology*, 17(2):183–191, 2003.

34   M. B. Jurado and M. Rosselli. The elusive nature of executive functions: A review of our current understanding. *Neuropsychological Review*, 17:213–233, 2007.

35   M. A. Just and P. A. Carpenter. A theory of reading: From eye fixations to comprehension. *Psychological Review*, 87(4):329–354, 1980.

36   W. Kintsch and T. Van Dijk. Toward a model of text comprehension and production. *Psychological Review*, 85(5):363–394, 1978.

37   A. M. Leslie. ToMM, ToBY, and Agency: Core architecture and domain specificity. In L. A. Hirschfeld and S. A. Gelman, editors, *Mapping the Mind: Domain Specificity in Cognition and Culture*, pages 119–148, Cambridge, 1994. Cambridge University Press.

38   J. Mandler. Narratology and cognitive science: A problematic relation. *Style*, 44(4):469–95, 2010.

39   J. Mandler and N. Johnson. Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*, 9(1):111–151, 1977.

40   A. Marini, S. Carlomagno, C. Caltagirone, and U. Nocentini. The role played by the right hemisphere in the organization of complex textual structures. *Brain and Language*, 93:46–54, 2005.

41   A. Marini, V. Galetto, E. Zampieri, L. Vorano, M. Zettin, and S. Carlomagno. Narrative language in traumatic brain injury. *Neuropsychologia*, 49:2904–2910, 2011.

42   A. Marini, I. Spoletini, I. A. Rubino, M. Ciuffa, G. Banfi, A. Siracusano, P. Bria, C. Caltagirone, and G. Spalletta. The language of schizophrenia: An analysis of micro- and macrolinguistic abilities and their neuropsychological correlates. *Schizophrenia Research*, 105:144–155, 2008.

43   G. McKoon and R. Ratcliff. Inference during reading. *Psychological Review*, 99:440–446, 1992.

44   D. C. Mitchell and D. W. Green. The effects of context and content of immediate processing in reading. *Quarterly Journal of Experimental Psychology*, 30(4):609–637, 1978.

45   M. R. Perkins, R. Body, and M. Parker. Closed head injury: assessment and remediation of topic bias and repetitiveness. In M. R. Perkins and S. J. Howard, editors, *Case Studies in Clinical Linguistics*, pages 293–320, London, 1995. Whurr.

46   S. Pinker. *How the Mind Works*. Norton, New York, 1997.

47   T. Shallice. Specific impairment of planning. *Philosophical Transaction of the Royal Society, Biological Science*, 25:199–209, 1982.

48   C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. The University of Illinois Press, Illinois, 1949.

49   D. Sperber and D. Wilson. Pragmatics, modularity and mind-reading. *Mind and Language*, 17:3–23, 2002.

50   D. Sperber and D.(1986) Wilson. *Relevance: Communication and Cognition*. Blackwell, Oxford, second edition, 1995.

51 D. Stuss and D. Benson. *The Frontal Lobes.* Raven, New York, 1986.

52 S. K. Tanskanen. *Collaborating towards Coherence. Lexical cohesion in English Discourse.* John Benjamins, Amsterdam/Philadelphia, 2006.

53 J. Tooby and L. Cosmides. The psychological foundations of culture. In Barkow J., L. Cosmides, and J. Tooby, editors, *The Adapted Mind*, pages 19–136, Oxford, 1992. Oxford University Press.

54 M. C. Welsh and B. F. Pennington. Assessing frontal lobe functioning in children: Views from developmental psychology. *Developmental Neuropsychology*, 4:199–230, 1988.

55 D. Wilson. Discourse, coherence and relevance: A reply to Rachel Giora. *Journal of Pragmatics*, 29:57–74, 1998.

56 D. Wilson and D. Sperber. Relevance theory. In L. R. Horn and G. Ward, editors, *The Handbook of Pragmatics*, pages 607–632, Oxford, 2004. Blackwell.

57 T. Zalla, C. Plassiart, B. Pillon, J. Grafman, and A. Sirigu. Action planning in a virtual context after prefrontal cortex damage. *Neuropsychologia*, 39:759–770, 2001.

Ontological Representations of Narratives: a Case Study on Stories and Actions*

Rossana Damiano[1,2] and Antonio Lieto[1]

1   **Dipartimento di Informatica**
     Università di Torino
     Torino, Italy
     `rossana@di.unito.it, lieto@di.unito.it`
2   **Centro Interdipartimentale di Ricerca su Multimedia e Audiovisivo**
     Università di Torino
     Torino, Italy

―――― **Abstract** ――――

In this paper, we describe the narrative ontological model encompassed in the Labyrinth system. The aim of the system is to allow users to explore a digital archive by following the narrative relations among the resources contained in it. Targeted at cultural heritage applications, the Labyrinth project relies on the notion of "cultural archetype", i.e., a core representation encompassing archetypical stories and characters, exploited as a conceptual framework for the access to archives of heterogeneous media objects.

In particular, we describe how the system leverages various types of ontological reasoning to let narrative relations emerge between artworks, and exemplify how these relations are exploited by the system to provide the user with a narrative conceptual framework she or he is familiar with in the exploration of the archive.

1   Introduction and Motivations

In recent years, the advent of digital media has enabled the publication on the Web of a huge quantity of resources, i.e., images, audiovisual objects, text documents and their combination. However, as the number of online digital contents increases, the way they are described is far from meeting the requirements of content-based access required by the general public: neither their description in terms of editorial metadata nor the tags added by users seem adequate to describe the content of media resources, and fall short of providing an effective access to digital media. The classification of resources in terms of stylistic features (layouts, patterns, colour profiles, etc.) is inadequate as well to the users' needs.

In the field of cultural heritage, in particular, a content based description of artworks is required. As shown by the media studies, when users tag artworks, they tend to describe the content of the artworks [33]. For visual arts, stories can provide an effective way to mediate between the users and the description of the artwork by using a conceptual model users are familiar with [4]. In many cases, the content of visual artworks can be described

---

in terms of some narrative situation, i.e., a basic framework where a character is caught in the act of doing something, using some instrument, possibly with the participation of other characters. The role of stories in the description of artworks is explicitly acknowledged by iconology. For example, consider a painting representing a mythological subject, such as Ariadne: the character of Ariadne is sometimes represented in the act of giving to Theseus a ball of thread (that he will employ to escape from labyrinth of Minos); or, she is depicted as abandoned by Theseus in the island of Naxos; or else, following a different myth, she is represented with Bacchus. In the Iconclass system for iconological classification[1], the first subject corresponds to the *id* 94M34 ("Ariadne gives Theseus a ball of thread"), the second one to the *id* 95B(ARIADNE)61 ("Ariadne left behind on the island of Naxos"), where Bacchus will find her later on (*id* 92L1211, "Bacchus finds Ariadne on Naxos") [17].

Thanks to the practice of imitation [19], in Western culture, the same subject is represented multiple times across authors and ages: in the neoclassical painting by Pelagio Pelagi ("Arianna dona il gomitolo a Teseo", 1814), in the painting by Jean–Baptiste Regnault ("Ariane et Thésée", 1827), etc. Moreover, the same stories are the subject of many other artworks, conveyed through different media: for example, consider "Ariadne auf Naxos" (1916), an opera by Richard Strauss, or the peplum film "Teseo contro il Minotauro" (Italy, 1960). In the digital age, remediation [3] contributes to keeping this practice alive, with narrative contents being adapted from texts, to films and comics and so on, meeting the expectations of different audience types and the distribution requirements posed by different devices. Finally, narrative situations are linked to each other to form larger stories, which provide further, more indirect connections between artworks: consider, for example, the "story of Ariadne" (encoded in Iconclass as "95B(ARIADNE) (story of) Ariadne"), which, as a more general class, encompasses the single stories mentioned above.

This paper describes an ontology-based approach to the description of the narrative content of artworks, implemented in the Labyrinth system. The aim of the Labyrinth system is to test the feasibility of using narrative concepts for the exploration of media archives, in the field of cultural heritage. Artworks, in fact, often have a narrative content, but span only single episodes of larger stories; moreover, they often do it through audio visual languages that are not available for text processing. The use of an ontological model, where narrative situations are described and grouped into stories, allows implicit relations to emerge among artworks through reasoning: for example, starting from their narrative features (represented characters, action, location, etc.), different artworks can be inferred as displaying the same episode, or different episodes of the same story. Narrative relations, then, are exploited to create navigation paths among the artworks, or to generate recommendations of similar contents.

Targeted at cultural heritage dissemination and digital publishing, the Labyrinth project relies on the notion of "cultural archetype", i.e., a core representation encompassing archetypal stories and characters, proposed as a conceptual framework for the access to archives of heterogeneous media objects. In the Archetype ontology, stories are represented as containing a set of actions, enacted by characters in a given location and time, and described according to a role–based schema. A story can encompass other simpler stories, which in turn are composed of increasingly simpler actions. The ontology has been designed with the goal of supporting reasoning on the relations among characters, actions and stories, while abstracting from different genres and media types. In this paper, we illustrate the Archetype Ontology and show, by resorting to examples, how the ontology supports narrative reasoning on the

---

[1] http://www.iconclass.org

relations between stories, actions and character. Finally, we illustrate how the narrative relations obtained through the reasoning process are employed in the system to support the navigation in the archive.

The paper is organised as follows: after surveying the related work (Section 2), in Section 3 we briefly describe the system within which the proposed model of story is being developed and tested. In Section 4, the conceptual model of the archetype ontology is described. Section 5 shows how the connections among artefacts and stories (or actions) is obtained. Section 6 illustrates how the inferences generated by the reasoner on the narrative model are employed to support the narrative based navigation among the artworks. Conclusions and future work end the paper.

2   Related Work

In the last decade, the access to cultural heritage and the distribution of media objects have moved toward a digital convergence [22]. In cultural heritage, this process has taken the form of digital platforms, such as online museums, cultural websites, etc. aimed at encouraging the access to the cultural heritage by the general public (consider, for example, the Europeana web portal).[2] In parallel, the advent of new media has pushed forward re-mediation practices [3]: according to the paradigm of re-mediation, the contents of one medium are re-focused onto another medium (like for example, the transposition of a novel into a film or the reuse of movie contents in videogame design). Despite the scenario described above, the convergent culture has not been effective in creating tools for organizing and accessing contents in the field of cultural heritage. In today's web, searching media objects, in fact, is still largely based on keywords and/or tags: the search outputs a list of objects (books, pictures, videos, etc.) but does not contain an explicit representation of the narrative relations they entertain with the input keywords.

The use of ontologies for the exploration of media archives has been explored by several research projects. A pioneering contribution in the use of ontologies to provide online access to cultural heritage is given by the CultureSampo project [20]. This project encompasses a set of domain ontologies, which provide the background against which cultural objects, encoded in different media formats, can be explored, tracking the connections among them [21] in terms of geographical and chronological relations, authorships, production processes, etc. The folkloric saga called Kalevala, also encoded in an ontological form, is employed to describe and connect the episodes referred to by artworks.

The DECHO system [1] relies on a conceptual model of the archeological domain to support the exploration of cultural heritage objects. Targeted at the integration of different data sources, the ontology has been developed on top of the CIDOC Conceptual Reference Model (CIDOC-CRM) [12]. In particular, the system encompasses digital images, 3D models of objects and environment, and narratives, using the ontology to establish connections among them. Mainly oriented toward the interaction with 3D virtual environment and objects, the DECHO system integrates in a unifying semantic framework an advanced 3D visualisation tool and a corpus of textual documentation about the displayed objects, including narratives.

Narrative is the focus of the Bletchley Park Text system [28], a semantic system designed with the goal of supporting the users in the exploration of online museum collections. Designed with the notion of the "guided visit" in mind, the system relies on an ontology of story, taken from the Story Fountain project [29]. The stories represented in the system are exploited to

---

[2] http://www.europeana.org

create relations between the entities contained in the online collections, allowing the user to query the system for a semantic path between entities. Similarly to [28], the Labyrinth project mainly relies on narrative concepts for the users' conceptualisation of resources [28]. However, the Labyrinth project is not targeted at the fruition of a specific (virtual or physical) collection: rather, it aims at exploiting narrative concepts to create an open system that leverages the reasoning capabilities of the Semantic Web technologies to let meaningful connections emerge within heterogeneous resources.

Concerning the use of ontologies to model narrative concepts, story ontologies have been proposed with two main goals, namely the purpose of classifying story types and the purpose of providing an underlying model for narrative annotation. A well known example of the first type of systems is the work by [15]. In this work, inspired by the work of Propp [32], an ontology of fairy tales, encoded in OWL, is exploited to model different plot types. The system uses the ontology to perform case-based reasoning: given a story plan, the system searches a similar plot in the ontology, measuring the semantic similarity of the given plot with the plot types encoded in the ontology. A natural language module, then, generates a textual version of the obtained plot, adapted to the input parameters (characters, situations, etc.) given by the user. In the same line, the work by [18] used automatic classification techniques to classify plot types. However, with the notable exception of the Opiate system [13], structuralist models not have received much attention in recent years, following the criticism that they do not provide the flexibility needed to face the challenges of new media.

Overcoming the differences among different media types and genres is a main challenge faced by the research in media annotation. In this field, story ontologies have been proposed as a way to provide a shared and interoperable model for annotation scenarios which are characterised by the presence of different types of narrative contents and rely on the paradigm of crowd–sourcing. A media-independent model is provided by the OntoMedia ontology, exploited across different projects (such as the Contextus Project [23, 24]) to annotate the narrative content of different media objects, ranging from written literature to comics and TV fiction. The OntoMedia ontology contains a very detailed model, tailored on story annotation, and mainly focused on the representation of events and the order in which they are exposed. In [24], it is employed to annotate common elements and plot across the different episodes of the Dr. Who sci-fi TV series. OntoMedia lends itself to the comparison of cross-media versions of the same story (for example, a novel and its filmic adaptation), while it does not cover in a detailed way the description of characters' behavior (intentions, roles, etc.).

The Cadmos project [25, 8] shares with these approaches the basic assumption that a media object can be segmented into meaningful narrative units and that, given some kind of formal or semi-formal annotation, these units can be accessed and navigated. In the Drammar ontology, the basic unit of the annotation is the "drama unit": based on the Aristotelean notion of unity of action, the unit contains one or more incidents, which can be either naturally occurring events, or intentional actions performed by the story characters in order to achieve some goals. A main innovation attained by the Drammar is the mapping of the annotation schema to linguistic resources (namely, FrameNet [2] and WordNet [27]), and the design of a meaning negotiation process [8] to let users select the appropriate concept for describing events in large common sense ontologies (such as YAGOSUMO [11]).

The SUMO ontology, although not specifically tailored on story modelling, has been employed for the task of story annotation and story generation. In [10], the axiomatic definition of processes, in SUMO, is exploited to reason on stories and to generate plots. This approach, although not directly relevant for story models, reveals the relevance of an accurate representation of actions (processes, in SUMO terminology) for story description.

The ontology encompassed by the Labyrinth system (Archetype ontology) incorporates some of the main tenets of the Drammar of ontology, such as the commitment towards a core story model, neutral with respect to different narrative theories and genres; differently from Drammar, which provides a character–based account of narrative concepts, it assumes that a library of basic stories, relevant to the cultural archetypes of Western culture, is edited by hand and exploited to link different media resources. Inspired by the approach proposed by [20], where a single folkloric saga has been employed as a red thread for the presentation of artworks, in the Labyrinth projects, the story model serves as a framework for the users' conceptualisation of artworks, in a domain (cultural heritage) where narrative is, across media and ages, a powerful metaphor for content description.

With respect to the story models encoded in [15, 23], which are oriented, respectively, to story classification and annotation, the narrative model of the Labyrinth system is oriented to the description of basic narrative situations, represented in terms of characters and objects participating to a narrative situation, and their organisation into larger stories according to a modular perspective. In this sense, it abstracts from the notions of plot types and genre, accounting for the distinction between entities and processes acknowledged by top level ontologies such as Dolce [26]. Being targeted at representing the narrative properties of artworks, the Archetype ontology has been designed with the goal of interoperability with the standard ontologies for media description, such as Media Ontology[3] and FRBR [30].

Finally, differently from other approaches issued by industrial and academic research (see, for example, Knowledge Graph project[4]), Labyrinth does not aim at creating a general infrastructure for representing semantic relations among media items, but limits its scope to a set of few relevant archetypes (the labyrinth, the hero, journey, etc.) which are pervasive in Western culture. The underlying assumption of this approach is that these archetypes (and their related stories) are limited in number and shared by the audience along geographical and temporal coordinates, providing a valid affordance to content access for the general public of new media.

3   Overview of the Labyrinth system

The Labyrinth system allows the user to explore a repository of media resources through the conceptual mediation of an "archetype". The user can see how the resources in the repository relate with the various element which compose the archetype model (places, stories, characters, objects, etc., described in Section 4), and how they are connected to each other through the links with the archetype they share (for example, resources displaying the same character or symbol, related to a certain archetype).

The interaction design of the system integrates a top-down, hypertextual exploration of the repository with a 3D environment (still under development). In the hypertextual mode, the system filters the contents according to their links to the reference archetype and shows how they are related with each archetype element (stories, characters, objects, locations, etc.). When a single artwork is reached, the user starts navigating the repository resource by resource, following the semantic relations between the artworks, in a way that resembles the walk through a maze. Semantic browsing and navigation are not limited to the explicit representation contained in the ontology, but leverage the inferences made by the ontology server.

---

[3] `http://dev.w3.org/2008/video/mediaann/mediaont-1.0/mediaont-1.0.html`
[4] `http://www.google.com/insidesearch/features/search/knowledge.html`

The Labyrinth system encompasses four main modules:

- The Ontology server maintains the ontology describing the archetypes, maps the media resources and their relations onto the ontological model, and provides the reasoning services.[5]
- The Media repository stores the resources and is managed through a relational database (DB).
- The Labyrinth Web application, written in Java, provides search and navigation functionalities by querying the ontology.
- The client side applications, i.e., the web site and the 3D plugin, support the interaction with the user.

Since the Labyrinth system merges the perspective of digital archives with the paradigm of new media, adding a new repository to the the system requires a two-step process, which includes data integration and editing: the data integration phase is accomplished through an internalization procedure, which translates the description of the resources to be incorporated (encoded in their metadata) into the ontological model of the system; in some case, the model must be edited to accommodate the new data.

The content creation pipeline, then, is the following: a domain expert examines the contents of the repository to be incorporated and suggests a list of archetypes which are relevant to the resources in the repository. Then, the ontology engineer edits the ontology to create the required archetypes (unless they are already present in the ontology). Adding new archetypes to the ontology, however, does not affect the archetype model, leaving the top-level classes of the ontology unmodified. Finally, the resources in the repository are mapped onto the archetypes contained in the system (and their audiovisual documentation is copied in the Media repository), through a set of mapping rules.

By using the online interface, the users can search and explore the resources contained in the repository based on the relations they have with the archetypes described in the ontology. The assumption is that "emergent" meaning relations, inferred by the semantic engine on the ontological model, can generate thematic (perhaps serendipitous) paths through the repository.

4 The archetype ontology

The Archetype Ontology (AO), designed for the Labyrinth system, relies on and incorporates, in a unifying model, multiple ontologies already available in the literature and representing different relevant aspects of the narrative of media objects. More precisely, the ontologies incorporated in AO are the following: the Ontology for Media Resource[6], a formal framework for describing media objects (e.g., images, videos etc.) according to their format (e.g., jpeg, avi etc.); the FRBR (Functional Requirements for Bibliographical Records) ontology[7], a framework for describing resources according to an abstract model; finally, part of the Drammar ontology [9], a core ontology for the representation of characters and actions in narrative units. Following a tradition dating back to Aristotle in drama studies and narratology [5, 31], the Drammar ontology acknowledges the primary role of character in

---

[5] Currently, the ontology server is provided by the Owlim RDF database management systems (`http://www.ontotext.com/owlim`).
[6] `http://dev.w3.org/2008/video/mediaann/mediaont-1.0/mediaont-1.0.html`
[7] `http://www.ifla.org/publications/functional-requirements-for-bibliographic-records`

■ **Figure 1** The top level of the Archetype ontology.

story, intended as the intentional agent of a sequence of causally related story incidents. The role of characters has emerged also in contemporary aesthetics, where it is considered the medium of audience involvement, thanks to the process of identification with characters [16, 7]

AO serves the explicit goal of representing the peculiar narrative aspects related to the definition of 'cultural archetypes'. Archetypes, in turn, serve the scope of describing the semantic relations occurring among a set of cultural and artistic artifacts. Therefore, for its own nature, the Archetype Ontology is not intended to cover the foundational aspects of narratives, but only the ones intrinsically connected with the concept of Archetype, intended as a core representation encompassing archetypical stories and characters. The ontology has been designed with the goal of supporting reasoning on these typ of narrative relations, that the Labyrinth system exploits to guide the user exploration of media archives.

The top level classes of AO are the represented in Figure 1. The **Archetypes** class contains the thematic archetypes to which a story can be referred; the **Artifact** class contains the media objects, organized according to the FRBR model; the **Dynamics** class (from Drammar) represents actions, processes and state of affairs involving the narrative entities; **Entity** contains the characters and objects represented in an artefact or involved in a story; **Story** represents a collection of stories; **Description Templates** (from the Drammar ontology) contains the role schema (**SituationSchema**) that can be filled by narrative entities (characters and objects) in a dynamics (i.e., a process or state), inspired by the "Situation Description" ontology pattern [14]; the **Format** class encodes the format and type of media resources; **Geographical Place** and **Temporal Collocation**, finally, represent the classes where it is possible to encode, respectively, the spatial and temporal information related to artefacts, stories and archetypes.

The above mentioned ontologies have been inserted in our model as follows: FRBR ontology has been used in order to describe the individuals belonging to the **Artifact** class. The four abstract levels for the resource description of the FRBR model [30] are: *work*, representing a certain intellectual or artistic creation (e.g., the *Faust*); *expression*, representing the different intellectual or artistic realization of a work (e.g., book, video etc.); *manifestation*, representing the physical embodiment of an expression (e.g., the Italian translation of the book of *Faust*) and *item*, the specific exemplar of a certain manifestation (e.g., the book number 32 of the Italian translation of the book of *Faust*). The incorporation
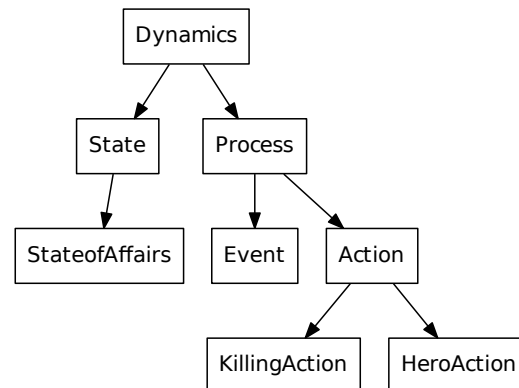
of this model allowed us to represent the status (in terms of ideation and production) of the media objects according to the commitments expressed in FRBR (however, in the Archetype Ontology, the "item" category was not used since this concepts is implicit in the membership of the class **Manifestation**). Furthermore, the Media Object ontology has been inserted at the level of **Format** class in our general model. Finally, the imported components of the Drammar ontology have been used to express the relations between the **Dynamics**, **Entity** and **Description Template** classes.

The "narrative branches" of the ontology are obtained by the following classes: **Archetype**, **Artifact**, **Story**, **Dynamics** and **Entity**. The **Archetypes** class currently contains two directed subclasses, namely **Hero** and **Labyrinth**, representing the two narrative myths ("Hero" and "Labyrinth" respectively) already encoded in the model. These classes have been specialized into other subclasses described with a set of necessary and sufficient conditions. For example: the class **MythologicalHero** has been defined as a subclass of **Hero** where the participating characters (or agents in Labyrinth terms) are restricted to the members of the **MythologicalEntity** class. Other properties (mainly corresponding to binary relations in first order logic) allow providing informative connections with other narrative classes. For example: the *isEvokedIn* property allows connecting archetypes and artefacts (e.g., with this property it is possible to state that the "Cnossos Labyrinth" archetype is evoked by a certain artefact, e.g., namely an artefact representing the characters of Ariadne and Theseus), the property *isRecalledIn* connects the archetypes with the stories (e.g., "Cnossos Labyrinth" can be connected to the story named "Ariadne gives Theseus a ball of thread").

The class **Artifact**, specialized, as mentioned above, according to the FRBR model, is connected with the other narrative classes by the following properties: *evokes*, inverse property of the above mentioned *isEvokedIn*, which connects **Archetypes** and **Artifact**; *displays*, connecting **Artifact** and **Entity** (e.g., this property allows asserting that a certain artefact, let suppose the painting "Arianna e Teseo", displays as characters the entities Ariadne and Theseus); *describeAction*, connecting Artifact and Dynamics (e.g., this property allows stating that the artefact "Theseus kills the Minotaur" describes the action of "killing"). In our model, the *describeAction* property has been defined as a sub-property of the more general property *hasPart*. This modeling solution has an impact on both the intended meaning of
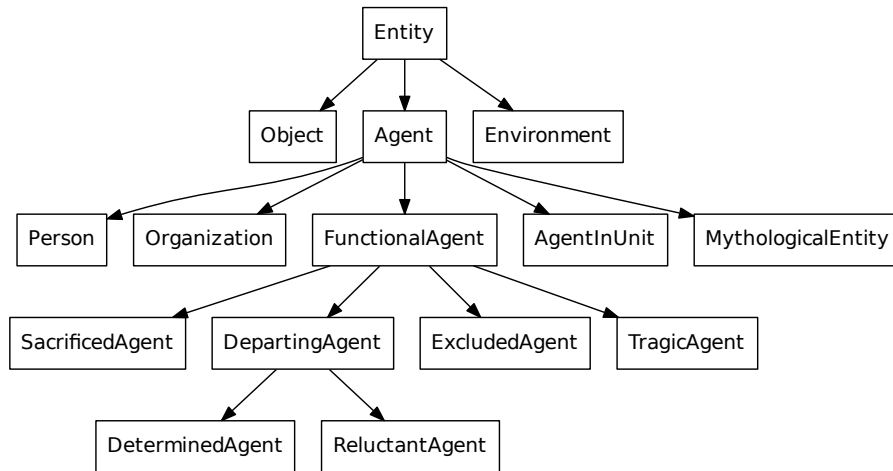
Figure 3 The Entity class.

describeAction (e.g., this property is intended as expressing a sense of "membership") and, therefore, on the reasoning processes coming up from the model. In fact, the statement that the painting "Theseus kills the Minotaur" describes the action of "killing" implies that this action is inferred as belonging to (via the part-of property) that artefact.

The class **Dynamics** has been imported from the Drammar ontology, where it represents the structure of the story incidents. It has been specialized as shown in Figure 2: its direct subclasses are **State** and **Process**, identifying different types of narrative situations. **Process** types have been divided in **Action** and **Event** and some actions, e.g., **Killing Action**, **Hero Actions** etc. have been constrained with necessary and sufficient conditions. For example: the class **Hero Action** is defined as the class of the Action instances having, as characters, some agent referred to the archetype of the Hero. The main properties connecting **Dynamics** with the other narrative classes in the ontology are: *hasCharacter*, connecting the **Dynamics** with the characters (**Entity**) participating in it, and *isDynamicsOf* (a transitive property connecting the **Dynamics**, e.g., actions or states, with the class **Story**). Even this property, as the *describeAction* property shown above, has been encoded as sub-property of *hasPart*. Therefore, it can be inferred that an action or a process are dynamics (intended as "part of") a certain story. Moreover, the stories are intended as composed by several dynamics.

The class **Story** is connected with the classes **Dynamics** (as illustrated before), **Entity** and **Archetype**. The connection with the class **Entity** is given by the already mentioned property *hasCharacter*, while the connection with the **Archetype** class is given by the property *recall* (e.g., it is possible to assert that a certain story "recalls" a certain archetype). At the current stage of development, the **Story** class is specialized with some subclasses defined according to necessary and sufficient conditions. The defined subclasses are **MythologicalStory** (or myths, defined as stories having as character some mythological entity), **KillingStory** (stories including a killing act), **CnossosStory** (stories of the Hellenic period recalling the Cnossos Labyrinth) and **StoryOfAriadne** (stories having as character Ariadne).

```
1   <owl:NamedIndividual rdf:about="&www;labyrinth#AriadneAndTheThread">
2    <rdf:type rdf:resource="&www;labyrinth#Story"/>
3       <hasCharacter rdf:resource="&www;labyrinth#Ariadne"/>
4       <hasCharacter rdf:resource="&www;labyrinth#Theseus"/>
5       <hasAction rdf:resource="&www;labyrinth#giving_the_thread"/>
6       <isPartOf rdf:resource="&www;labyrinth#MinotaurStory"/>
7       <recalls rdf:resource="&www;labyrinth#CnossosLabyrinth"/>
8       <recalls rdf:resource="&www;labyrinth#Theseus_Hero"/>
9       <hasTimePeriod rdf:resource="&www;labyrinth#Year2000BC"/>
10   </owl:NamedIndividual>
```

Figure 4 RDF description of the story "AriadneAndTheThread" ("Ariadne and the ball of thread") in the AO ontology.

Finally, the class **Entity** describes the agents (and objects or places) that have some narrative role within actions, stories and archetypes, and that are represented in some **Artefact**. It has been specialized according to the following hierarchy: **Agent**, **Object** and **Environments** are the direct subclasses of the root. **Agent** further specializes into the following subclasses: **Person**, **Organization**, **Mythological Entity** and **FunctionalAgent**. This latter class has as subclasses a set of different classes inspired by Propp's theory of functional roles in tales [32] (merged with some elements from [6]), such as **TragicAgent**, **DepartingAgent**, etc. These specific agent types are obtained by posing necessary and sufficient conditions on the type of actions they perform, by exploiting the relation between the **Agent** class and the **Dynamics** class. Figure 3 shows an overview of the Entity taxonomy.

5 Applying Narrative Properties to Artwork Representation

The representation of stories and actions in the ontology is functional to the description of the narrative elements of the artefacts incorporated in the system. In this paragraph, we show how stories and actions are represented in the ontology and how this information is applied to the description of the artefacts through a three-step process.

Let us consider, for the sake of simplicity, the example of the story "Ariadne and the thread", illustrated in Figure 4. It is characterised by the performance of an action (here, the action of giving a ball of thread) and by a set of characters (Ariadne and Theseus). In the ontology, this story corresponds to an individual ("AriadneAndTheThread") belonging to the **Story** class (line 2), and is described as follows:

- the *hasCharacter* property connects the story with its charactes, Ariadne and Theseus (lines 3–4)
- the *hasAction* property (line 5) connects the story with the actions composing it, i.e., "giving_the_thread" (see below the description of this action) ;
- the *partOf* property describes the story as a subpart of another story, "MinotaurStory" (line 6);
- *recalls* relates this story with the archetype of the labyrinth ("CnossosLabyrinth", instance of the **Labyrinth** class, line 7) and with the archetype of the hero (Thesues_Hero, instance of the **Hero** class, line 8); both **Hero** and **CnossosLabyrinth** are subclasses of the **Archetype** class;
- *hasTimePeriod* locates the narrated time into a time period ("Year2000BC", line 9).

```
1 <!-- http://www.di.unito.it/labyrinth#giving_the_thread -->
2
3 <owl:NamedIndividual rdf:about="&www;labyrinth#giving_the_thread">
4   <rdf:type rdf:resource="&www;labyrinth#Action"/>
5 </owl:NamedIndividual>
6
7 <!-- http://www.di.unito.it/labyrinth#GivingProcessSchema -->
8
9 <owl:NamedIndividual rdf:about="&www;labyrinth#GivingProcessSchema">
10  <rdf:type rdf:resource="&www;labyrinth#ProcessSchema"/>
11   <Frame rdf:datatype="&xsd;string">Giving</Frame>
12   <predicate rdf:datatype="&xsd;string">Giving</predicate>
13   <hasRole rdf:resource="&www;labyrinth#Donor"/>
14   <hasRole rdf:resource="&www;labyrinth#Recipient"/>
15   <hasRole rdf:resource="&www;labyrinth#Theme"/>
16 </owl:NamedIndividual>
```

■ **Figure 5** RDF description of the action of giving (a ball of thread) in the AO ontology.

Notice that value of the *hasAction* property of the story description corresponds to the URI "giving_the_thread" , an instance of the **Action** class (see Figure 5, lines 3–5). This action is described by an instance of the class **ProcessSchema** (subclass of the **SituationSchema** class), part of a design pattern [14] which represents an action as a process having a set of roles attached to it, filled by the characters involved in the process. The **ProcessSchema** class relates the action to a Framenet frame (here, "giving", line 11) and to its roles (here, "Donor", "Recipient" and "Theme", lines 13–15). The roles are filled by the agents and objects which play these roles in the action: "Ariadne" as the "Donor", "Theseus" as the "Recipient", and the "ball of thread" as the "Theme" (for the sake of brevity, role fillers not shown in the Figure). The DescriptionTemplate class has been taken from the Drammar ontology as well as the pattern of which it is part (further details are contained in [8]).

While the representation of stories and actions is assumed to be encoded top down in the ontological model, the representation of the artefacts is obtained through a three step process. First, an internalization phase (*internalization*) imports the encoded metadata information about an artefact into the ontology. Then, the connection (*mapping*) of the imported artefact with the ontological model is performed. Finally, narrative connections are established. All phases are accomplished via rules encoded in SWRL format, in order to guarantee the portability of the system. The internalization process is as follows: when a new resource is added to the system, the ontological base is updated: a new individual is created to represent the artwork, and a set of assertions describing it are added to the ontological base. The system assumes that the resources given as input are described according to the Dublin Core (DC) metadata schema, encoded in RDF/XML.

As an example of the internalization phase, consider the "creator" DC element, which usually contains the reference to the artist who created a given artwork. The internalization procedure searches the ontological base for an author having the same name as the one contained in the creator element. If an instance representing the author of the artwork is not found, a new instance is created, with the *name* data property set to the value of the creator descriptor. Then, the artwork is connected to its author through the appropriate property (*hasCreator*). Similar rules are applied to all description elements. The internalization mechanism has a strategic importance, since it extends the task of populating the ontology to non–experts, by allowing them to annotate the information about the artefact in the Dublin Core format, and leaving to the system the task of encoding it into the ontology.

```
1  <!-- http://www.di.unito.it/labyrinth#
               Arianna_dona_il_gomitolo_di_filo_a_Teseo -->
2
3  <owl:NamedIndividual
     rdf:about="&www;labyrinth#Arianna_dona_il_gomitolo_di_filo_a_Teseo">
4    <rdf:type rdf:resource="&www;labyrinth#Manifestation"/>
5    <hasResourceType rdf:resource="&www;labyrinth#Image"/>
6    <evokes rdf:resource="&www;labyrinth#CnossosLabyrinth"/>
7    <evokes rdf:resource="&www;labyrinth#Theseus_Hero"/>
8    <displays rdf:resource="&www;labyrinth#Ariadne/>
9    <displays rdf:resource="&www;labyrinth#Theseus"/>
10   <hasGeographicalLocation rdf:resource="&www;labyrinth#Italy"/>
11   <describesAction rdf:resource="&www;labyrinth#giving_the_thread"/>
12   <ma-ont:hasCreator rdf:resource="&www;labyrinth#Pelagio_Palagi"/>
13   <ma-ont:hasPublisher
        rdf:resource="&www;labyrinth#Pinacoteca_Nazionale_di_Bologna"/>
14 </owl:NamedIndividual>
```

▪ **Figure 6** Description of an artwork "Arianna dona il gomitolo a Teseo" ("Ariadne gives Theseus a ball of thread") in the ontology after the internalization and mapping phase.

In the mapping phase, the metadata of the internalized resource are searched for keywords associated with the components of each archetype. For example, the title of the resource is searched for characters' names, which often occur in the title of artworks. Temporal and geographic information is sought for in "date" and "coverage". As an example of these rules, consider the rule that examines the keywords contained in the elements of type "subject" (which have already been internalized in the previous step) in order to find a connection with the archetype of the labyrinth: if the "labyrinth" or "maze" keywords are found in the subject descriptors of the artwork, the rule is applied and the artwork will be connected with the archetype of the labyrinth via the *evokes* property.

The internalization and mapping phases output a description of the artefact where the metadata have been internalized (such as creator, subject, etc.) and the connection with archetypes has been established. For example, consider the artwork represented as in Figure 6. In this case, the representation of the Italian painting "Arianna dona il gomitolo a Teseo" ("Ariadne gives Theseus a ball of thread"), by Pelagio Pelagi (1814) describes the artwork as an image, whose author is the painter Pelagio Pelagi, and whose subject displays two characters, Adriadne and Theseus. The painting is represented in the ontology as an instance of the **Manifestation** class (Figure 6, lines 3–4). This instance has several properties that relate it with other individuals in the ontology:

- *hasResourceType* describes the media type (image) in line 5;
- *evokes* connects the painting with the archetypes of the Labyrinth and of the Hero, lines 6–7;
- *displays* connects the painting with each entity which appears in it, i.e., Ariadne and Theseus (lines 8–9);
- *hasGeographicalLocation* relates the painting to the place where it is located ("Italy", line 10);
- *describesAction* relates the painting with the action type it represents ("giving_the_thread", line 11);

```
1  <!-- http://www.di.unito.it/labyrinth#
                Arianna_dona_il_gomitolo_di_filo_a_Teseo -->
2
3  <owl:NamedIndividual
     rdf:about="&www;labyrinth#Arianna_dona_il_gomitolo_di_filo_a_Teseo">
4    <rdf:type rdf:resource="&www;labyrinth#Manifestation"/>
5    <rdf:type rdf:resource="&www;labyrinth#artefact"/
6    <hasResourceType rdf:resource="&www;labyrinth#Image"/>
7    <evokes rdf:resource="&www;labyrinth#CnossosLabyrinth"/>
8    <evokes rdf:resource="&www;labyrinth#Theseus_Hero"/>
9    <displays rdf:resource="&www;labyrinth#Ariadne/>
10   <displays rdf:resource="&www;labyrinth#Theseus"/>
11   <hasGeographicalLocation rdf:resource="&www;labyrinth#Italy"/>
12   <describesAction rdf:resource="&www;labyrinth#giving_the_thread"/>
13   <describesAction rdf:resource="&www;labyrinth#Killing_the_Minotaur"/>
14   <describesAction rdf:resource="&www;labyrinth#Fighting"/>
15   <ma-ont:hasCreator rdf:resource="&www;labyrinth#Pelagio_Palagi"/>
16   <ma-ont:hasPublisher
       rdf:resource="&www;labyrinth#Pinacoteca_Nazionale_di_Bologna"/>
17   <hasPart rdf:resource="&www;labyrinth#AriadneAndTheThread"/>
18   <hasPart rdf:resource="&www;labyrinth#MinotaurStory"/>
19 </owl:NamedIndividual>
```

Figure 7 Description of the artwork "Arianna dona il gomitolo a Teseo" ("Ariadne gives Theseus a ball of thread") in the AO ontology, with properties added by the reasoner.

- *ma-ont:hasCreator*, taken from the Media Ontology, connects the painting with its author, Pelagio Pelagi (line 12);

- *ma-ont:hasPublisher* connects the painting with the institution by which it has been put online (line 13);

Upon the internalized and mapped representation of the artwork, a set of narrative rules in then applied. These rules consider the stories associated with the archetypes and check if their characters and actions are referred to by the artefact. By doing so, narrative associations are inferred only after the archetype association is performed, as in a classical cascade model.

For example, let us consider now the story "AriadneAndTheThread" represented in Figure 4: this story has Theseus, Ariadne, and the Minotaur as characters, and includes the action of giving the thread to Theseus by Ariadne. Since the artefact "Arianna da' il gomitolo a Teseo" ("Ariadne gives Theseus a ball of thread") shares with the "AriadneAndTheThread" story both the characters and the action, then the latter story is recognized to be a narrative "part of" the artwork (see Figure 7, line 17). The larger story, "MinotaurStory" (of which the "AriadneAndTheThread" story is a part), is also recognised as a part of the artwork (line 18). Therefore an "augmented" representation of the artefact with a narrative information is obtained (as shown in the Figure 7), which includes, via narrative rules, also the connection to the actions contained in the related stories (lines 12 and 13): "giving_the_thread" (from "AriadneAndTheThread") and "Killing_the_Minotaur" (from the "MinotaurStory").

6 Using narrative inferences in navigation

After the phases of internalization and mapping, "AriadneAndTheThread" and the application of narrative rules, the system performs several inferences on the artefacts represented in the ontology base. These inferences are exploited by the Labyrinth system for two main goals: on the one side, allowing the user to navigate among the artworks by following the semantic relations encoded in the ontology, even when they are not explicitly stated but inferred; on the other side, to classify the artefacts (and the other entities referred by it) with respect to the model encoded in the ontology, so that possible, alternative perspectives emerge on the artwork and can be proposed to the user's conceptualisation.

As an example of the first type of inference, consider the story termed "Ariadne and the ball of thread", described in Section 5 (Figure 4). The description of the artwork used as example (i.e., the painting "Arianna dona il gomitolo di filo a Teseo", Figure 7) states that the painting refers to this story. However, the painting also refers to the more general "Story of Ariadne" (**StoryOfAriadne**), encoded in the ontology as the class of stories having Ariadne as character (in OWL terms, the necessary and sufficient condition attached to the class is: *Story and (hasCharacter value Arianna)*): the fact that the painting also refers to the "Story of Ariadne", although not stated in the description, is automatically inferred by the system given the description of the story "Ariadne and the ball of thread". So, when the user is presented with the painting, the accompanying information is augmented with the inferred relations with more general stories: beside the "Story of Ariadne", for example, the system encompasses other similar classes, for example "Mythological Story", etc.

By the same reasoning type, other stories having Ariadne as character will be recognised as belonging to the **StoryOfAriadne** class: for example, the story termed "Bacchus and Ariadne" is also classified as belonging to this class, since it features Ariadne as character. The system uses this information for generating links between the artworks and for generating recommendations of similar contents (a painting which represents Bacchus and Ariadne, for example, or the libretto of the work "Ariadne auf Naxos" by Richard Strauss).

Useful inferences for the conceptualisation of stories also concern the relation between story and actions. In our model, stories can be subpart of other stories and are related to the actions they contain via a specialisation of the hasPart property (as described in Section 4). So, for example, the story "Ariadne and the ball of thread" is not only a specific type of story ("story of Ariadne") but also a subpart of the "story of the Minotaur". This mereological relations are supposed to hold also among the dynamics (e.g., actions, processes etc.) occourring in a story as well as among stories. In this way, it is possible to model the fact that some actions are part of other, more general, actions; that some actions are part of certain stories; that stories can be part of other stories. As a consequence of this modeling solution, higher-level stories (e.g., the "story of the Minotaur") are automatically connected with the actions which appear in the sub stories which are part of it. So, for example, the "story of the Minotaur" will be inferred as containing all the actions which appear in its sub stories. Therefore, from the description of "Ariadne and the ball of thread", it will be inferred as containing the action of giving, where Ariadne gives the ball of thread to Theseus; from the description of the story "Theseus and the Minotaur", it will acquire the action of killing, where Theseus kills the Minotaur and so on.

Thanks to these inferences, by gathering smaller stories (or episodes) in larger stories, actions are gathered as well, without the need, for the encoder, to explicitly track the relation between the actions in the episodes and the actions in the larger story. The navigation among the different artefacts through inferred actions and stories represents one of the innovative

Figure 8 Example of the narrative based semantic exploration.

aspects of our proposal. In cultural heritage, this is a benefit for user access, because users may not know how the single episodes are linked to each other as part of the larger story: by exploring the repository, they become aware of the relation between the episodes through the artworks.

An example of this narrative, semantic driven, navigation among the artifacts is given in Figure 8. In this case, we hypothesize that the user is exploring the Actions related to the archetype of the Hero. For the action "Killing", a set of narrative information is recovered from the ontological model, using both the explicit and the inferred relations. Namely: information about the possible classifications of the selected actions ("Hero Action", "KillingAction" or, more in general, "Event"), about the stories related to that action (e.g., the "Minotaur Story", the story of "Theseus killing the Minotaur" – subpart of the previous one – and the story of "Hercules and the Lion") and about some characters involved, with different roles, in that action (e.g., Marat, Minotaur, Hercules and Theseus). Finally some artefacts are also obtained, which represent directly, as in the case of the painting called "Death of Marat", or indirecty, as in the painting "Minotauromachia" by Picasso, the action of killing: by doing so, the system enhances the semantic grouping and retrieval of information in a way that is not possible with classical relational databases.

7   Conclusions and Future Work

Formal ontologies can be a powerful tool for intelligent information systems aimed at improving the narrative navigation of digital artefacts. Beside the narrative model explicitly encoded in the ontology, automatic reasoning processes provide useful insights on the relations connecting different media resources having, by and large, narrative content. In cultural heritage, this amounts to providing the users with a conceptual framework – stories – they are familiar with, open to heterogenous contents over media and ages.

In this paper, we presented Labyrinth, an ontology based system designed for the enhancement of the semantic exploration of digital media archives. Given the information encoded in the ontology, the user can explore a repository based on the relations that link the repository contents (i.e., a set of media resources) with a given archetype, which has a prominently narrative nature. Thanks to the representation of the complex interplay of the concepts of story, characters and actions in the ontology, the use of reasoning tools lets a set of relations emerge among the resources in the archive, thus contributing to the user involvement in the exploration experience.

The system is currently being developed and was tested on a small corpus of resources of different type and format, in order to asses the functioning and feasibility of the approach.

As future work, we envisage the validation of the ontology on a larger set of media resources, and the testing of the proposed representation and inferences on real users, according to the paradigm of user studies.

References

1 D. G. Aliaga, E. Bertino, and S. Valtolina. DECHO—a framework for the digital exploration of cultural heritage objects. *Journal on Computing and Cultural Heritage*, 3(3):12, 2011.

2 C. F. Baker, C. J. Fillmore, and J. B. Lowe. The Berkeley FrameNet project. In Christian Boitet and Pete Whitelock, editors, *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics*, pages 86–90. Association for Computational Linguistics, 1998.

3 J. D. Bolter and R. Grusin. *Remediation: Understanding new media.* MIT Press, 2000.

4 J. Bruner. The narrative construction of reality. *Critical Inquiry*, 18(1):1–21, 1991.

5 I. Bywater. *Aristotle, On the Art of Poetry.* Clarendon, Oxford, 1920.

6 J. Campbell. *The hero with a thousand faces.* New World Library, 2008.

7 N. Carroll. *Beyond Aesthetics: Philosophical Essays.* Cambridge University Press, Cambridge, 2001.

8 M. Cataldi, R. Damiano, V. Lombardo, and A. Pizzo. Lexical mediation for ontology-based annotation of multimedia. In *New Trends of Research in Ontologies and Lexical Resources*, Theory and Applications of Natural Language Processing, pages 113–134. Springer, 2013.

9 M. Cataldi, R. Damiano, V. Lombardo, A. Pizzo, and D. Sergi. Integrating commonsense knowledge into the semantic annotation of narrative media objects. In R. Pirrone and F. Sorbello, editors, *AI*IA 2011: Artificial Intelligence Around Man and Beyond, XIIth International Conference of the Italian Association for Artificial Intelligence, Palermo, Italy, September 15–17, 2011. Proceedings*, number 6934 in Lecture Notes in Computer Science, pages 312–323. Springer, 2011.

10 J. Cua, E. Ong, R. Manurung, and A. Pease. Representing story plans in SUMO. In Paul Cook and Anna Feldman, editors, *Proceedings of the NAACL HLT 2010 Second Workshop on Computational Approaches to Linguistic Creativity*, pages 40–48. Association for Computational Linguistics, 2010.

11 G. De Melo, F. Suchanek, and A. Pease. Integrating YAGO into the Suggested Upper Merged Ontology. In *20th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2008), November 3–5, 2008, Dayton, Ohio, USA, Volume 1*, pages 190–193. IEEE, 2008.

12 M. Doerr. The CIDOC conceptual reference module: an ontological approach to semantic interoperability of metadata. *AI Magazine*, 24(3):75, 2003.

13   C. R. Fairclough. *Story Games and the OPIATE System: Using Case-Based Planning for Structuring Plots with an Expert Story Director Agent and Enacting them in a Socially Simulated Game World.* PhD thesis, University of Dublin, Trinity College, 2004.

14   A. Gangemi and V. Presutti. Ontology design patterns. In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, pages 221–243. Springer, 2009.

15   P. Gervás, B. Díaz-Agudo, F. Peinado, and R. Hervás. Story plot generation based on CBR. *Knowledge-Based Systems*, 18(4):235–242, 2005.

16   A. Giovannelli. In sympathy with narrative character. *Journal of Aesthetics and Art Criticism*, 67(1):83–95, 2009.

17   A. Grund. ICONCLASS: On subject analysis of iconographic representations of works of art. *Knowledge Organization*, 20(1):20–29, 1993.

18   K. Hartmann, S. Hartmann, and M. Feustel. Motif definition and classification to structure non-linear plots and to control the narrative flow in interactive dramas. In G. Subsol, editor, *Virtual Storytelling. Using Virtual Reality Technologies for Storytelling Third International Conference, ICVS 2005, Strasbourg, France, November 30 – December 2, 2005. Proceedings*, number 3805 in Lecture Notes in Computer Science, pages 158–167. Springer, 2005.

19   G. Highet. *The classical tradition: Greek and Roman influences on Western literature.* Oxford University Press, USA, 1949.

20   E. Hyvönen, E. Mäkelä, T. Kauppinen, O. Alm, J. Kurki, T. Ruotsalo, K. Seppälä, J. Takala, K. Puputti, H. Kuittinen, K. Viljanen, J. Tuominen, T. Palonen, M. Frosterus, R. Sinkkilä, P. Paakkarinen, J. Laitio, and K. Nyberg. CultureSampo: A national publication system of cultural heritage on the Semantic Web 2.0. In L. Aroyo, P. Traverso, F. Ciravegna, P. Cimiano, T. Heath, E. Hyvönen, R. Mizoguchi, E. Oren, M. Sabou, and E. Simperl, editors, *The Semantic Web: Research and Applications, 6th European Semantic Web Conference, ESWC 2009, Heraklion, Crete, Greece, May 31–June 4, 2009, Proceedings*, volume 5554 of *Lecture Notes in Computer Science*, pages 851–856. Springer, 2009.

21   Eero Hyvönen, Eetu Mäkelä, Tomi Kauppinen, Olli Alm, Jussi Kurki, Tuukka Ruotsalo, Katri Seppälä, Joeli Takala, Kimmo Puputti, Heini Kuittinen, Kim Viljanen, Jouni Tuominen, Tuomas Palonen, Matias Frosterus, Reetta Sinkkilä, Panu Paakkarinen, Joonas Laitio, and Katariina Nyberg. CultureSampo—Finnish culture on the semantic web 2.0. thematic perspectives for the end-user. In J. Trant and D. Bearman, editors, *Museums and the Web 2009: Proceedings*, pages 15–18. Archives & Museum Informatics, 2009.

22   H. Jenkins. *Convergence Culture: Where Old and New Media Collide.* NYU Press, 2008.

23   M. O. Jewell, K. F. Lawrence, M. M. Tuffield, A. Prugel-Bennett, D. E. Millard, M. S. Nixon, M. C. Schraefel, and N. R. Shadbolt. OntoMedia: An ontology for the representation of heterogeneous media. In *Multimedia Information Retrieval Workshop, MMIR 2005, Salvador, Brazil, 19 August 2005*, 2005.

24   K. F. Lawrence. Crowdsourcing linked data from Shakespeare to Dr Who. Poster at the Web Science International Conference, 2011.

25   V. Lombardo and R. Damiano. Semantic annotation of narrative media objects. *Multimedia Tools and Applications*, 59(2):407–439, 2012.

26   C. Masolo, S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. WonderWeb deliverable D18: Ontology Library, 2009. Final report for IST Project 2001-33052 WonderWeb: Ontology Infrastructure for the Semantic Web.

27   G. A. Miller. WordNet: a lexical database for English. *Communications of the ACM*, 38(11):39–41, 1995.

28   P. Mulholland and T. Collins. Using digital narratives to support the collaborative learning and exploration of cultural heritage. In *13th International Workshop on Database and Expert Systems Applications (DEXA 2002), 2–6 September 2002, Aix-en-Provence, France*, pages 527–531. IEEE, 2002.

29   P. Mulholland, T. Collins, and Z. Zdráhal. Story fountain: intelligent support for story research and exploration. In N. Jardim Nunes J. Vanderdonckt and C. Rich, editors, *Proceedings of the 2004 International Conference on Intelligent User Interfaces, January 13–16, 2004, Funchal, Madeira, Portugal*, pages 62–69. Association for Computing Machinery, 2004.

30   E. T. O'Neill. FRBR: Functional requirements for bibliographic records. *Library resources & technical services*, 46(4):150–159, 2002.

31   G. Prince. *A Dictionary of Narratology.* University of Nebraska Press, Lincoln, NE, new edition, 2003.

32   Vladimir Yakovlevich Propp. *Morphology of the Folktale.* University of Texas Press, Austin, TX, 2nd edition, 1968. transl. L. Scott.

33   J. Trant and B. Wyman. Investigating social tagging and folksonomy in art museums with steve.museum. Presented at the Collaborative Web Tagging Workshop at WWW2006, Edinburgh, Scotland, 2006.

Story Comparisons: Evidence from Film Reviews[*]

Bernhard Fisseni[1,2], Aadil Kurji[3], Deniz Sarikaya[1], und Mira Viehstädt[1]

1   Fachbereich Mathematik
    Universität Hamburg
    Hamburg, Germany
    `{deniz.sarikaya,mira.viehstaedt}@stud.uni-hamburg.de`
2   Institut für Germanistik
    Universität Duisburg-Essen
    Essen, Germany
    `bernhard.fisseni@uni-due.de`
3   Department of Philosophy
    University of Bristol
    Bristol, UK
    `ak12004@bristol.ac.uk`

Abstract

Interested in formally modelling similarity between narratives, we investigate judgements of similarity between narratives in a small corpus of film reviews and book–film comparisons. A main finding is that judgements tend to concern multiple levels of story representation at once. As these texts are pragmatically related to reception contexts, we find many references to reception quality and optimality. We conclude that current formal models of narrative can not capture the task of naturalistic narrative comparisons given in the analysed reviews, but that the development of models containing a more reception-oriented point of view will be necessary.

1   Background and Research question

We are interested in story similarity, which has been approached both from a cognitive point of view (e.g., [12, 3]) and from a formal point of view (e.g., [13, 5, 6]), but also is implicitly an important question for all formal approaches to narrative. Similarity is the natural subject of discussion in comparisons of narratives. We discuss preliminary experiences when analysing data from two practical applications of comparisons between narratives: comparisons between books and their adaptations on the one hand, and reviews that contain remarks that compare a certain film with other films, where different relationships obtain between the films, mainly: remake and just alleged similarity. Although we cannot claim that our data are sampled to be representative, we think that the points we make are exemplary and can plausibly be generalised to other data of this kind.

---

2   Investigating Narrative similarity

When we want to compare stories formally, we need at least the following two components: (a) **a representation of the narratives** that captures the relevant aspects, and (b) **a metric** to determine the 'distance' between two representations.

Determining the metric and defining a framework are interdependent decisions. For both, the most important step is to determine which information enters the comparison, as for every formal framework like Plot Units [12] or the Doxastic Preference Framework [14], only certain information goes into the representation, and hence into the comparison. To take Plot Units as an easy, but relatively representative target, only such events that involve a 'mental state' and ultimately an emotional change can be part of the model of a story; hypothetical events, or non-anthropomorphised non-emotional events (e.g., volcano eruptions) etc. are not part of the formal model, only potential emotional or motivational consequences. With our analysis, we want to identify the aspects of narrative important for comparisons.

Data and Procedure.   We intended to collect 'natural' data that (a) were in the domain of story comparison, (b) but were not prepared in an academic context, because we wanted to avoid 'self-observation' on our side. Data was collected between 5 and 12 December 2012 and consists of two parts, each with about 100 extracts.[1] First, a collection of 25 texts in German, which compare books and films (17 compared a book with 'its' film, the rest considered films and a theatre play). The data compare narratives in a rather informal, non-academic context.[2] These were collected on the internet using simple keyword search, and include texts from forums (like: `http://www.dvd-forum.at`, `http://de.answers.yahoo.com`, overall 11), blogs (5), semi-professional review sites (like: `http://www.moviepilot.de`, `http://www.negativ-film.de`), from Wikipedia (1) and a fan wiki. These data are generally not from professional writers.

The second part of our data consists of comparisons of films in English and German, which were found starting from IMDB. The texts were generally from film review websites, and the authors can be assumed to have at least a semi-professional background.[3]

Mainly to facilitate the qualitative evaluation and 'get a feeling for the data', we annotated the data. We extracted spans (about 200, of 3 to 180 words, with a median of 30) of texts and annotated these with tags corresponding to the levels Schmid's 'ideal-genetic model' [17] (which is a four-level model of narration, adding a level 'below' the story and between story and discourse), adding a level for the story world and pragmatic effects of texts, and noting interactions between levels. Regarding interaction, consider (5-b) and (5-c). (5-b) contains two relatively unrelated claims (new language, speakability). (5-c) illustrates the more common case: Several levels are treated at once and are seen as interrelated: What is talked about is the story world (and its presentation), but also about the aesthetic effect and entertainment value of these.

---

[1] Analysis is still ongoing and data occasionally corrected, so counts should be taken with a grain of salt.

[2] To retain non-academic setting, we excluded 3 texts from the analysis because they were explicitly tagged as (academic or school) 'homework', but kept a *Schülertext* ('pupils' text') from a newspaper by 8th-formers discussing two film versions of *Pride and Prejudice.*

[3] (a) *Avatar* (2009), *Pocahontas* ('myth' and film), *Dances with Wolves* (1990) (of the first 40 reviews on `http://www.imdb.com`, 32 contained comparisons and were accessible with respect to network and language) (b) *West Side Story* (1961, dir. Robert Wise) and *Romeo and Juliet* (play and film versions). (c) *Infernal Affairs* (2002) / *The Departed* (2006) (4); and (d) *Abre los Ojos* (1997) / *Vanilla Sky* (2001) (5). In the last two, the second film was a remake of the first, and also the relationship between *West Side Story* and *Romeo and Juliet* is obvious.

Which levels of the story are accessed in the reviews?   Current formal frameworks of narrative model the intradiegetic [7] level of narration ('What happens?'), and if they are computational, often also taking into account the story world (characters, entities and the relationships between them, as far as relevant to the story) [9]. We find that in the reviews and comparison, some comparisons are on this level, but then often goes 'deeper' towards abstraction of patterns or a metaphoric or allegoric interpretation as in (1), or just themes the author of the reviews sees but which belong to a 'deeper' level of interpretation.[4]

(1) [*Gone with the Wind*, book–film] *Im Buch geht es um den Untergang einer Gesellschaft, um ein untypisches Bild der Sklaverei, um einen Krieg, um den Überlebswillen, um den Werdegang einer Frau. Wo hingegen der Film eher (nur) das Frauenbild der 30er Jahre behandelt.*
The book is about the demise of a society, an untypical depiction of slavery, a war, the will to survive, the development of a woman. While the movie rather (only) treats the image [society had] of women in the 30s.
http://www.hochzeitsplaza.de/hochzeits-forum/off-topic/off-topic/87701-vom-winde-verweht-ist-hier-noch-jemand-ein-scarlett-o-hara-fan/

A level that is present in most comments (145) in our little corpus is the perspective of the recipient. Comments referring to the reception of the story are by far the most frequent (out of 83 extracts somehow referencing what happens, 67 also relate it to reception). Surprisingly sometimes, the plot level is contrasted not only with a presentation/discourse level, but the 'same' story is not the same, as in (2).

(2) [Regarding Grenouille in *Perfume*, book/film] *Zwar deutet das Filmende eine andere Motivation für seinen schlussendlichen Selbstmord an [. . .] aber im Handlungsablauf ist der Roman an dieser Stelle eigentlich ziemlich genau umgesetzt worden.*
The ending of the film indicates a different motivation for his final suicide [. . .] but regarding the course of action, the novel was converted quite faithfully.
http://www.gutefrage.net/frage/das-parfum---wesentliche-unterschiede-zwischen-buch--film-

The story level is also presented as a means towards the reconstruction of the story world (15 times clearly, more often in allusions that need further analysis), which is what cognitivist narratologists take the interpretation of the intradiegetic level to be (e.g., [11]), either its causal connections, its plausibility or its aesthetics. Comments explaining why a certain scene is important and what it contributes to the understanding of the story world such as (3) come closest to an intradiegetic metric of similarity. Except for one of these comparisons are from book–film comparisons; this is plausible. Intuitively, such observations need a very great similarity: It is moot to observe such things, e.g., about films that are just vaguely similar. But note that a similar role is also ascribed to discourse/presentation features and other levels (13 out of 20 references) as in (4). Such interaction between levels is well known to translators, of course; compare, e.g., Dusi's remark [2, p. 9] that one needs to translate not only semiotic levels, but also the relations between them.

(3) a. [*The Reader*: book/film] *Diese Stelle ist insofern wichtig, da hier das erste mal deutlich gezeigt wird, dass Hanna Analphabetist ist und zu welchen Problemen dies führen kann.*
This passage is important because it is shown clearly for the first time that Hanna is illiterate and which problens result from this. *http://www.hundertachtzehn.com/603/kritik-der-vorleser-der-film.html*
b. [*Harry Potter 7*] Once they get there [Sirius Black's house], the film leaves out several pieces of the book. Firstly, Harry, Ron and Hermione hear an interesting story from Kreacher the house elf that is left out of the film. After this, Kreacher becomes much kinder to the three young wizards. This story includes some relevant information about Sirius' brother.
http://www.bukisa.com/articles/399514__harry-potter-and-the-deathly-hallows-differences-between-the-book-and-the-movie#ixzz2LHVtLY2T

---

[4] We do not want to judge the adequacy of these analyses!

(4) *Im Buch ist die Sprache außerdem viel komplexer, Grenouilles Innenwelt wird einem viel klarer und Auffälligkeiten sind teilweise subtil, während der Film viel mit Erschrecken arbeitet.*
In the book, the language is much more complex, Grenouille's interior world is presented much more articulately, and abnormities are partly subtle, while the film works much more with shocking the reader.

http://de.answers.yahoo.com/question/index?qid=20061121043445AAiQ5Zr

Task-Relatedness. It is immediately plausible that the pragmatic goal of the comparison has an influence on the structure of the comparison. Film reviews are inherently concerned with the effect on the audience, and therefore reception-oriented comments are a natural level of judgement. From our results, we assume that, although arguably more 'natural' than story analogy experiments as those reported by Gentner et al. [8], which show that subjects prefer (diegetically-oriented) relational mapping as a measure for similarity, our data are less well-suited to investigate the intradiegetic level. We do find analogy mappings for films or books on almost any level, e.g. on the reception-oriented level; e.g., (5-c) refers to the concept of *alienness* and how it is realised in different films.

3   Preliminary Conclusions

The conclusion we have to draw is such that there is a gap between what current formal models of narrative can analyse and the tasks brought about by naturalistic narrative comparisons. To find a formal model that is well-suited to the analysis task we have to research, and develop, systems which are more reception-oriented. Without extensions and connection to other levels and 'deeper interpretations' relating to the actual world of the recipients, current formal models of narrative cannot inform such comparisons in a substantial way, even though they are certainly quite natural cases of story comparisons. (This criticism does not affect the usefulness of current formal models for computer games or other 'simple' applications where the recipient becomes part of the story world, or where aesthetic criteria are unimportant, as in some retrieval tasks.) Our data provides evidence that besides intradiegetic models, we also need models of the extradiegetic and reception-oriented aspects of narrative. While this is not a new suggestion [10, 1], only few systems seem to take up the idea (cf. [15, §2.4.3], where integrating a user model is discussed, but references are limited to suspense generation). The current trend towards machine-learning (e.g., [15, 4, 16]) tends to shift the attention away from such high-level tasks.

For the future, we plan to extend the data analysis; ultimately we want to integrate reception-oriented criteria into formal frameworks in the hope to approximate an adequate analysis of naturalistic narrative comparisons. We agree with [10, 7] that an 'interdisciplinary' approach is needed, i.e., formal approaches which incorporate more than just the story level of analysis, merging the extradiegetic with intradiegetic analysis. For example, a formal epistemic framework which models not only what is known, and when, by characters, but also the change of knowledge (and beliefs) of narrator(s) and reader(s), and the effects of these on the latter.

References

1 Paul Bailey. Searching for storiness: Story-generation from a reader's perspective. In M. Mateas and P. Sengers, editors, *Narrative Intelligence: Papers from the AAAI Fall Symposium*, number FS-99-01 in AAAI Technical Reports. AAAI Press, 1999.

2 Nicola Dusi. *Il cinema come traduzione.* UTET Università, Torino, 2003.

3 Brian Falkenhainer, Kenneth Forbus, and Dedre Gentner. The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 20:1–63, 1989.

4 Mark Alan Finlayson. *Learning Narrative Structure from Annotated Folktales.* PhD thesis, Massachusetts Institute of Technology, 2011.

5 Bernhard Fisseni and Benedikt Löwe. Which dimensions of narratives are relevant for human judgments of story equivalence? In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 114–118, İstanbul, 2012.

6 Bernhard Fisseni and Benedikt Löwe. Event-mappings for comparing frameworks for narratives, submitted.

7 Gérard Genette. *Discours du récit.* Points, Paris, 2007.

8 Dedre Gentner, Mary Jo Rattermann, and Kenneth D. Forbus. The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, 25:524–575, 1993.

9 Pablo Gervás. Story generator algorithms. In Peter Hühn, Jan Christoph Meister, Wilhelm Schernus, John Pier, Wolf Schmid, and Jörg Schönert, editors, *The Living Handbook of Narratology.* Hamburg University Press, 2012. version 6 May 2013. URL: `http://hup.sub.uni-hamburg.de/lhn/index.php/Story_Generator_Algorithms`.

10 Pablo Gervás, Birte Lönneker-Rodman, Jan Christoph Meister, and Federico Peinado. Narrative models: Narratology meets artificial intelligence. In Roberto Basili and Alessandro Lenci, editors, *Workshop Toward Computational Models of Literary Analysis, May 22nd, 2006, Genoa (Italy)*, pages 44–51, 2006.

11 David Herman. *Basic Elements of Narrative.* Wiley-Blackwell, Malden, MA, 2009.

12 Wendy G. Lehnert. Plot units and narrative summarization. *Cognitive Science*, 5(4):293–331, 1981.

13 Benedikt Löwe. Methodological remarks about comparing formal frameworks for narratives. In Patrick Allo and Giuseppe Primiero, editors, *Third Workshop in the Philosophy of Information, Contactforum van de Koninklijke Vlaamse Academie van België voor Wetenschappen en Kunsten*, pages 10–28, Brussel, 2011. KVAB.

14 Benedikt Löwe, Eric Pacuit, and Sanchit Saraf. Identifying the structure of a narrative via an agent-based logic of preferences and beliefs: Formalizations of episodes from CSI: Crime Scene Investigation™. In Michael Duvigneau and Daniel Moldt, editors, *Proceedings of the Fifth International Workshop on Modelling of Objects, Components and Agents. MOCA'09*, number 290/09 in FBI-HH-B, pages 45–63, 2009.

15 Inderjeet Mani. *Computational Modeling of Narrative.* Number 5 in Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2012.

16 Neil McIntyre. *Learning to Tell Tales: Automatic Story Generation from Corpora.* PhD thesis, University of Edinburgh, 2011.

17 Wolf Schmid. *Elemente der Narratologie.* de Gruyter, Berlin, 2nd edition, 2008.

A   Additional Data

(5) [*Avatar*]

    a. Avatar is a simple story of war versus peace, human versus alien, a modern species versus an indigenous tribe. It draws influences from numerous films (apart from Cameron's own), most notably Kevin Costner's Dances With Wolves (1990), Edward Zwick's The Last Samurai (2003), and Hayao Miyazaki's Princess Mononoke(1997) and Castle In The Sky (1986).
       `http://filmnomenon2.blogspot.de/2009/12/avatar-2009.html`

    b. It invents a new language, Na'vi, as Lord of the Rings did, although mercifully I doubt this one can be spoken by humans, even teenage humans.
       `http://rogerebert.suntimes.com/apps/pbcs.dll/article?AID=/20091223/LETTERS/912239997`

    c. More importantly, Cameron gives us an alien world in the true sense of the meaning. In most science-fiction films and television – the various Star Trek series being particularly guilty offenders – alienness never amounts to anything more than extras with a few funny facial appliances covering their noses and foreheads. In these there is frustratingly little effort made to conceive of something that is truly alien and goes beyond the standard human-like anthropomorphism.
       `http://0to5stars-moria.ca/sciencefiction/avatar-2009.htm`

A Paradigm for Eliciting Story Variation*

Bernhard Fisseni[1,2] and Faith Lawrence[3]

1   Fachbereich Mathematik
     Universität Hamburg
     Fachbereich Mathematik
     Hamburg, Germany
2   Institut für Germanistik
     Universität Duisburg-Essen
     Essen, Germany
     `bernhard.fisseni@uni-due.de`
3   Digital Humanities
     King's College London
     London, UK
     `faith.lawrence@kcl.ac.uk`

──── **Abstract** ────

The understanding of story variation, whether motivated by cultural currents or other factors, is important for applications of formal models of narrative such as story generation or story retrieval. We present the first stage of an experiment to elicit natural narrative variation data suitable for evaluation with respect to story similarity, to qualitative and quantitative analysis of story variation, and also for data processing. We also present few preliminary results from the first stage of the experiment, using *Red Riding Hood* and *Romeo and Juliet* as base texts.

1   Introduction

Adaption, re-adaption and remakes have become a common part of popular culture [11]. Between this and the growing mainstreaming of remixed and transformative works [10], the co-existence of multiple variations of a story is becoming[1] widely accepted [9, 19]. Story[2] variation, or distinctness, and stahlory similarity are two aspects of the same question [12, 6, 13, 5]: When does a narrative cease being a version of a story but is instead seen

---

[1]   Not for the first time if we take oral traditions and folklore into account.
[2]   In this paper *story* and *narrative* are often used interchangeably, although *story* properly refers to the 'greater' culturally situated artefact rather than the material manifestation for which we prefer *narrative*. Where we wish to refer to the semantic level of the text for which the term *story* (in the sense of Todorov's *histoire*) is often used in computational literature, we have used *plot*. This was done to reflect that *story* seemed to be the most appropriate term to use in communication with test subjects. That *story/narrative* are understood in a broader sense is important because in a too narrow understanding, the question may seem appropriate: Isn't a narrative already different if we change a comma? We have not drawn a distinction between *version* and *variation* for similar reasons.

as a story in its own right [18]? An underlying question is: Do the levels of plot (*story* in Todorov's sense) and discourse interact, or are they independent?

This issue is highly relevant for story generation and story retrieval as well as any other practical applications of formal models of narrative [1, 2, 14, 15, 4]. Selecting test material poses significant challenges: While the textual corpus of the world is extensive, such data is generally too diverse to allow an easy comparison with respect to research questions, and frequently too complex to be analysed in depth. Conversely, working with texts created for the purpose of experimentation and analysis allows control over variables such as length, complexity and composition but has a high risk of losing the characteristics of 'natural' discourse found in organically created texts.

This paper introduces the first stages of an experiment designed to elicit story variations in English and German, to (a) **create a corpus** of related texts that are more suitable to the needs of researchers than what one finds 'in the wild', and subsequently (b) allow for and carry out **analysis**. We aim at the following kinds of analysis: (i) qualitative (and ultimately quantitative) 'human' corpus analysis for similarities and difference, of variations to the paired summary, to each other and to similar texts; also, (ii) annotation of the data allowing to test annotation systems and computational analysis tools on a known corpus of manageable size. By controlling specific axes of the product (length, story derivation, variation), but otherwise allowing the authors freedom in their process, we hope to ensure that the corpus is interesting and viable for researchers while retaining as much authenticity as possible. Ultimately, we expect that our data may become part of a 'story bank', often named as a desideratum in computational modelling of narrative (e.g., in recent CMN workshop announcements), and establish a methodology for expanding and analysing such a collection.

2   Experiment

A preliminary survey was carried out to explore detail recollection of three narratives: *Romeo and Juliet* (RJ), *Little Red Riding Hood* (LRR) and *Harry Potter: The Prisoner of Azkaban.* Volunteers in both Germany and the UK rated a series of potential events to indicate whether they occurred in the given text. Results indicated that (a) few details of the stories were retained and (b) it was unlikely to be productive if we asked for subtle variations that went beyond the main characters and plot line. It was decided to focus on RJ and LRR as well known works with comparatively simple primary plots. Although arguably both cautionary tales, they represent different basic narrative types: the fairy tale and the tragic love story.

The main experiment was divided into two stages, the first of which we detail in this paper; it involved the collection of the corpus of story and variation summaries. In the second stage of the experiment, we will elicit multi-dimensional similarity ratings for a selection of the collected data regarding similarity to the 'original' story, but also to other narratives presumably containing similar plots or motifs, such as the stone filling seen in both LRR ([7, #26], cf. fn. 6) and *The Wolf and the Seven Young Kids*, ([7, #5]).

Methodology.   Test subjects were invited to write a short summary (100–300 w) of their selected story. Once the base summary was submitted, the participant was given one of the variation constraints and asked to write a second text of around the same length taking this change into account.[3] The experiment was carried out online and volunteers were mainly

---

[3] It was made clear that the second text need not follow the structure of the first.

recruited from creative writing and other amateur author groups. While presented as a scientific experiment, volunteers were encouraged to to see it as a playful creative challenge (see Extract 1). They were allowed to complete the experiment at their own pace and were not prevented from refreshing their memories about the plot at any time.[4]

(1) **Instruction (short extract):**[5] This is an experiment to collect data on story variation. It does not require or test any particular level of intelligence, education or writing ability. We hope it will be fun and we are very grateful for your help.
For this experiment you will be asked to write a summarised version of *Romeo and Juliet* and/or *Little Red Riding Hood* under a constraint you will be given.

> *An example constraint which is not in the experiment might be*:
> all human characters are animals and vice versa.

For the experiment, six short instructions for writing variations ('constraints') were prepared for each narrative (see below). These constraints were chosen to include significant changes to the setup of the characters (LRR1, LRR2, LRR 5, RJ1) and their properties (LRR3, LRR4, LRR6, RJ2, RJ4, RJ5, RJ6) or introduced a narrative change, or 'twist', to the plot (LRR5, RJ2, RJ5). The choice of constraints was randomised and participants were given the option to turn a constraint down. Rejections were recorded and only three rejections in a row were possible. Once they had the variation summary for the assigned constraint, participants could continue with another variation, or change to the other story.

Little Red Riding Hood Constraints: (LRR1) The character of the wolf is not in the story. (LRR2) The character of the Huntsman/woodcutter is not in the story. (LRR3) Little Red Riding Hood's grandmother lives with Little Red Riding Hood and Little Red Riding Hood's family. (LRR4) One or both of the main characters (Little Red Riding Hood, Little Red Riding Hood's grandmother) are male. (LRR5) Little Red Riding Hood's grandmother died before the story starts. (LRR6) The main characters are political, geographical or commercial entities.

Romeo and Juliet Constraints: (RJ1) The character of Mercutio is not in the story. (RJ2) Romeo and Juliet are not in love with each other and are forced to marry each other against their will. (RJ3) The Capulets and the Montagues are good friends. (RJ4) Romeo falls in love with another character than Juliet or Rosalind. (RJ5) Juliet reveals her secret marriage with Romeo to her parents. (RJ6) The main characters are political, geographical or commercial entities.

Participants.   From the initial call for volunteers there were 32 responses, 8 in English and 24 in German (see table below); 10 test subjects submitted one summary, 4 submitted two, 1 and 2 five and six, respectively. Due to the differential between the number of responses in German and English, it was not deemed possible to draw any comparison between the two groups at this time. Further English volunteers are currently being sought.

		German		English	
		Male	Female	Male	Female
Little Red Riding Hood	Baseline Summary Only	1	3	1	
	Summary & Variations	2	8		4
Romeo and Juliet	Baseline Summary Only			1	1
	Summary & Variations	1	2	1	
No Summary		3	5		

(One volunteer submitted responses to both narratives so was included twice.)

---

[4] Some volunteers mentioned reading synopses of the story on Wikipedia.
[5] A German and an English version of the instructions was available; we only give English examples.

3   Observations

Due to the distribution of the responses received to date, it is not yet sensible to give an elaborate analysis. We cite some 'paradigmatic' examples, mainly focusing on LRR.

Usefulness of the Summary.   Collecting summaries, not only variations, is especially important in the case of LRR, as it is a story with many existing variants (see, e.g., [17, *Rotkäppchen*]).[6] As expected based on our preliminary survey, the baseline summaries do not completely agree with respect to the detail: Many test subjects do not mention the 'punishment' episode and the death of the wolf (*liberation*: 19y : 2n; *punishment*: 10y : 11n; *wolf dead*: 12y : 9n), and one of them has the liberation of the grandmother take place before LRR gets swallowed. We are not aware of a published variant that have liberation but not punishment, so that this may be an indication that the 'resurrection' is the more important scene (and furthermore it is plausible to assume that the wolf dies of a cut stomach).

Simple and Complex Solutions.   Comparing the story variations that we received, it was clear that the constraints prompted very differing levels of transformation. In addition to this, it was noticeable that some test subjects chose very simple solutions to the problem of integrating the proposed change, while others took the opportunity to change much more than 'necessary'. The instructions did not specify a preference for either solution, so it will be interesting to look into this further.

In the case of (LRR1), the absence of the wolf can result in a removal of the main story line (in one variant, explicitly 'nothing' happens) or in a simple exchange of the aggressor (once the grandmother, once a Bambiraptor), which keeps the main story line, but there are solutions in between (e.g., the grandmother beats LRR); similarly so for (LRR4), where some test subjects implement the sex change by simply exchanging the pronouns; one volunteer noted that this felt 'like cheating'. Others change the story completely: in one variant, the wolf character is exchanged for a beautiful wench (German: *"Maid"*), who seduces LRR, eating the cake intended for the grandmother and stealing the box of tools LRR was bringing to him; finally, she robs the grandmother of her money. Constraint (LRR6) and (RJ6) required the replacement of the main characters with commercial or geographic entities. Here the difficulty lay in: (a) signalling the character mapping to the readers – even though this was not demanded – and (b) giving analogues to eating and swallowing. For the first question, test subjects (4 out of 4, one only for LRR itself, sc. LRR as the German social democrats / 'reds') choose to use names that playfully point to the original characters, such as "Redhood Bank" [English], or "Lupuria" and "Omar" (in German, the latter more or less homophonous with the colloquial word for grandmother, *Oma*).

4   Preliminary Conclusions and Continuation

Based on our preliminary analysis, we conclude that the paradigm is suitable to elicit variations of texts from test subjects. However, due to the range of variation caution must be exercised in collation and due care taken in similarity judgements using multidimensional

---

[6] In Germany, the version collected by the Brothers Grimm [7, # 26] (or a variant of it, such as Bechstein's [3]) is the most popular. Compared to the earliest published version by Perrault [16] it also contains the liberation of LRR and her grandmother from the wolf's stomach, and the punishment of the wolf by filling his stomach with stones and his subsequent death (as in *The Wolf and the Seven Young [Goat] Kids*); it also lacks a 'moral'. For research on the relation between versions, see [17, *Rotkäppchen*]

rating. We also note that certain changes seem to always co-occur, e.g., setting stories with economic/geographic entities in the modern world. Whether they are 'causally' related, possibly due to the disenchantment of the tale by the insertion of 'realistic' elements [8], is outside our current remit. This co-occurrence also means that some combinations of properties do not occur in current corpus, although this may be 'corrected' by future expansion or limited alteration on the side of the experimenters. Ultimately, we expect the data set to be useful for learning or testing of algorithms modelling narrative similarity. We look forward to presenting the full results of the experiment in the near future.

References

1   David K. Elson and Kathleen R. McKeown. A platform for symbolically encoding human narratives. In Brian S. Magerko and Mark O. Riedl, editors, *Intelligent Narrative Technologies: Papers from the AAAI Fall Symposium. November 9–11, 2007, Arlington, Virginia*, number FS-07-05 in AAAI Technical Reports. AAAI Press, 2007.

2   David K. Elson and Kathleen R. McKeown. Extending and evaluating a platform for story understanding. In Sandy Louchart, Manish Mehta, and David L. Roberts, editors, *Intelligent Narrative Technologies II. Papers from the AAAI Spring Symposium*, number SS-09-06 in AAAI Technical Reports. AAAI Press, 2009.

3   Hans-Heino Ewers, editor. *Ludwig Bechstein, Deutsches Märchenbuch*. Reclam, Stuttgart, 1996.

4   Mark Alan Finlayson. *Learning Narrative Structure from Annotated Folktales*. PhD thesis, Massachusetts Institute of Technology, 2011.

5   Bernhard Fisseni and Benedikt Löwe. Which dimensions of narratives are relevant for human judgments of story equivalence? In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 114–118, İstanbul, 2012.

6   Bernhard Fisseni and Benedikt Löwe. Event-mappings for comparing frameworks for narratives, submitted.

7   Jacob Grimm and Wilhelm Grimm. *Kinder- und Hausmärchen*. Dieterichsche Buchhandlung, Göttingen, 1857.

8   Vanessa Joosen. Disenchanting the fairy tale: Retellings of "Snow White" between magic and realism. *Marvels & Tales*, 21(2):228–239, 2007.

9   Mikel J. Koven. Folklore studies and popular film and television: A necessary critical survey. *Journal of American Folklore*, 116(460):176–195, 2003.

10   Lawrence Lessig. *Remix: Making Art and Commerce Thrive in the Hybrid Age*. Bloomsbury Academic, 2008.

11   Kathleen Loock and Constantine Verevis, editors. *Film Remakes, Adaptations and Fan Productions: Remake/Remodel*. Palgrave Macmillan, 2012.

12   Benedikt Löwe. Methodological remarks about comparing formal frameworks for narratives. In Patrick Allo and Giuseppe Primiero, editors, *Third Workshop in the Philosophy of Information, Contactforum van de Koninklijke Vlaamse Academie van België voor Wetenschappen en Kunsten*, pages 10–28, Brussel, 2011. KVAB.

13   Benedikt Löwe. Methodological remarks about comparing formal frameworks for narratives. In Patrick Allo and Giuseppe Primiero, editors, *Third Workshop in the Philo-*

sophy of Information, Contactforum van de Koninklijke Vlaamse Academie van België voor Wetenschappen en Kunsten, pages 10–28, Brussel, 2011. KVAB.

14 Inderjeet Mani. *Computational Modeling of Narrative.* Number 5 in Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2012.

15 Neil McIntyre. *Learning to Tell Tales: Automatic Story Generation from Corpora.* PhD thesis, University of Edinburgh, 2011.

16 Charles Perrault. *Histoires ou contes du temps passé, avec des moralités.* Claude Barbin, Paris, 1697.

17 Walter Scherf. *Märchenlexikon.* Directmedia, Berlin, 2003.

18 Sandra K. D. Stahl, Elaine Jahner, Barbara Babcock, Barre Toelken, and Dell Hymes. Scarface vs. scar-face: The problem of versions. *Journal of the Folklore Institute*, 18(2/3):125–150, 1981.

19 Andrew Wright. Jane Austen Adapted. *Nineteenth-Century Fiction*, 30(3):421–453, 1975.

Propp's Morphology of the Folk Tale as a Grammar for Generation*

Pablo Gervás

Instituto de Tecnología del Conocimiento
Universidad Complutense de Madrid
Madrid, Spain
`pgervas@sip.ucm.es`

─── **Abstract** ───────────────────────────────────

The semi-formal analysis of Russian folk tales carried out by Vladimir Propp has often been used as theoretical background for the automated generation of stories. Its rigour and its exhaustive description of the constituent elements of Russian folk tales, and the enumeration of the patterns they follow, have acted as inspiration for several story generation systems, both sequential and interactive. Yet most of these efforts have attempted to generalize Propp's account to types of stories beyond the corpus that it arose from. In the process, a number of the valuable intuitions present in the original work are lost. The present paper revisits Propp's morphology to build a system that generates instances of Russian folk tales. Propp's view of the folk tale as a rigid sequence of character functions is employed as a plot driver. Unification is used to incrementally build a conceptual representation of discourse by adding to an ongoing draft story actions that instantiate the character functions. Story actions are defined by pre and post conditions on the state of the plot to account for the causal relations crucial to narrative. The potential of the resulting system for providing a generic story generation system is discussed and possible lines of future work are discussed.

1 Introduction

At the start of the 20th century, Vladimir Propp identified a set of regularities in a subset of the corpus of Russian folk tales collected by Afanasiev. Propp set out to study a subset of the corpus already classified as fairy tales, and concentrated on 100 of those tales to carry out this study. Over these tales he carried a systematic analysis in terms of character functions, understood as acts of the character, defined from the point of view of its significance for the course of the action. The conclusions of his study where that, for the given set of tales, the number of such functions is limited, the sequence of functions was always identical, and all these fairy tales could be considered instances of a single structure. His book "Morphology of the Folk Tale" [19] describes this set of character functions, the sequence in which they appear, and the overall structure of this archetype of fairy tale. Propp's work was intended to provide insights for the description and classification of folk tales. It has in fact been used in this way by many researchers [26, 14, 17].

However, the fact that it decomposes a tale into restricted set of elementary components, and outlines a procedure for putting them together to construct further tales has made it very appealing for researchers hoping to construct systems capable of generating stories automatically, both for sequential stories [24, 25, 8, 10, 26] and interactive ones [9, 7]. Yet most of these efforts have attempted to generalize Propp's account to types of stories beyond the corpus that it arose from, or combine it with additional techniques that had not been considered by Propp. A brief review of some of these efforts included in section 2.3 discusses the way Propp's work has been extended and adapted in some existing storytelling systems. In the process of these extensions and adaptations, a number of the valuable intuitions present in the original work are lost. As the field of computational narratology matures, it has become generally accepted that Propp's formalism can only be stretched so far, and that using it beyond its intended setting leads to severe limitations to story generation abilities of the resulting system. The story generation systems based on Propp's formalism have not managed to provide generic story telling capabilities. As a result both Propp's formalism and the goal of achieving generic story telling capabilities stand discredited.

In view of this situation, research in story generation has a chance to shift towards generation of quality stories for very specific domains [1]. It is in this last direction that Propp's work may find a different application niche. One of the reasons that made Propp's work so atractive to researchers in story generation is that Propp actually describes how his formalism might be used for the generation of tales. Seen in this light, Propp's formalism constitutes a blue-print for a story generation system intended to reproduce a particular model of story, while strongly adhering to specific genre and domain conventions. It is in this spirit that the present paper revisits Propp's morphology as a story generation procedure, exploring its potential for building a system that generates instances of Russian folk tales faithful to Propp's description. In this endeavour, a principle of economy is followed, considering the extension of the system with additional technologies only where their absence would clearly result in poorer stories. Additionally, an attempt has been made wherever possible to model Propp's formalism for this task in a declarative manner. This is with a view to hopefully replace in the future these declarative descriptions of Propp's formalism for Russian folk tales with a different set of descriptions, possibly capturing different types of story.

2 Previous Work

Before the proposed system can be described, a number of issues addressed by previous work must be presented: basic elements of Propp's morphology, Propp's description of how his morphology could be used to generate stories, Propp's influence on existing automated storytellers, and relevant insights from existing story generators even though not explicitly considered by Propp.

2.1 Elements of Propp's Formalism Relevant for Computational Implementation

The collection of tales that Propp focuses on involves stories built on combinations of a number of narrative ingredients: a protagonist sets out on a journey, usually triggered by a lack in his immediate environment or a villainy performed upon it, faces a villain, and in the process gets helped by a magical agent. A possible complication considered is the presence of a an additional character that competes with the protagonist for the role of hero of the story, which involves additional ingredients such as a gradual unveiling of the hero's real role

in the story, from initial presentation in disguise to the obtention of a reward towards the end, and usually involving recognition as a result of success on a difficult task.

The two corner stones of Propp's analysis of Russian folk tales are a set of roles for characters in the narrative (which he refers to as *dramatis personae*), and a set of character functions. These two concepts serve to articulate the morphology as an account of the elementary structure of the tales. Both of these concepts are constructed specifically for the family of tales begin considered. Therefore the set of roles includes fundamental elements such as the hero (who sets out on a journey), the dispatcher (who dispatches the hero on his journey), the villain (that the hero faces during the story), the donor (who provides the magical agent to the hero), the false hero (who competes with the protagonist for the role of hero of the story). The set of character functions includes a number of elements that account for the journey, a number of element that detail the involvement of the villain, including the villainy itself, some possible elaborations on the struggle between hero and villain, and a resolution, a number of elements that describe the dispatching of the hero, a number of elements that describe the acquisition of a magical agent by the hero, a number of elements concerned with the progressive unveiling of the hero's role in opposition to the false hero.

The sequence of character functions described by Propp is supposed to apply to all stories of the type described, so that any story will include character functions from this sequence appearing in the given order. With respect to the relative ordering, some deviation is allowed in that tales may depart from it by shifting certain character functions to other positions in the sequence.

Character function are sometimes repeated three times. This is a widely spread feature fo fairy tales, called *trebling*, where upon three instances of a particular event occur in sequence (stepmother tries to kill Snow White in three different ways but only the last one succeeds, Cinderella attends the Prince's ball under disguise on three consecutive nights but only on the last one does she forget to leave before midnight,...). For the purposes of the system described in this paper, this type of refinement may be left for consideration at a later stage.

Character functions in a given narrative are related to one another by long range dependencies related to motivation and correference.

Propp says:

> "The majority of character's acts in the middle of a tale are naturally motivated by the course of the action, and only villainy, as the first basic function of a tale, requires a supplementary motivation." [19, p. 75]

The concept of motivation that is referred to here concerns the network of causal relations between the different events of a story that a reader usually provides during comprehension [23]. This network representation determines the overall unity and coherence of the story. When considering the procedural generation of tales based on this model, motivation introduces a significant problem. The selection of what particular instantiation of a character function to use at a particular point of the tale must take into consideration that the new character function instance appear appropriately motivated by the preceding selections already made. This is a fundamental aspect for the success of the result as a story. As shown later, in order to account for this problem additional computational mechanisms need to be added.

Some character functions are implicitly linked to one another. Propp mentions two types of link between character functions:[1] elements which are always linked with varieties

---

[1] Propp's own abbreviations for specific types of his character functions are used to allow reference to the original work.

corresponding to one another (alternative instantiations of struggle and victory, such as H^1 – fight in an open field – always connected to I^1 – victory in an open field, or of villainy and its liquidation, such as A^{11} – enchantment – linked to K^8 – the breaking of a spell – ...); and elements that act as necessary preconditions to others (second element cannot happen unless the first one is already present) but allow for variation (the hero can only be rescued from pursuit if a pursuit has commenced, but rescue can take several forms regarless of how the pursuit started).

These links are mostly concerned with particular instantiations of certain character functions being linked to instantiations of character functions that went before them. This is one of the ways in which overall coherence of the tale can be ensured: characters kidnapped at the beginning are freed towards the end, and so on. A computational procedure must take these links into account when deciding which characters to assign to particular roles in each new character function added to a story. If the sister of the hero was bewitched at the start, it is she that needs to be released from the spell towards the end.

Character functions are so named because, in Propp's understanding, they represent a certain contribution to the development of the narrative by a given character. When he talks about the set of characters of the story (or *dramatis personae*), Propp constantly reccurs to a set of labels to describe particular roles played by characters in tales. They are gathered together in chapter VI where he discusses the distribution of functions among dramatis personae. For simplicity I will refer to these as *role names*, though Propp does not. Some examples of these roles are: the *villain*, the *donor* (who provides the hero with a magical agent), the *helper* (usually a magical agent, that helps the hero carry out his tasks), the *dispatcher* (who sends the hero on his mission), the *hero* (the protagonist of the story), and the *false hero* (who maliciously sets himself up to usurp the protagonist as hero of the story).

Propp defines these only in terms of the set of character functions that can be grouped around each one of them, as involving the same character. This set he refers to as the *sphere of action* for a particular role name. In the description of each character function in chapter III, Propp mentions how the character fulfilling a particular named role is involved in the various actions that can instantiate that character function (the villain carries out the villainy, the dispatcher sends the hero on his mission, the hero departs from home,...). If a procedural solution is sought that attempts to model closely the vision of tales that Propp had, these narrative roles must be explicitly defined, and some means of explicitly defining their participation in each type of character function should be provided, to ensure that these participations are instantiated by particular characters in a coherent manner throughout the tale.

Finally, Propp considers how more complex stories can be seen to conform to his morphology. To achieve this, the proposed sequence is considered as the skeleton for a single narrative thread, and complex stories may be composed of several such threads combined in different ways. In Propp's terminology, these narrative threads are called *moves* of a tale. Even though part of the same story, different moves may each involve different heroes or villains. The present paper is concerned with the construction of single move tales. The construction of multiple move tales may be addressed at a later stage both as a way of combining single moves and as a refinement of the construction procedure to produce individual moves suitable for combination with others.

2.2   Propp's Description of Tale Generation

Propp provides in his book a very clear description of how his morphology could be used for story generation:

"In order to create a tale artificially, one may take any A, then one of the possible B's then a $C\uparrow$, followed by absolutely any D, then an E, the one of the possible F's, then any G, and so on. In doing this, any elements may be dropped, or repeated three times, or repeated in various forms. If one, then distributes functions according to the dramatis personae of the tale's supply of by following one's own taste, these schemes come alive and become tales. Of course, one must also keep motivations, connections, and other auxiliary elements in mind" [19, pp. 111–112]

In addition to this clearly procedural description he provides a number of constraints that a potential storyteller should obey and an emumeration of the points where a storyteller has freedom to decide.

The constraints on the story teller are:

1. "The storyteller is constrained [. . .] in the overall sequence of functions, the series of which develops according to the above indicated scheme." [19, p. 112]
2. "The storyteller is not at liberty to make substitutions for those elements whose varieties are connected by an absolute or relative dependence." [19, p. 112]
3. "In other instances, the storyteller is not free to select certain personages on the basis of their attributes in the event that a definite function is required." [19, p. 112]

The points where Propp considers that a storyteller has a certain freedom are:

1. "In the choice of those functions which he omits, or, conversely, which he uses" [19, p. 112]
2. "In the choice of the means (form) through which a function is realized." [19, p. 112]
3. in the assignment of story characters to particular slots in functions: "If one then distributes functions according to the dramatis personae of the tale's supply or by following one's own taste, these schemes come alive and become tales" [19, pp. 111–112] and "The storyteller is completely free in his choice of the nomenclature and attributes of the dramatis personae. Theoretically the freedom here is absolute." [19, pp. 112–113]
4. "The story teller is free in his choice of linguistic means." [19, p. 113]

On the third point, Propp follows on to discuss in rather vague terms that people do not make wide use of this freedom, preferring to let personages recur much as functions do. So there is a typical villain, a typical donor. . . Given the level of uncertainty involving this description, it has been decided not to consider it in the present system. The fourth point surely underlies Propp's decision not to address linguistic issues in his morphology at all. We follow this decision in deciding not to address the linguistic rendering of the tales in the initial implementation of our system.

The remaining insights are considered in a computational implementation in section 3.

2.3   Propp in Existing Automated Storytellers

Lang [13] developed the Joseph system to produce instances of stories akin to those in the Afanasiev corpus of Russian tales that Propp used as inspiration, but Lang departed from Propp's formalism in favour of a story grammar closer to Thorndyke's model [22], coupled with a complex network of logical procedures for modelling time and rational intention. Lang explicitly identifies some of the difficulties inherent in trying to formalize Propp's account with a view to generation [19, section 2.1.1]. The grammar used by Lang represents a story as a sequence of episodes, each one involving an initial event, a reaction by the protagonist to the event, and an outcome.

Turner [24] mentions Propp's work as an inspiration for his thesis, precisely for its potential as a story generator (as described in the introduction). But he also specifically claims that there were no traces of Propp's work in the final version of his story writing program. For his MINSTREL system, Turner claims to have been inspired by Propp's work, but then developed a system that combined case-based reasoning and planning to produce stories about King Arthur and his knights, without resorting to Propp's ideas at all.

Peinado [8] developed a description logic ontology for Propp's set of character functions, but then focused on exploring the potential of description logic ontologies for providing a knowledge intensive case-based solution for tale generation, based on reusing structure from existing tales into new ones. This ontology is a valuable resource that could have been used to implement further systems based on Propp's formalism. However, description logic ontologies have proven to be very good at representing the world as it is, and not so good at representing a world liable to change. Representation of change is a fundamental aspect of narrative. For this reason we have preferred to rely on a different representation mechanism for the effort reported in this paper.

Grasbon and Braun [9] and Fairclough and Cunnigham [7] adapted some of Propp's ideas to the realm of interactive story telling. They rely on character functions based on Propp's work to develop a story engine for interactive narrative. They extend the concepts of Propp with the idea of polymorphic functions, which can have different outcomes depending on user interaction.

The Proppian fairy tale Markup Language (PftML) [16] is an XML application developed by University of Pittsburgh's researchers based on Propp's work. PftML utilizes a Document Type Definition (DTD) to create a formal model of the structure of Russian magic tale narrative and to help standardize the tags throughout a corpus when analyzing it. PftML has been tested on a subset of the same Russian language corpus from which Propp drew has been used, as an empirical test of the conclusions of Propp's initial analysis against the original data. PftML constitutes a formal grammar for Propp's morphology of the folk tale. However, it is not well geared towards generation, and it misses many of the subtleties uncovered in section 2.1.

Fairclough and Cunningham [6] implement an interactive multiplayer story engine that operates over a way of describing stories based on Propp's work, and applies case-based planning and constraint satisfaction to control the characters and make them follow a coherent plot. They define a plot as a series of character functions and a series of complication-resolution event pairs, where a complication occurs whenever a character performs a function that alters the situation of the hero. A case based reasoning solution is used for storyline representation and adaptation. They use 80 cases extracted from 44 multi-move story scripts given by Propp. There are stories composed of one, two or more moves. A case is a move, seen as a story template, to be filled in by a constraint satisfaction system that chooses which characters perform the functions.

2.4 Relevant Insights of Existing Story Generators

There are a large number of story generators in existence, relying on a multiplicity of techniques. For the sake of brevity only those specific ingredients of some of them that are relevant for the solution proposed in this paper are listed here.

Although Propp never described his formalism as a grammar for stories, it has often been described as such [24]. The popularity of grammars as a representation mechanism peaked in the seventies and early eighties as a result of Noam Chomsky's work on formal grammars. A popular technique for modelling common phenomena was to develop a grammar for them.

This lead to the concept of a story grammar, pioneered by Rummelhart [21] and later taken up by Thorndyke [22] and many others. The very concept of story grammar was questioned by Black and Willensky [3], and a debate around story grammars went on for many years and lead to the discreditation of the concept as an actual model of human cognitive processing of stories. Nevertheless, story grammars remained a popular technique with researchers in story generation. The Joseph system [13] and the BRUTUS system [4] were based on story grammars. They both produced a succesful number of stories of high quality. In this sense, the concept of story grammar for the generation of stories can be considered validated as a sound and successful technology.

A different concept related with the implementation of narrative systems is that of story actions as operators that change the world. Actions in a story are applicable if certain conditions hold in the state of the world before they happen, and after they happen they change the state of the world. This idea has been represented by defining actions with an associated set of preconditions and another of postconditions or effects. This approach to defining actions is important because it constitutes a possible way of capturing the causal dependencies that constitute a fundamental ingredient of narrative as it is understood by people [23]. It has become popular in story generation through the numerous research efforts that use planning techniques [20, 2, 11], which are inherently based on this concept. Even systems based on alternative generation technologies include the possibility of associating pre and postconditions to actions, such as the information on emotional links between characters considered in the MEXICA system [18] or the preconditions added to the representation of actions in the Joseph system [13].

3   A Computational Solution for Proppian Story Generation

The first step for considering Propp's formalism as a computational procedure would be to define specific representations for the concepts involved. In the description of Propp's formalism given in section 2.1 we have relied on two different concepts that would need to be assigned a conceptual representation:

- *character function* (a label for a particular type of acts involving certain named roles for the characters in the story, defined from the point of view of their significance for the course of the action)
- possible instantiations of a character function in terms of specific *story actions*, involving a number of *predicates* describing events with the use of *variables* that represent the set of characters involved in the action

To fully capture Propp's restrictions (constraint 3), story actions will also include non-narrative predicates which encode constraints on the specific choice of dramatis persona that can fill particular argument slots in the predicates of the story action; for instance, the fact that the author of a villainy must be the villain.

The sequence of character functions chosen as backbone for a given story we will refer to as a *plot driver*.

The set of story actions available for instantiating a given character function, as defined by Propp, includes several variants concerning the form of the action. For instance, a villainy can take the form of kidnapping a person, seizing a magical agent, ruining the crops,... Each of these would be represented in our proposal by an action with a set of preconditions and a set of postconditions. To keep track of the effects of this actions as they are added to the story, some form of representation of the context must be employed. As the simplest possible solution, a representation of the context is considered as a set of states, each one

■ **Table 1** Examples of story actions.

Character function	villainy	liquidation
Preconditions	married H Y	married H W
	hero H	sundered H W
	villain X	hero H
Action	makes_disappear X Y	resume_marriage H W
Postconditions	victim Y	
	sundered H Y	

representing the state of the world before a certain story action took place. A *state* of the world is represented as a set of predicates describing the facts that hold in that state. The sequence of states for a given story we call a *fabula*.

We represent a *story action* as a set of predicates that describe an instance of a character function. Links with preceding story actions are represented as dependencies of the story action with predicates that need to have appeared in previous story actions (preconditions). Therefore a story action involves a set of preconditions (predicates that must be present in the context for continuity to exist), and a set of postconditions (predicates that will be used to extend the context if the action is added to it). Some additional predicates not corresponding to events in the story are added to encode the sphere of action to which each story action belongs. These predicates explicitly link the corresponding narrative role to a particular variable in the story action The predicates in a story action are defined over free variables as arguments. This ensures that relative instantiation of the various arguments in the predicates of a story action is coherent, as discussed later. Table 1 includes example of story actions linked by preconditions.

Each successive state in a fabula contains all the predicates arising from the preceding actions that have not been retracted by a story action since they occurred. This is difficult to read. Also it is difficult to define over such a structure significant metrics on measures such as number of predicates in which a certain character appears (which we will need to consider when measuring the structural quality of a story). For this purpose we define a final structure called a *flow* for a story, which is simply an ordered sequence of all the predicates in the fabula, such that each one appears only once, and grouped into subsets according to the particular state of the fabula in which they were first introduced.

Based on this representation, the procedure originally sketched by Propp can be subdivided into the following stages, each one of which will be addressed by a different module in our proposed system:

- employ an algorithmic procedure for generating a sequence of character functions considered valid for a tale (*plot driver generator*)
- given a valid sequence of character functions, progressively select instantiations of these character functions in terms of story actions (*fabula generator*)
- given a fabula where all variables have been replaced by constants, produce a flow for the story (*flow generator*)

For each of these stages a computational decision procedure must be selected. We are considering a possible computational implementation. For this purpose we intend to consider in the first instance the simplest representation and the simplest procedures compatible with

acceptable results. To this end, a number of computational options for some of these modules have been considered, together with a knowledge engineering effort to produce the required resources. The results have been empirically tested for fulfillment of Propp's constraints. The following section report on the development, the evaluation procedures, and the results of the tests.

3.1   The Implementation

The computational solution described in this paper has been partially implemented as a working prototype written in Java and operating over a small set of resources defined as plain text files. This development involved a very small effort of simple coding of the overall algorithmic procedures, but a considerable effort of knowledge engineering over the set of resources.

3.1.1   Plot Driver Generators

Three possible implementations have been considered for plot drivers:

- a baseline plot drivers that randomly selects character functions, not necessarily in sequence, up to a randomly decided number,
- a plot driver that follows a canonical sequence of character functions, deciding at each stage whether to add it to the plot driver or not,[2] or
- a grammar based plot driver which generates based on a gramar automatically extracted from a subset of the schemes analyzed by Propp in Appendix III.

In the second case, a reference sequence of character functions is constructed following the matrix employed by Propp in Appendix III for tabulating his analyses of stories from his corpus. This sequence includes several possible placements of certain character functions in the sequence, to capture the accepted possibilities for inversion.

The grammar for the third case is built automatically from a set of Propp's schemes annotated as a grammar by considering subgroupings of character functions that recurred frequently, and using Propp's own terminology for different segments of a tale. An example of grammar is given in Table 2.

Performance results for the three plot driver generators considered are given in section 3.3.

3.1.2   Fabula Generators

A fabula generator receives a plot driver and selects story actions for the character functions given in it. To do this, the fabula generator has to define a fabula, a sequence of states that contain a chain of instances of character functions ideally somehow linked by having their preconditions fulfilled by the context. The initial state by default incorporates all predicates of the first action, and each valid action added to the fabula generates a new state that incorporates all predicates of the previous state, plus the predicates of the new action.

A mapping is established between the set of story actions and the set of character functions, so that each of the available story actions is considered a possible instantiation of a given character function.

---

[2] With the exception of the villainy/lack character functions, for one of the two is always added to a story to ensure story interest. This follows Propp's own suggestion [19, p. 102].

◼ **Table 2** A Simple Grammar.

```
FOLKTALE              =   COMPLICATION DONOR COURSE_OF_ACTION CLOSURE
FOLKTALE              =   COMPLICATION DONOR COURSE_OF_ACTION
FOLKTALE              =   COMPLICATION COURSE_OF_ACTION CLOSURE
FOLKTALE              =   COMPLICATION COURSE_OF_ACTION
FOLKTALE              =   ↑ DONOR T COMPLICATION DONOR COURSE_OF_ACTION CLOSURE

COMPLICATION          =   TRIGGER
COMPLICATION          =   TRIGGER MEDIATION ↑
COMPLICATION          =   TRIGGER MEDIATION
COMPLICATION          =   TRIGGER M MEDIATION

TRIGGER               =   A
TRIGGER               =   a

MEDIATION             =   C
MEDIATION             =   B C

DONOR                 =   F
DONOR                 =   F G
DONOR                 =   D E
DONOR                 =   D E F
DONOR                 =   D E F G
DONOR                 =   D E G F

COURSE_OF_ACTION      =   TASK
COURSE_OF_ACTION      =   ↓ PURSUIT
COURSE_OF_ACTION      =   K ↓ PURSUIT
COURSE_OF_ACTION      =   ↓ PURSUIT o
COURSE_OF_ACTION      =   o DONOR K T PURSUIT
COURSE_OF_ACTION      =   CONFRONTATION K ↓
COURSE_OF_ACTION      =   CONFRONTATION K ↓ PURSUIT
COURSE_OF_ACTION      =   K ↓
COURSE_OF_ACTION      =   CONFRONTATION
COURSE_OF_ACTION      =   CONFRONTATION K

PURSUIT               =   Pr Rs
PURSUIT               =   Pr DONOR Rs

CONFRONTATION         =   H I
CONFRONTATION         =   I

TASK                  =   M N

CLOSURE               =   W
CLOSURE               =   Q W
CLOSURE               =   Q Ex U W
CLOSURE               =   T W
CLOSURE               =   X U W ↓ X
```

To evaluate whether the preconditions of a story action are satisfied by the context, they are unified with the set of predicates that hold in that state. This serves two purposes:

- if the preconditions are not satisfied, an alternative story action will be considered
- unification allows any of the free variables in these preconditions to unify with those in the predicates holding in the fabula

A story action is considered a valid extension of a given fabula if the set of its preconditions can be successfully unified with the predicates in the latest state of the fabula. Once the story action is added, the next state is built by extending the preceding state with the action and the postconditions of the story action.

When the preconditions unify with the state in the fabula, any replacement of free variables in the preconditions is carried over to the rest of the story action before it is added to the context. This ensures that the story action become coherent with the rest of the predicates in the fabula, creating continuity.

This enables the system to model long range dependencies between character functions. If the choice for a character function such as liquidation of misfortune or lack depends on which particular story action was chosen to instantiate the character function for lack, this procedure will both block non appropriate instantiations for liquidation (as their preconditions will not be satisfied) and will ensure the appropriate assignment of variable names to ensure coherence (for instance, that the person that was kidnapped at the beginning be freed towards the end). The additional predicates encoding the sphere of action to which each story action belongs enforce a correct distribution of functions over dramatis personae. Overall, the use of unification models Propp's constraints 2 and 3.

However, a requirement of strict unification narrows down the set of options for extending to a very small set of story actions. given the current size of the set of available story action (described in section 3.2), it was considered advisable to allow a certain relaxation of this constraint. This corresponds to an operation of accommodation [15], in which if some of the preconditions unify and some do not, those that do not can be added to the context together with the action and its postconditions. Given a story action and a fabula, the less preconditions that need to be accommodated, the more appropriate the story action is considered as an extension of the fabula.

Two fabula generators have been considered:

- a purely random fabula generator that, for a character function given by a plot driver, picks at random one of the possible story actions available for that character function (considered as baseline)
- a fabula generator that adds the story action that best unifies with the context (allowing for accommodation)

Performance results for these fabula generators considered are given in section 3.3.

3.2   The Knowledge Engineering Effort

Although the grammar proved easy to write, developing and debugging the set of story actions required for the 31 character functions proved to be an onerous task.

Propp never exhaustively described his set of instances of character functions nor his set of schemes for the structure of tales, and it is difficult to match the numbering of the tales that he uses as reference with existing compilations of translated (into English) tales from the Afanasiev corpus. There is therefore no obvious source for constructing a set of resources

to match Propp's descriptions.[3] In order to develop a set of story actions for our system, we followed a detailed procedure to guarantee close conformance to Propp's specification and reasonable coverage of his set of character functions.

A set of story actions conforming to the described representation was built following Propp's descriptions of character functions (Chapter III) and matching the set of abbreviations proposed for the analysis formalism (Appendix IV); ensuring that dependencies indicated by Propp are represented as preconditions

We then tested that this set of story actions could be used to construct fabulas matching both the structure and the specific choice of instances of each character function given in the set of schemes listed by Propp in his book. The resulting fabulas showed acceptable continuity (or impression thereof under some possible interpretation). Success in this respect was taken to constitute a certain degree of validation of both the set of story actions engineered (in as much as they provide enough links to guarantee continuity in terms of co-occcurrence of variables across the whole set of predicates for the fabula) and the unification procedure employed (validation procedure for extensions enforces continuity).

Finally we tested that this set of story actions could be used to construct fabulas for the original schemes described by Propp in Appendix III but allowing freedom of choice with respect to the specific choice of story action. Success in this respect was taken to constitute a certain degree of validation of both the set of story actions engineered (enough articulation to provide variation over a set of possible instantiation choices for each character function) and the unification procedure employed (validation procedure for extensions still enforces continuity even when free choice of instance is allowed).

A final set of 280 story actions was obtained corresponding to the restricted set of 24 story actions that Propp used to describe the schemes given in Appendix III (Propp explains that functions of the preparatory section were not included for lack of space).

3.3 Evaluation

Given that the development effort has focused at a very abstract level of representation, evaluation has to be considered at a corresponding level to provide valid feedback for the improvement of the system. As the linguistic modelling of the stories has not been addressed, evaluation by human volunteers is plagued with difficulty. Introducing some kind of rapidly constructed stage for rendering the results as text by providing text templates for each story action (as done in some existing story generators [18]) is likely to introduce noise in terms of elements present in the text and not necessarily produced by the system. Asking human evaluators to rate the quality of an abstract representation as produced by the system runs the risk of judgements being clouded by the difficulty of interpreting the representation.

Additionally, evaluations by humans necessarily have to be restricted to a small number of instances of system output. The choice of which particular instances to test is left to the designer of the experiment, and there is a risk of focusing on examples that are not representative of system performance overall.

As an alternative, quantitative procedures have been defined to measure the specific qualities desired for each stage of the representation, at a corresponding abstract level. These procedures can be applied to a large number of system results, providing a measure of the

---

[3] This is probably one of the reasons why no one has yet attempted to build a story generator for Russian fairy tales. Given the need to engineer the set of knowledge resources from scratch, most researchers have preferred to start from alternative material that was either more easily available or closer to their native intuitions of folklore or narrative.

■ **Table 3** Results for plot drivers.

Random	Sequence	Grammar
59.2	100	83.48

quality of system output at the working level of abstraction and applicable to a broad range of system results, leaving no doubt as to their significance over the complete set of outputs.

3.3.1 Plot Driver Generators

Plot driver generators must obey constraint 1, as described in section 2.2. To establish the extent to which the various implementations fulfill this constraint, a measure of conformance to a reference sequence has been defined. The key measure to consider is, given a certain character function appearing in a candidate plot driver, how many of the functions preceding/following it in the plot driver are contained in the part of the reference sequence that goes before/after (the best scoring of) its appearances in the reference sequence. This value is normalised as a percentage over the length of the plot driver. This measure is 100 if all character functions before and after the one considered have the same relative order in the reference sequence. The measure for a complete driver is taken as the average value for all its functions. This is 100 if the plot driver satisfies perfectly the order in the reference sequence and degrades towards 0 if some of its character functions appear out of place with respect to the given sequence.

Results for the three different plot driver generators that have been tried are reported in Table 3. Each of the alternative implementations was run 100 times and values were averaged over the results.

These results confirm as expected that the random approach results in plot drivers that do not conform to Propp's requirement, that strict adherence to Propp's instructions results in perfect conformance. The results for the grammar approach are more surprising. The grammar obtained from the set of schemes described by Propp performs considerably better than the random baseline, but almost 20 points below the version following the sequence strictly. This is due to the fact that the set of schemes contains several exceptions to the general rules. As the grammar is extrapolated from actual tales, this result suggests that allowing a certain flexibility with respect to Propp's rules is likely to produce less rigid stories that resemble more closely the kind that might be produced by humans. However, testing this hypothesis is beyond the scope of the present paper.

3.3.2 Fabula Generators

Fabula generators must obey constraints 2 and 3, and the links between instances of character functions described in section 2.1.

In order to evaluate the extent to which the different construction algorithms fulfill these requirements, a number of metrics have been defined over the flow for a story, which is the closest representation produced by our system of the final shape of the story. Consideration of the referential chains in the flow (the set of predicates in which each character occurs) is significant of the degree of continuity achieved. These measures include: the maximum length of referential chain (%MLRC), the minimum length of referential chain (%mLRC), and the average length of referential chain (%avLRC). All of these measures are normalised

Table 4 Results for fabula generators.

	FL	%MLRC	%mLRC	%avLRC
PR	35.52	11.98	2.6	4.39
UA	37.38	47.58	2.57	7.64

Table 5 Example story.

hero 296	lack 74 75	test 284 296	disguised 296
captive 295	*	donor 284	artisan 613
asks 295 296 297	dispatches 261 296	*	apprentice 296 613
*	seeker_hero 296	finds 296 453	unrecognised 296
not_perform_service 296	banished 261	follows 296 453	*
negative_result 296	victimised_hero 261	at_target_location 296	false_hero 616
*	transported_to 261 296	*	claims 616 617
villain 73	*	arrives 296 612	unfounded 617
maims 73 74	sets_out 296	location 612	*
victim 74	*	home 612	returns 296

over the length of the flow (FL) and expressed as a percentage, to allow comparison across tales of different length.

The fabula generators are tested over plot drivers produced by generators that follow Propp's sequence strictly. Results for their evaluation are given in Table 4. Each of the alternative implementations was run 100 times and values were averaged over the results.

These results confirm that the use of unification and accommodation as techniques for ensuring continuity over stories do result in longer referential chains, and in stories that more clearly involve specific characters over a period of time.

4 Discussion

An example of a story is given in Table 5. The example appears broken down over four columns, separated by * into groups corresponding to single story actions. The arguments of predicates are represented by numbers, each corresponding to a character, location or object in the story. The story concerns character 296, who behaves badly at the start of the story, is banished, is tested by a donor, finds a trail that leads him home, arrives disguised as an apprentice to an artisan, suffers an impostor and returns. The example was picked out from the combined results of a plot driver relying on Propp's sequence and a fabula generator relying on unification with accommodation.

The story evidences some of the problems with literal implementation of Propp's algorithm: long range dependencies are captured by the unification mechanism when the two character functions involved appear in the plot driver (hero sets out and returns), but the mechanism for generating the driver does not take them into account. As a result, for instance, the villain and the false hero go unpunished in this case. This suggests that some computational means of taking the long range dependencies into account must be included to improve the performance of the system. The grammar approach has a potential for representing implicitly these long range dependencies, though empirical testing of this hypothesis is again beyond the scope of this paper.

With respect to the usefulness of Propp's various concepts, some of the roles described for dramatis personae seemed less relevant to the actual set of character functions than others (*princess*, for instance, was much more loosely linked to specific functions than *hero* or *villain*). Equivalent roles not mentioned by Propp, such as the victim of the villainy, were more useful in establishing continuity.

Preliminary results indicate that the overall quality of the generated stories is highly dependent on the quality of the set of story actions. This evidence confirms the limitations of knowledge-based computational solutions, but it also highlights a potential for generalization for the approach followed in the paper. Although the current initiative strove to build a system as faithful as possible to Propp's formalism, all references to Propp's material in the resulting implementation occur only within the set of plain text files that constitute the knowledge resources. The choice and names of the character functions, and the restrictions imposed on how they may combine are captured in the grammar. The set of story actions and the relationships between stated in terms of preconditions are also written in a separate text file. The actual Java code that exploits these resources is independent of Propp's particular solution for Russian folk tales. As a result, it would be possible to write an alternative set of resources to use a completely different grammar, over elements of a similar nature but which need no longer be called character functions, and which combine according to different rules, or which get instantiated with a completely different set of story actions, and assigned to different characters.

Because of this property, the approach presented in the paper has the potential for being ported to different domains (for instance, to science fiction stories by changing the set of story actions and the set of characters) or adapted to account for different structural analyses of narrative (by changing the grammar into one that covers, for instance, Campbell's account of the hero's journey [5] or Lakoff's account of the structure of fairy tales [12]).

5   Conclusions

The theoretical account of Russian fairy tales provided by Vladimir Propp has been revisited as potential source for a procedure for story generation. Although many research efforts in story generation have considered Propp as an inspirational source, none of them has explored the actual procedures explicitly described by Propp. By considering the simplest possible implementation of these procedures, a framework for story generation has been developed that takes full advantage of the intuitions behind Propp's account but which is built in a modular and declarative manner so that particular details arising from Russian folk tales can later be replaced with material from alternative knowledge sources.

The approach suffers from the limitations inherent to any knowledge intensive approach, in the form of a heavy knowledge engineering effort required to kick start the necessary set of resources. Preliminary results show strong coupling between the quality of these resources and the quality of the resulting stories. This can be seen as a weakness in terms of a deep adaptation curve for any new domain or new application, but also as a significant strength, in terms of the possibility of extending it to other domains and of targeted refinement until a desired level of quality is reached.

The various refinements and possible extensions described through the paper will be considered as future work.

References

1 Nicholas D. Allen, John R. Templon, Patrick Summerhays McNally, Larry Birnbaum, and Kristian Hammond. StatsMonkey: A data-driven sports narrative writer. In Mark Alan Finlayson, editor, *Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium*, number FS-10-04 in AAAI Technical Reports, pages 2–3. AAAI Press, 2010.

2 Byung-Chul Bae and R. Michael Young. A use of flashback and foreshadowing for surprise arousal in narrative using a plan-based approach. In Ulrike Spierling and Nicolas Szilas, editors, *Interactive Storytelling, First Joint International Conference on Interactive Digital Storytelling, ICIDS 2008, Erfurt, Germany, November 26-29, 2008, Proceedings*, number 5334 in Lecture Notes in Computer Science, pages 156–167. Springer, 2008.

3 John B. Black and Robert Wilensky. An evaluation of story grammars. *Cognitive Science*, 3:213–230, 1979.

4 Selmer Bringsjord and David A. Ferrucci. *Artificial Intelligence and Literary Creativity: Inside the Mind of BRUTUS, a Storytelling Machine*. Erlbaum, 1999.

5 Joseph Campbell. *The Hero with a Thousand Faces*. Princeton University Press, Princeton, second edition, 1968.

6 Chris Fairclough and Pádraig Cunningham. A multiplayer case based story engine. In Quasim H. Mehdi, Norman E. Gough, and Stéphane Natkin, editors, *4th International Conference on Intelligent Games and Simulation (GAME-ON 2003), 19–21 November 2003, London, UK*, pages 41–46, London, 2003. EUROSIS.

7 Chris Fairclough and Pádraig Cunningham. A multiplayer O.P.I.A.T.E. *International Journal of Intelligent Games & Simulation*, 3(2):54–61, 2004.

8 Pablo Gervás, Belén Díaz-Agudo, Federico Peinado, and Raquel Hervás. Story plot generation based on CBR. *Knowledge-Based Systems*, 18:235–242, 2005.

9 Dieter Grabson and Norbert Braun. A morphological approach to interactive storytelling. In Monika Fleischmann and Wolfgang Strauss, editors, *CAST 2001. Living in Mixed Realities: Conference on Artistic, Cultural and Scientific Aspects of Experimental Media Spaces, September 21–22, 2001, Schloss Birlinghoven, Sankt Augustin*, netzspannung.org event, pages 337–340, 2001.

10 Shohei Imabuchi and Takashi Ogata. Story generation system based on propp theory as a mechanism in narrative generation system. In Masanori Sugimoto, Vincent Aleven, Yam San Chee Chee, and Baltasar Fernández-Manjón, editors, *2012 IEEE Fourth International Conference On Digital Game And Intelligent Toy Enhanced Learning, DIGITEL 2012, Takamatsu, Japan, March 27-30, 2012*, pages 165–167, Los Alamitos, CA, USA, 2012. IEEE Computer Society.

11 Arnav Jhale and R. Michael Young. Cinematic visual discourse: Representation, generation, and evaluation. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(2):69–81, 2010.

12 George P. Lakoff. Structural complexity in fairy tales. *The Study of Man*, 1:128–150, 1972.

13 R. Raymond Lang. *A Formal Model for Simple Narratives*. PhD thesis, Tulane University, 1997.

14 Piroska Lendvai, Thierry Declerck, Sándor Darányi, Pablo Gervás, Raquel Hervás, Scott A. Malec, and Federico Peinado. Integration of linguistic markup into semantic models of folk narratives: The fairy tale use case. In Nicoletta Calzolari, Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors, *Proceedings of the Seventh conference on International Language Resources and Evaluation*, pages 1996–2001, Paris, 2010. European Language Resources Association.

15 David Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8:339–359, 1979.

16   Scott A. Malec. Proppian structural analysis and XML modelling. Presented at *Computers, Literature and Philology (CLiP 2001)*, 2001.

17   Scott A. Malec. AutoPropp: Toward the automatic markup, classification, and annotation of Russian magic tales. In Sándor Darányi and Piroska Lendvai, editors, *Proceedings of the First International AMICUS Workshop on Automated Motif Discovery in Cultural Heritage and Scientific Communication Texts*, pages 112–115, Szeged, 2010. University of Szeged, Faculty of Arts, Department of Library and Human Information Science.

18   Rafael Pérez y Pérez. *MEXICA: A Computer Model of Creativity in Writing*. PhD thesis, University of Sussex, 1999.

19   Vladimir Yakovlevich Propp. *Morphology of the Folktale*. University of Texas Press, Austin, TX, 2nd edition, 1968. transl. L. Scott.

20   Mark Owen Riedl and R. Michael Young. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:217–268, 2010.

21   David E. Rumelhart. Notes on a schema for stories. In Daniel G. Bobrow and Allan Collins, editors, *Representation and Understanding: Studies in Cognitive Science*, pages 211–236, New York, 1975. Academic Press, Inc.

22   Perry W. Thorndyke. Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology*, 9:77–110, 1977.

23   Tom Trabasso, Paul van den Broek, and So Young Suh. Logical necessity and transitivity of causal relations in stories. *Discourse Processes*, 12(1):1–25, 1989.

24   Scott R. Turner. *MINSTREL: a computer model of creativity and storytelling*. PhD thesis, University of California at Los Angeles, Los Angeles, CA, USA, 1993.

25   Cati Vaucelle. Générateur d'histoires. Master's thesis, Université Paris VIII, 2000.

26   Takenori Wama and Ryohei Nakatsu. Analysis and generation of Japanese folktales based on Vladimir Propp's methodology. In Paolo Ciancarini, Ryohei Nakatsu, Matthias Rauterberg, and Marco Roccetti, editors, *New Frontiers for Entertainment Computing*, number 279 in IFIP International Federation for Information Processing, pages 129–137. Springer, 2008.

Computationally Modeling Narratives of Social Group Membership with the Chimeria System*

D. Fox Harrell, Dominic Kao, and Chong-U Lim

Computer Science and Artificial Intelligence Laboratory (CSAIL)
Massachusetts Institute of Technology
Cambridge, MA, USA
`{fox.harrell,dkao,culim}@mit.edu`

Abstract

Narratives are often used to form, convey, and reinforce memberships in social groups. Our system, called *Chimeria*, implements a model of social group membership. Here, we report upon the Chimeria Social Narrative Interface (*Chimeria-SN*), a component of the *Chimeria* system, that conveys this model to users through narrative. This component is grounded in a sociolinguistics model of conversational narrative, with some adaptations and extensions in order for it to be applied to an interactive social networking domain. One eventual goal of this work is to be able to extrapolate social group membership by analyzing narratives in social networks; this paper deals with the inverse of that problem, namely, synthesizing narratives from a model of social group membership dynamics.

1998 ACM Subject Classification H.4.3 Communications Applications, H.5.1.b Artificial, augmented, and virtual realities, J.4.c Sociology, K.4.2 Social Issues, K.4.3.b Computer-supported collaborative work

Keywords and phrases computational narrative, cognitive categorization, social classification, social group membership and naturalization, social media

Digital Object Identifier 10.4230/OASIcs.CMN.2013.123

1 Introduction

Everyone belongs to social groups based on factors such as musical preference, fashion, gender, or race. Narratives are often used to form, convey, and reinforce memberships in such social groups. Furthermore, a robust model of group membership can be an important aspect for modeling many everyday forms of narrative. Additionally, when taking a cognitive science approach to computationally modeling narrative, it is important to attend not only to canonical forms of narrative, such as produced in literature, but also to everyday forms of narrative exchanged in social groups such as narratives of personal experience and life stories. Such everyday forms of narrative are common objects of study in the field of sociolinguistics [10, 15, 12]. Here, we augment such research with insights from cognitive linguistics, computer science, and sociology of classification.

In this paper, we discuss the *Chimeria Social Narrative Interface* (*Chimeria-SN*), a narrative generation component of a larger system called *Chimeria*. *Chimeria* implements dynamic computational models of social group membership and narratives associated with group membership. Similar to other research using virtual environments and games to empirically study social phenomena, such as the game *Prom Week*, *The Restaurant Game*,

and experiments of the Virtual Human Interaction lab (VHL) [13, 14, 1], our aim is to provide a testbed for studying aspects of social and computational identity. Toward this end, we have constructed a computational environment in which narratives of social group membership can be simulated and analyzed for both social scientific understanding and creative expression. The underlying model, grounded in cognitive science accounts of categorization, is capable of representing issues such as naturalization (becoming a category member over time), marginalization (becoming a boundary category member), and passing (being a member of one category, but appearing to be a member of another). Broadly, *Chimeria* serves our aims of both evoking narrative experiences of social group membership and enabling the creation of such experiences by anyone.

2   Theoretical Framework

Chimeria was developed as a part of an ongoing research endeavor called the Advanced Identity Representation (AIR) Project. The AIR Project seeks to develop new models of social identity in computational media to be deployed in technologies like interactive narratives, videogames, and social networks. Social identity can be conveyed through "digital identities" [2] using avatars, social networking profile posts, images, and so on. However, digital identities are limited technically in their expressivity and seldom explicitly mitigate against or model identity-related social ills (e.g., prejudices, stereotypes, etc.) [3, 5, 8], a topic addressed by the AIR Project. Below, we describe relevant research that undergirds the *Chimeria* system.

In [5], it is argued that in many forms of everyday communication, narrative provides a deep and satisfying sense of involvement. Sociolinguist William Labov [10], conducted empirical studies of narratives of personal experience, which can be formally represented as in [4]. Sociolinguist Charlotte Linde built on this work to relate narrative to social identity with "life stories" [12]. Since many everyday forms of narrative are now externalized through social media, we take a data-driven approach utilizing social networking profiles as a site where narratives of personal experience and life stories are performed. Social networking profiles are important sites for both generating expressive narrative content, and for the analysis of social categorization phenomena. *Chimeria* relates a formalization of a sociolinguistics model of narrative to social identity, implemented as a succession of posts on social networking profiles that collectively convey personal experiences related to social group membership.

Chimeria's generated narrative, implemented in *Chimeria-SN*, is based upon sociolinguist Livia Polanyi's model of narratives in conversational storytelling. In particular, we model what Polanyi defines as "story sequences" [15], in which multiple members in the conversation contribute individual stories towards the construction of a single, overarching narrative. Polanyi identifies two types of constraints in such conversational narratives: 1) linguistic constraints, which outline narrative structure using a past time storyworld with main line event clauses and contextualizing state clauses and 2) contextual constraints, which focus on making the narrative relevant, coherent, and accessible to recipients (recipient-design). In *Chimeria*, we used these constraints to inform construction of the model implemented in our narrative generation system, adapted to fit the context of a social network.

As a data-driven application of our model of group membership, we use *musical identity* (e.g., being a fan of a certain genre) as a test case. The music that people listen to is a vehicle for conveying "Music In Identities (MII)", wherein music is viewed as a "means for developing other aspects of our personal identities, including gender identity; youth identity; national identity; and disability and identity" [6]. Our test case narrates changes of social group membership related to musical identity as expressed via preferences in a social network.

3   Implementation

Chimeria dynamically models group membership and marginalization, and presents narratives generated from that model in a novel social networking interface. It consists of two components: (1) the *Chimeria Engine*: a dynamic algorithmic model of users' degrees of membership in multiple groups, and (2) the *Chimeria Social Narrative Interface* (*Chimeria-SN*): a narrative social networking interface for expressing experiences of membership and marginalization in social groups as represented using social media.

The *Chimeria Engine* models users' category memberships as gradient values in relation to the membership values of more central members, enabling more representational nuance than binary statuses of member/nonmember [3, 9, 11]. These are calculated from music artist "likes" (binary indications of positive valuation) on the user's Facebook profile, from which we extrapolate artists' moods (e.g., cheerful, gloomy, etc.), themes (e.g., adventure, rebellion, etc.), and styles (e.g., film score), which are used to express the identity of the user.

Chimeria-SN is a streamlined, aestheticized social networking interface, consisting of a dynamic collage of photos representing the user's musical taste preferences (Fig. 1), and a feed of posts that appear in an adjacent vertical timeline (Fig. 2). *Chimeria-SN* generates narratives in a simulated social networking environment that incorporates aspects of a user's real-world identity collected from a user's Facebook profile (e.g., real name, pictures, wall-posts) using the Facebook Graph API. The system reacts to the user by generating interaction posts from computer-controlled users who make up the user's simulated social circle.

Figure 1 A Screenshot of Chimeria-SN.

Figure 2 A Sample Chimeria-SN Wall Post.

Figure 3 The Chimeria System Architecture.

Fig. 3 gives a system overview and outlines the process by which *Chimeria-Engine* retrieves moods, themes and styles associated with music artists and genres. *Chimeria-Engine* first uses the user's music artist "likes" to find their preferred (e.g., Pop/Rock), and oppositional (imitation category, e.g., Jazz) genres. The user's initial genre group membership is calculated using the overlap of common moods and themes between the user's music artist

"likes" and the genre's. User actions (such as liking a post from a particular genre) causes the *Chimeria-Engine* to modify membership dependent on the intensity of the action.

3.1   Sample Narrative

The story in Figure 4 below was generated for a user that has the following Facebook music likes: *Ke$ha*, *Taylor Swift*, and *Justin Bieber*. The "Author Legend" represents the group membership of post authors, while the "Wildcard Legend" indicates dynamic insertions by *Chimeria-SN*. Each post appears on the user's wall in a manner similar to that shown in Figure 2.

Author Legend: N (Neutral Author), P (Pop/Rock Member Author), J (Jazz Member Author)

Wildcard Legend: ■ (Artist Wildcard), ■ (Genre Wildcard), ■ (Mood or Theme Wildcard)

Post: Good to see you on Chimeria! You'll be prompted with a series of posts like this one. Your actions in subsequent posts will determine the course of the story. Good luck, and have fun! (N, Story Entrance)

Post: Check out this Wayne Shorter music video.† (J, Sub-Story Entrance), Post: Wayne Shorter is wonderfully amiable. ;) (J, Contextualizing State), Post: Please tell me you enjoyed that Wayne Shorter song.* (Liked) (J, Main Line Event), Post: So you're a closet jazz lover. (P, Sub-Story Exit)

Post: Eric Clapton always makes my day better! (P, Sub-Story Entrance), Post: Eric Clapton makes great passionate music! (P, Contextualizing State), Post: I don't understand how anybody could not like Eric Clapton.*† (Disliked) (P, Main Line Event), Post: Way to back down on your pop rock roots. (P, Sub-Story Exit), Post: You seem meandering lately. (N, Sub-Story Exit)

Post: You know who's great? Keith Jarrett. (J, Sub-Story Entrance), Post: Keith Jarrett is too brooding for my taste. (P, Contextualizing State), Post: Here's a YouTube recommendation from yours truly..*† (Liked) (J, Main Line Event), Post: I never pegged you for a jazz nut. (P, Sub-Story Exit)

Post: I feel like listening to Van Halen. (P, Sub-Story Entrance), Post: Van Halen is so awfully energetic. (J, Contextualizing State), Post: Found this while surfing YouTube, it's awesome :)*† (Disliked) (P, Main Line Event), Post: You said you were into pop rock you poser. (P, Sub-Story Exit), Post: You seem atmospheric lately. (N, Sub-Story Exit)

Post: Well that's odd. You said you liked Ke$ha, but now you're a huge fan of jazz. (N, Story Exit)

* Dialog box with two buttons for the user to "Like" or "Dislike" this post
† Dialog box containing an embedded YouTube video (always a music video of the topic artist)

■ **Figure 4** Chimeria Sample Narrative.

Narrative generation uses a narrative structure specified by a simplified finite state machine, called the Linear Event Structure Machine (or probabilistic bounded transition machine) [5, 7]. Narrative structures are instantiated by a database of narrative templates (content-clauses), each filled in with artist, genre, theme, mood, and style content describing musical items from the Rovi Cloud Services API. An example clause template is:

```
<main-line-event-clause>
  <id>102</id>
  <author>original</author>
  <category-membership-test>any</category-membership-test>
  <naturalization-trajectory-test>any</naturalization-trajectory-test>
  <content>What did you think of that *mO *gO song?</content>
  <intensity>10</intensity>
</main-line-event-clause>
```

The 'id' uniquely identifies the content, and 'author' indicates the originating category. The 'category-membership-test' specifies the current degree of gradient membership within a social group the content-clause narrates (e.g., central, peripheral, or non-membership). The 'naturalization-trajectory-test' specifies trajectories of social group membership across all groups the content-clause narrates. The 'content' is the exact text of the post to be displayed, including wildcard (indicated by a "*" prefix) replacements using element types of artist, genre, theme, mood, or styles from the Rovi Cloud Services API. "O", "U", or "P" in

wildcards references a user's original social group, current profile, or imitation social group (internally referred to as a Passing social group) respectively.

Furthermore, topics are initiated by sub-story entrance clauses (as seen in the sample narrative in Fig. 4), which can then later be referenced using the *topic wildcard. Since topics themselves can be wildcards (e.g., *aP for a music artist from the imitation social group) any contextual information about them is retrieved at run-time (using wildcards such as *topic-m and *topic-video to find moods and YouTube music videos respectively).

4   Discussion

Building upon Polanyi's work, with *Chimeria-SN* we have developed a model of interactive conversational narrative. Some adaptations were necessary in order for the model to be computationally implementable and so that it could be used in an interactive framework. Table 1 contains a comparative summary between *Chimeria*'s narrative model and Polanyi's. *Chimeria*'s grounded narrative model extends Polanyi's model by involving the user in an ongoing social network narrative adapted to user taste, which necessarily substitutes virtual affordances for physical ones.

Table 1 Comparison between *Chimeria* & Polanyi's models of Conversational Narrative.

Component	Parallels Polanyi's Model	Diverges from Polanyi's Model
Linguistic constraints	Event propositions occur at unique discrete moments. Structured using main line event clauses, contextualizing state clauses and evaluative meta-information. Stories have a "point" (e.g., tale of imitation).	User responses to posts directly affect a dynamic narrative (a narrative referring to both the past and present time in the storyworld). In other words, the user has agency regarding story trajectory.
Contextual constraints	Story is relevant and recipient-designed for the user based on expression of musical taste. Stories consist of entrance and exit clauses for overall coherence.	Story recipients acknowledge tellings through posts on the *Chimeria-SN* wall rather than physically.
Story Sequence & Sub-stories	Sub-stories are individual, self-contained stories from multiple storytellers (but one primary narrator for each sub-story), which collectively form a story sequence. Evaluation is internal to the storyworld clauses.	Occurs in virtual and simulated social network (e.g., affordances available in wall posts like emotes, punctuation, videos etc.) instead of in physical space (e.g., body language, tonality, etc.) which includes a larger set of story recipients in general.

Chimeria-SN's narrative model is grounded in Polanyi's results. This means that we must strive for a high degree of fidelity between our implementation and Polanyi's empirical results, with only necessary adaptations driven by the differences between real-world conversation and our social networking domain. Furthermore, our model should continuously be reconciled with the latest empirical sociolinguistics results on conversational narrative in social networks.

5   Concluding Reflections

Social group memberships are important aspects of societies. Stories of social group membership are important for constituting our social fabrics. As stated above, robust model of social group membership can be an important aspect for modeling everyday forms of narrative. Reciprocally, narrative generation can be an effective means of conveying a dynamic model of social group membership for both research and applications such as interactive narratives and videogames. *Chimeria* implements dynamic computational models of social group membership and conversational narratives, and we hope that it provides a useful testbed

in which narratives of social group membership can be both simulated and analyzed. An advantage of our approach is that in future work, we believe that social group memberships of users could be extrapolated from analyzing narratives in social networks, the inverse of our current aim of synthesizing narratives from a model of social group membership dynamics. Our longer term hope is that, by computationally modeling issues such as naturalization, marginalization, and passing, we can contribute to scientific approaches to issues of social empowerment and diversity most often served by research in the humanities, arts, and social sciences.

References

1 Jeremy N. Bailenson. Doppelgängers: A new form of self. *The Psychologist*, 25(1):36–38, 2012.

2 Nick Bostrom and Anders Sandberg. The future of identity, (2011) report commissioned by the united kingdom's government office for science, 2011.

3 Geoffrey C. Bowker and Susan Leigh Star. Sorting things out: Classification and its consequences. *Inside Technology*, 1999.

4 Joseph Goguen. Notes on narrative. `http://charlotte.ucsd.edu/~goguen/courses/275f00/narr.html`, 2001.

5 Joseph A. Goguen and D. Fox Harrell. Style: A computational and conceptual blending-based approach. In Shlomo Argamon, Kevin Burns, and Shlomo Dubnov, editors, *The Structure of Style, Algorithmic Approaches to Understanding Manner and Meaning*, pages 291–316. Springer, 2010.

6 David J. Hargreaves, Dorothy Miell, and Raymond A. R. MacDonald. What are musical identities, and why are they important. In Raymond A. R. MacDonald, David J. Hargreaves, and Dorothy Miell, editors, *Musical identities*, pages 1–20. Oxford University Press, 2002.

7 D. Fox Harrell. Walking blues changes undersea: Imaginative narrative in interactive poetry generation with the griot system. In Hugo Liu and Rada Mihalcea, editors, *Computational Aesthetics: Artificial Intelligence Approaches to Beauty and Happiness*, number WS-06-04 in AAAI Technical Reports, pages 61–69, 2006.

8 D. Fox Harrell. Designing empowering and critical identities in social computing and gaming. *CoDesign: International Journal of CoCreation in Design and the Arts*, 6(4):187–206, 2010.

9 D. Fox Harrell. Toward a theory of critical computing. *CTheory, Code Drift: Essays in Critical Digital Studies*, cds006, 2010.

10 William Labov. *Sociolinguistic patterns*, volume 4. Philadelphia: University of Pennsylvania Press, 1972.

11 George Lakoff. *Women, fire, and dangerous things: What categories reveal about the mind.* Cambridge University Press, 1990.

12 Charlotte Linde. *Life stories.* Oxford University Press, 1993.

13 Josh McCoy, Mike Treanor, Ben Samuel, Brandon Tearse, Michael Mateas, and Noah Wardrip-Fruin. Authoring game-based interactive narrative using social games and *Comme il Faut.* Presented at the 4th International Conference & Festival of the Electronic Literature Organization, 2010.

14 Jeff Orkin and Deb Roy. The restaurant game: Learning social behavior and language from thousands of players online. *Journal of Game Development*, 3:39–60, 2007.

15 Livia Polanyi. *Telling the American story: A structural and cultural analysis of conversational storytelling.* MIT Press, 1989.

Narrative Similarity as Common Summary

Elektra Kypridemou and Loizos Michael

School of Pure and Applied Sciences
Open University of Cyprus
Nicosia, Cyprus
`elektra.kypridemou@st.ouc.ac.cy, loizos@ouc.ac.cy`

Abstract

The ability to identify similarities between narratives has been argued to be central in human interactions. Previous work that sought to formalize this task has hypothesized that narrative similarity can be equated to the existence of a common summary between the narratives involved. We offer tangible psychological evidence in support of this hypothesis. Human participants in our empirical study were presented with triples of stories, and were asked to rate: (i) the degree of similarity between story A and story B; (ii) the appropriateness of story C as a summary of story A; (iii) the appropriateness of story C as a summary of story B. The story triples were selected systematically to span the space of their possible interrelations. Empirical evidence gathered from this study overwhelmingly supports the position that the higher the latter two ratings are, the higher the first rating also is. Thus, while this work does not purport to formally define either of the two tasks involved, it does argue that one can be meaningfully reduced to the other.

1 Introduction

Stories play a central role in human knowledge, understanding and reasoning, and are key in explaining human behavior and social communication [1, 24].

According to Schank and Abelson [24], human knowledge is functional, in the sense that all knowledge is encoded as stories along with mechanisms to construct, store and retrieve them. Hence, human memory is a collection of stories that we experienced, heard or composed in the past. Searching in memory is like searching for stories within one's own collection of stories. Moreover, existing stories in memory form our beliefs and are the ingredients of new ideas.

Following that view, understanding and behavior depends upon our past experiences, stored as old stories. In order to understand a new story, we have to find an old story that is *similar* to the new one. Thus, understanding, for a listener, means mapping the speaker's stories onto the listener's stories. This is why different people may interpret the same story in different ways. Analogously, when we face a new unfamiliar situation, we have to determine which of the old situations — stored as stories — in our memory is most relevant to the situation at hand. The familiar situation that is most similar to the new one determines our behavior in the new situation. As a result, intelligent behavior lies upon our capability of finding a relevant past experience that will help us make sense of a new experience [24].

But how does one determine the extent to which stories are considered to be similar?

1.1   Similarity: Existing Approaches

Similarity is a very powerful construct in psychology, entering into the analysis of many diverse phenomena, such as creative and scientific discovery, problem-solving, categorization, decision-making, learning and transfer [10, 27]. In this section we briefly describe several approaches to similarity as proposed in the literature.

i) The **geometric approach** determines similarity using the *mental distance models* (see, e.g., [20, 25, 26]). According to that view, concepts are represented as points within a multi-dimensional mental space and similarity between concepts as the inverse of the metric distance between these points. Thus, the less the distance between the point representations of two concepts A and B, the more similar the concepts are. As the distance between two points A and B within the space increases, it becomes easier for us to detect that the concepts A and B are different.

The geometric approach treats similarity as a symmetric relation, in the sense that the similarity of A to B equals the similarity of B to A. However, Tversky [27] argued against the symmetry assumption, providing empirical evidence for asymmetric similarities, and proposed his own model for similarity.

ii) The **featural approach**, represented by Tversky's classic *contrast model* [27], proposes that concepts are represented as selections of features, and similarity is described as a feature-matching process. Specifically, the similarity between two concepts is computed based on numbers of shared and not shared features (expressed as a linear combination of the measures of their common and distinct features). The greater the size of the set of the common features of two concepts A and B, and the smaller the sizes of the two sets of distinct features $A \setminus B$ and $B \setminus A$, the greater the similarity between the two concepts A and B. This model also accounts for the asymmetry assumption of similarity, since the negative effects of the two complement sets $A \setminus B$ and $B \setminus A$ are not equal. Hence, if we ask how similar A is to B, then the set $B \setminus A$ counts much more than the set $A \setminus B$.

Later on, Gentner [7] indicated that although Tversky's contrast model seemed to be correct for literal similarity comparisons, it did not provide a good account for other types of comparisons, such as analogies. Hence, Gentner [7] introduced the following approach, providing a theoretical framework for analogy.

iii) The **structural (or relational) approach** is based on Gentner's work on *Structure Mapping Theory (SMT)* [7, 8, 10]. Two concepts, that function as wholes, may be treated as analogous when they share some essential relations, even if they may have a lot of distinct features. Because of this singularity of analogy, structure-mapping, in contrast to the featural approach, treats commonalities and differences as dependent features. Concepts are considered as structures of object attributes and relations, and comparisons are made by mapping the structures of the two concepts. The main idea of this theory is that analogy is characterized by mapping relations between objects, rather than merely mapping the attributes of the objects. Thus, comparisons rely only on the syntax of the knowledge representations and not on their specific content. Furthermore, the structural approach allows analogies to be distinguished clearly from literal similarity and other kinds of concept comparisons, as we shall discuss in more detail in Section 2.2.

iv) A more recent approach, the **tansformational approach** [11, 12] considers concepts as representations, and suggests that similarity depends on the ease of transformation between these representations. Specifically, the less the number of steps needed to transform a representation into another representation, the higher the similarity between the two concepts is. Thus, the more dissimilar the entities, the more transformations are needed. Hahn *et al.* empirically tested the view of similarity as transformational distance [11, 12],

and proposed the *Representational Distortion* as a specific example of this approach [11]. However, Larkey and Markman, while testing the similarity judgments for geometric objects, found some evidence against this approach [14].

1.2 Computational Modelling of Narratives

Previous attempts to create computational models of narratives include Propp's narrative functions (or narratemes) [21], Rumelhart's Story Grammars [22], Lehnert's Plot Units [15, 16], Löwe's Doxastic Preference Framework (DPF) [17], Elson's Story Intention Graphs (SIG) [4] and Chambers' Narrative Event Chains (NEC) [2]. Moreover, several recent computational studies provided algorithms that manage to make story comparisons [5, 13] and recognize narrative similarity [18] and analogy [3].

The current study attempts to define similarity of narratives by means of the concept of summary. To this end, it has been conjectured [18, 19] that *similarity between two stories is effectively equivalent to saying that the two stories have a common summary*; i.e., an abstraction that is appropriate for both stories. In particular, the more appropriate this common summary is for the two stories, the more similar the two stories are. This statement forms the main hypothesis that we seek to empirically examine in this work.

2 Background

Our investigation of the above hypothesis does not presuppose any particular framework of story understanding, nor any particular approach to defining similarity. Nonetheless, we find it useful to adopt certain notions and terminology from the SMT [7, 8] and to compare some of our obtained results to those obtained under the SMT. We shall, thus, present the SMT in more detail in this section.

2.1 Structure Mapping Theory for Analogy

When we interpret analogies such as "A is like B", we draw inferences about a concept A (target), based on our knowledge of another concept B (base), which serves as the source of our knowledge. The central idea of the SMT [7, 8] is the definition of the analogy as "an assertion that a relational structure that normally applies in one concept can be applied in another concept" [7]. Thus, the essence of an analogy between two concepts A and B is that they share a common structure. This structure is the dominant aspect of concepts A and B, even though these concepts may differ in many other aspects. For example, we might say that a child's mind is like a sponge. One could easily interpret this analogy, drawing the inference that a child's mind absorbs a lot of information. Hence, we use our knowledge about sponges in order to draw inferences about a child's mind. However, we do not transfer all our knowledge about sponges to the child's mind. If we did so, then we could argue that the child's mind is yellow with holes and holds water. But we do not! Although the child's mind and the sponge do not share other common features, this does not seem to count against the analogy. This is why Gentner criticized Tversky's *contrast model* [27] as not appropriate for analogies [7].

Continuing with the sponge-mind example, it seems that people, seeking to identify analogies, somehow know which features must be transferred and which not. The SMT is capable to explain such behavior, by suggesting that i) we tend to focus on relational information and ignore the distinctive attributes of objects in A and B ii) we prefer to focus

Table 1 Similarity types according to Gentner's work [7, 8, 9].

Similarity type	Object Attributes	Relational Predicates	
		Low-order	Higher-order
Literal Similarity (LS)	✓	✓	✓
Analogy (ANA)	×	✓	✓
Surface Similarity (SS) or Mere Appearance	✓	✓	×
First Order Relations (FOR)	×	✓	×
Objects Only (OO)	✓	×	×
Anomaly (ANO)	×	×	×

on interconnected systems of relations and favor these relations in our interpretations (see the mapping principles as described later in this section).

According to the structural approach, knowledge is represented as a propositional network of nodes and predicates [7, 23]. Nodes represent concepts, while predicates express propositions about these concepts. Predicates may be either object attributes, taking one argument (e.g., YELLOW(x)), or relationships, taking two or more arguments (e.g., SMALLER(x, y)). A further syntactic distinction of predicates is also made: first-order predicates take objects as arguments (e.g., HURT(y, x) and HATE(x, y)), while higher-order predicates take propositions as arguments (e.g., CAUSE[HURT(y, x), HATE(x, y)]) [7].

Using these notions, we briefly introduce the two main mapping principles of the SMT: i) relations, rather than object attributes, are mapped from the base to the target concept, and ii) a relation that belongs to a system of mutually interconnected relations is more likely to be mapped than an isolate relation (known as the *Systematicity Principle*) [7].

2.2 Similarity Types

The SMT clearly distinguishes analogy from other kinds of concept comparisons. Table 1 summarizes several similarity types, proposed by Gentner *et al.* [7, 8, 9], based on the possible combinations of the predicates that two concepts may share.

To apply Gentner's work to story comparisons, Gentner, Rattermann and Forbus [9] considered i) object attributes as the characters, objects and locations of the story, ii) first-order relations as the events, actions and other relations between the objects of the story (e.g., X talking to Y), and iii) the higher-order relational structure as the causal (or other types of) relations in the story's plot. According to these assumptions, they created pairs of stories by generating several versions of a base story, which differed in the level of similarity they shared with the original story.

In the present study we borrow the names of the similarity types proposed by Gentner (Table 1). However, while applying these types into story pairs, we interpret predicates in a slightly different manner than Gentner *et al.* [9] did (see Section 3.2.2).

3 Empirical Method

In order to test our main hypothesis that story similarity relates to the existence of a common summary, we designed an online questionnaire and we invited people to rate pairs of stories in terms of their similarity or in terms of how appropriate a third story was as a summary of each of the first two stories.

3.1 Participants

Individuals were invited via e-mail to participate voluntarily to the study, by completing an online questionnaire. No financial or other compensation was given for participation. Although more than a hundred people started completing the questionnaire, participants who left before answering at least one question of the main part of the experiment were excluded from our sample. Hence, the final sample of our study comprised 52 adults (21 male, 31 female), aged 18 to 65 years (one reported age 66+). Participants were mainly residents of Greece (26) and Cyprus (22) and they all spoke Greek as their native language.

3.2 Measures

Before proceeding to the detailed description of the various stages of the questionnaire construction, we give a brief overview of the procedure followed, and we introduce the notation that we will use. To prevent any later confusion, we find it useful to discriminate early between the **groups** of stories initially constructed, the **trials** of the questionnaire, and the **triples** of stories used during the analysis stage.

First, we created a pool of selected narratives, serving as the source for our stories (Section 3.2.1). We then edited some of these stories in order to form pairs (A, B) of stories $S \in \{A, B\}$ with varying degrees of similarity (Section 3.2.2). For each story pair (A, B) we also generated 4 other stories $C_{i,j} \in \{C_{11}, C_{10}, C_{01}, C_{00}\}$, which we considered as summaries of stories A and B (Section 3.2.3). As a result, during the story preparation stage, we formed 16 **groups** comprising 6 stories each $(A, B, C_{11}, C_{10}, C_{01}, C_{00})$. These groups of stories served as a source for the 16 **trials** of the questionnaire (Section 3.2.4), where each trial comprised 4 stories $(A, B, C_{11}$, along with one of $C_{10}, C_{01}, C_{00})$ selected from its corresponding group. Later on, in order to test our main hypothesis (Section 4.2), for each trial of the questionnaire we formed two **triples** $(A, B, C_{i,j})$, one with $C_{i,j} = C_{11}$ and one with $C_{i,j} \in \{C_{10}, C_{01}, C_{00}\}$, resulting in a total of 32 triples for each participant who fully completed the questionnaire. For participants who completed the questionnaire only partially, we gathered as many triples as possible given the participants' responses.

3.2.1 Selection of a Pool of Original Stories

Since we wanted the stories of the experiment to be as naturalistic as possible and not to appear like artificial stories made especially for use in the lab, we decided to mainly use existing stories, rather than to produce our own. In order to find such original stories, we searched for literature books and collections of myths and fairy tales in the library. We also searched for online videos and texts of narrations, and we recorded oral narrations.

During the selection of the original stories, we mainly took into account two factors: i) the reputation of the story, and ii) the content of the story. Regarding the first factor, to ensure that the participants' previous knowledge of the stories would not bias their responses, we avoided using well-known stories (e.g., the French novella "The little prince" of Antoine de Saint-Exupéry). Concerning the second factor, we gave preference to stories with didactic or entertaining content, which we considered as interesting enough for the participants to read. Having in mind the fact that participation to the questionnaire was a voluntary and time consuming process, such a restriction of the space of different story types was necessary to motivate people to fill in the questionnaire and also to prevent them from leaving the process early. Hence, we avoided stories that merely described a sequence of events without any further purport, and favored those where the author's intention was to pass a deeper message, or even entertain the reader, through the narration. Given these constraints, we

created a pool of dozens of original stories, or excerpts of them, both worth reading and not well known.

3.2.2   Generation of Story Pairs According to Similarity Types

Following the story selection, we edited the original stories by adjusting their length. Edited stories were (i) not too long for the participants to read, but also (ii) long enough, in order to keep the style of the original version and not to appear artificial.

We then created story pairs (A,B) with varying degrees of similarity, inspired by the similarity types used in previous work [7, 8, 9] based on the SMT. However, since our aim was to cover the entire range of similarity degrees between stories A and B, rather than to contrast the several similarity types proposed in the literature, we only used Literal Similarity (LS), Analogy (ANA), Surface Similarity (SS) and Anomaly (ANO). Consequently, we defined the variable *sim-type*(A, B), taking values in $\{LS, ANA, SS, ANO\}$, to indicate the similarity degree assumed by the experimenters for each story pair (A,B).

To create story pairs of the above types of similarity, we interpreted the SMT's predicates (cf. Table 1) as follows: i) object attributes served as the *characters* of the stories, ii) first-order relations served as *relations between the characters*, which could either state a relationship (e.g., X neighbor of Y) or an emotional relation (e.g., X loves Y), while iii) higher-order relations served as the *overall structure of the plot*, in which the relationships and the main interactions between the characters and their actions are represented as an interconnected system (e.g., X lies to Y repeatedly; someday X really needs Y's help; Y do not believe X's need; Y leave X helpless; great disaster happens to X).

Given the above interpretations, we created 4 pairs of stories (A,B) for each of the 4 similarity types, resulting in a total of 16 pairs. To do so we chose a story A from the pool of our selected original stories and we paired it with another story B, which was either generated by the experimenters or, preferably, also extracted from the pool, so as to match the constraints of each similarity type.

Story pairs of the LS type comprised two versions of practically the same story. Hence, we produced story B by expressing story A using different wording, and changing one or two minor details having little or no importance for the overall meaning of the story (e.g., iron wheel vs. stone wheel, for a story where the wheel weight was essential, rather than its material).

Pairs belonging to the ANA type were stories with different characters, where the same basic relations held between them. In the most dissimilar cases of ANA pairs, the overall structures of the two stories differed in some isolated relations, but never in the most interconnected ones.

For example, the structure described earlier (X lies to Y repeatedly; someday X really needs Y's help; Y do not believe X's need; Y leave X helpless; great disaster happens to X) was a common structure of an ANA pair. For story A, X was a shepherd and Y were his fellow villagers, while for story B, X was a grandmother and Y were her children and grandchildren. Some other elements which differed between the stories were: the lie itself, the outcome of lying, the real need of X, and the disaster that happened to X. For example, in story A, the lie was that a wolf attacks the sheep, the outcome was to laugh at the villagers who went to help, the need appeared when a wolf really attacked the sheep, and the disaster was the loss of the sheep. In story B, the lie was that the old woman pretends to be very sick, the outcome was to gain her family's company by 'forcing' them to visit her, the need was when the old woman was really feeling sick, and the disaster was the woman's death.

■ **Table 2** Guidelines used for the construction of the story pairs (A,B) and their summaries $C_{i,j}$ for each similarity type $sim\text{-}type(A, B) \in \{LS, ANA, SS, ANO\}$.

	(A,B)	C_{11}	C_{10} (resp., C_{01})	C_{00}
LS	Same stories with few minor details differing.	Specific non-shared detail(s) of A and B missing.	C_{11} plus the detail(s) of story A (resp., B).	C_{11} with a key element of the plot changed or missing.
ANA	Stories share a common structure of their plot except for some isolated relations. Different characters.	The common structure of A and B. Characters expressed in an abstract way compatible with both A and B.	C_{11} plus one of the missing isolated relations specified only in A (resp., B).	C_{11} with a key element of the plot missing or replaced to be incompatible with both A and B.
SS	Identical stories up to one point. Different endings.	A summary of the common part of A and B. Ending is missing.	C_{11} plus a summary of the ending of story A (resp., B).	C_{11} plus an ending different than both A's and B's.
ANO	Different stories with a common extremely abstract structure.	The common abstract structure of A and B.	C_{11} plus one of the abstract elements specified in A (resp., B).	C_{11} with one of the elements changed to be incompatible with both A and B.

In order to create story B for the ANA pairs, we searched for stories in the pool and, if we could not find any appropriate story, we generated our own B.

For the SS story pairs, we chose a story A and we generated a story B, identical to A up to some point, but with a completely different ending. The ending was edited in a way that was crucial for the overall meaning of the story (e.g., at the end the king admired and appointed the hero as his main advisor vs. the king disapproved the hero's behavior and banished him: the choice characterizes the whole behavior of the hero as good vs. bad).

Finally, the stories of the ANO pairs were completely different stories that shared some common structure. This common structure was extremely abstract, so that it missed most of the key elements for understanding the overall meaning of the stories (e.g., an older man who lies reveals a secret to a younger man). Concrete characters and relations between the characters (e.g., the secret, the reason the man lies, the reason the young man asks for the secret, and the relation of the two men) differed between the two stories. Similarly to ANA types, we generated our own stories B for the ANO pairs when we could not find any appropriate story B in the pool.

Overall, 20 out of 32 stories of the questionnaire came from the pool of selected original stories (16 as story A, 4 as story B). The remaining B stories were either edited versions of the corresponding original stories A (for the LS and SS types), or stories generated by the experimenters to meet the specific criteria of the ANA and ANO similarity types. Most of the original stories used (14 out of 20) were found in books already translated in Greek, with 6 of them being fairy tales and myths from all over the world (Japan, China, Arabia, Lithuania, Burma, Hungary). Among the remaining original stories, 3 came from online texts, 2 from online videos, and 1 from oral narration.

3.2.3   Generation of Summaries According to Summary Types

For each of the 16 story pairs (A,B) we created 4 different summaries $C_{i,j} \in \{C_{11}, C_{10}, C_{01}, C_{00}\}$, and for each triple ($A$,$B$,$C_{i,j}$,) we defined the variable $sum\text{-}type(C_{i,j}, S)$, to indicate the appropriateness of summary $C_{i,j}$ for story $S \in \{A, B\}$. Our intention here was to cover the range of possible summaries in a semi-systematic manner. Thus, we created i) a 'common' summary C_{11}, with the aim of being an appropriate summary for both the stories A and B, ii) a summary C_{10}, with the aim of being appropriate for story A and inappropriate for story B, iii) a summary C_{01}, with the aim of being inappropriate for story A and appropriate for story B, and finally, iv) a summary C_{00}, with the aim of being inappropriate for both the stories A and B.

Given the above definition of $C_{i,j}$, the variable $sum\text{-}type(C_{i,j}, A)$, indicating the appropriateness of summary $C_{i,j}$ for story A, had the value 'good' if $i = 1$, and the value 'bad' if $i = 0$. Accordingly, the variable $sum\text{-}type(C_{i,j}, B)$, indicating the appropriateness of summary $C_{i,j}$ for story B, had the value 'good' if $j = 1$, and the value 'bad' if $j = 0$. Hence, i can be seen as an index of the appropriateness of summary $C_{i,j}$ for story A, and j as an index of its appropriateness for story B. In general, the variable $sum\text{-}type(C_{i,j}, S)$ had the value 'good' if either ($S = A$ and $i = 1$) or ($S = B$ and $j = 1$), and the value 'bad' otherwise.

We considered a summary $C_{i,j}$ as inappropriate, when some essential feature of the story S was either missing or changed in the summary. However, appropriate summaries could include an unimportant detail of a story as well. Thus, we did not consider the succinctness of the summary as a defining characteristic of the appropriateness of the summary. Table 2 describes the methodology we adopted in order to create the 4 summaries $C_{i,j}$, for each of the 4 similarity types. C_{11} served as a baseline for creating the other 3 summaries. For C_{10} and C_{01} we modified C_{11} so as to be slightly more appropriate than C_{11} for one of the two stories and slightly less appropriate for the other story, while for C_{00} we modified C_{11} so as to be slightly less appropriate than C_{11} for both A and B. Our concern was to avoid creating summaries that would be *obviously* inappropriate for one of the two stories. This choice was guided by the fact that our main hypothesis states that the existence of an appropriate common summary indicates a high similarity pair. Consequently, low values of summary scores would be less useful in validating or falsifying this hypothesis during our analysis.

3.2.4   Questionnaire Construction

Each trial of the questionnaire included the presentation of 4 stories (A, B, their common summary C_{11}, and one of the non-common summaries), along with 5 questions, asking participants to rate the similarity between stories A and B, and the appropriateness of the other two stories as summaries for each of the stories A and B.

In the place of the non-common summary we chose C_{00} to appear in half of the cases, and C_{10} and C_{01} to appear in one fourth of the cases each. To achieve these frequencies, we created quadruples (C_{10}, C_{01}, C_{00}, C_{00}) and we used the Latin Squares method to counterbalance their order. We created 4 conditions (Table 3) and we randomly allocated participants to each of these conditions (see Section 3.3). As a result, for each of the 4 similarity types, in addition to the common summary C_{11}, each participant was presented with each of summaries C_{10} and C_{01} once, and with summary C_{00} twice. Moreover, the order of the 16 trials of the questionnaire, as well as the order of the two summaries of each trial, was randomized.

Using the open source survey application LimeSurvey 2.00+, we implemented the above methods by constructing an online questionnaire, which we describe in more detail next.

■ **Table 3** Summaries presented across the four Conditions (in addition to C_{11}, which was presented in each Condition).

Trial	Similarity type	Condition				Trial	Similarity type	Condition			
		1	**2**	**3**	**4**			**1**	**2**	**3**	**4**
1	LS	C_{10}	C_{01}	C_{00}	C_{00}	9	SS	C_{10}	C_{01}	C_{00}	C_{00}
2	LS	C_{01}	C_{00}	C_{00}	C_{10}	10	SS	C_{01}	C_{00}	C_{00}	C_{10}
3	LS	C_{00}	C_{00}	C_{10}	C_{01}	11	SS	C_{00}	C_{00}	C_{10}	C_{01}
4	LS	C_{00}	C_{10}	C_{01}	C_{00}	12	SS	C_{00}	C_{10}	C_{01}	C_{00}
5	ANA	C_{10}	C_{01}	C_{00}	C_{00}	13	ANO	C_{10}	C_{01}	C_{00}	C_{00}
6	ANA	C_{01}	C_{00}	C_{00}	C_{10}	14	ANO	C_{01}	C_{00}	C_{00}	C_{10}
7	ANA	C_{00}	C_{00}	C_{10}	C_{01}	15	ANO	C_{00}	C_{00}	C_{10}	C_{01}
8	ANA	C_{00}	C_{10}	C_{01}	C_{00}	16	ANO	C_{00}	C_{10}	C_{01}	C_{00}

3.3 Procedure

An e-mail was sent to a number of people, inviting them to complete the questionnaire and to further disseminate the invitation. In order to encourage participation, in the invitation e-mail we informed the recipients that during completion of our on-line questionnaire they were going to read interesting stories and myths from all over the world. They were also informed that the study concerns the similarity of narratives and summarization, but they were given no further information on the specific hypothesis of our study.

Participants who entered the questionnaire were assured that responses were collected anonymously. After certifying that they are adults and native Greek speakers, they were given navigation instructions, and they were instructed how to safely interrupt the process in case they would wish to continue later. Afterwards, they provided information on their demographic characteristics (gender, age, educational level, country of residence) and they reported whether they enjoy reading stories, myths, fairy tales and literature, and how often they read such texts.

The practice section followed, where an example of a trial was given in order to clarify the process. Participants were told that they are going to read two stories, A and B, and rate their similarity in a scale from 0 (no similarity) to 5 (high similarity). For illustrative purposes, the stories A and B used in the example were considerably shorter than the stories of the experiment. Some help was given, where we stated explicitly the differences between the two stories and asked participants whether they consider these differences (along with others that they may have identified) as significant enough so as to rate the two stories with a low degree of similarity, or they consider the differences meaningless for the general meaning of the stories, and so they would rate the two stories with a high score. After registering their responses, we informed them that two more stories would appear on the screen in order to evaluate how good they consider each of the new stories as a summary for each of the stories A and B, in a scale from 0 (very bad summary) to 5 (very good summary). The stories A and B, along with two example summaries were presented. As previously, some help was provided, stating that some statements of the original stories are missing in the summaries, and asking participants to evaluate whether they consider these statements as important enough so as to rate the summary as bad, or not. Finally, participants were given the chance to read the instructions once more, in case the did not feel confident with the process. At no point in this practice section did we suggest what an appropriate score would be for the example questions.

Following the practice section, a screen with some guidelines appeared. Participants were first encouraged to answer according to their own judgment, since there are no "right" or "wrong" answers. They were then advised to rate summaries by considering which facts of the stories A and B are included and which are missing, and ignore possible syntactic or stylistic aspects of the stories. We also asked participants to be as consistent as possible in their ratings during the whole experiment. Finally we advised them not to judge a summary solely by its length, since length does not necessarily indicate a summary's appropriateness. After reading these guidelines, we asked them to make sure that they are in a quiet environment, since the main part of the experiment was about to start.

At the beginning of the main part of the experiment, each participant was randomly allocated to one of the four conditions of the experiment. Then, the first randomly selected trial started. Participants were blind to the labeling of the story pairs (A, B) according to the 4 similarity types, and the labeling of the summaries $C_{i,j}$ according to the 4 summary types. Hence, they could only see 4 stories presented on the screen, without having any cue about the experimenters' assumptions (i.e., $sim\text{-}type(A, B)$ and $sum\text{-}type(C_{i,j}, S)$). The first pair of stories (A, B) appeared and participants were asked to rate their similarity on a scale from 0 (no similarity) to 5 (high similarity). After registering their score, which was recorded as $sim\text{-}score(A, B)$, two summaries appeared on the screen in random order: the common summary C_{11}, along with one of the summaries C_{10}, C_{01}, C_{00}, according to the trial and the participant's condition (Table 3). Participants were asked to judge how good they considered each of the two summaries for each of the stories A and B, on a scale from 0 (very bad summary) to 5 (very good summary). These ratings were recorded as $sum\text{-}score(C_{i,j}, A)$ and $sum\text{-}score(C_{i,j}, B)$ respectively. As a result, for each trial we recorded 1 similarity score and 4 summary scores. After completing the first trial the participants proceeded to the next one and continued scoring as previously, until all 16 trials were completed. Finally, participants were thanked and were invited to leave comments about the survey.

4   Empirical Results

Although only 28 out of 52 participants fully completed the questionnaire, we decided to keep incomplete responses in our analysis, since the experimental design allowed us to do so (the trials of the questionnaire were randomly ordered and independent of one another). For those who fully completed the questionnaire the whole procedure lasted on average 87 minutes. The long duration indicates that participants spent enough time to read the instructions and that they read the stories carefully before they gave their scorings. Table 4 displays the demographic characteristics of our sample. The fact that most of the participants reported that they like reading stories and they do so frequently (see Table 4), indicates that our sample had the appropriate skills and experience for completing the task of our study.

Regarding our experimental design, our data suggests that participants were well distributed over the 4 conditions (12 in the first, 10 in the second, 16 in the third and 14 in the fourth condition). The Latin Squares method also worked efficiently since, in addition to the common summary C_{11} which was presented in all 551 trials, the second summary was chosen 25% (137/551) of the times to be summary C_{10}, another 25% (137/551) of the times to be summary C_{01}, and 50% (277/551) of the times to be summary C_{00} (see 'Overall' row, Table 5).

Table 4 Number of Participants (NOP) for each Age Range, Educational Level, Preference and Frequency of Reading Stories.

Age		Educational Level		Reading Stories		
Range	NOP	Degree	NOP		Like	Do not Like
18–20	2	High School	7	Never	0	1
21–30	19	Bachelor's	24	Rarely	5	6
31–40	9	Master's	6	Sometimes	14	0
41–50	6	Ph.D.	15	Often	20	0
51–60	10			Always	6	0
>60	6					

4.1 Participant Scores Compared to Experiment Type

As a preliminary analysis to validate our experimental setting, we investigated whether the types of similarity and summaries that we had considered, indeed matched the scores given by participants. Participant similarity ($sim\text{-}score(A, B)$) and summary ($sum\text{-}score(C_{i,j}, S)$) scores confirmed the experimenters' assumptions regarding the similarity ($sim\text{-}type(A, B)$) and summary ($sum\text{-}type(C_{i,j}, S)$) types, respectively. The similarity scores $sim\text{-}score(A, B)$ not only matched the similarity level assumed for each similarity type $sim\text{-}type(A, B)$, but also, considering the error bars of Figure 1 (Left), clearly differentiated the LS and ANA (average similarity scores 4.64 and 3.44, respectively) from the SS and ANO types (average similarity scores 2.07 and 1.64, respectively).

Moreover, the 4 summary types $sum\text{-}type(C_{i,j}, S)$ were also confirmed by the summary scores $sum\text{-}score(C_{i,j}, A)$ and $sum\text{-}score(C_{i,j}, B)$ given by the participants for the appropriateness of summary $C_{i,j}$ for stories A and B, respectively (Figure 1, Right). Participants judged the common summary C_{11} as appropriate for both stories A and B. They also judged summary C_{10} as more appropriate for story A than for B, and summary C_{01} as more appropriate for story B than for A. Finally, they judged summary C_{00} with a relatively low degree of appropriateness for both stories A and B (Figure 1, Right; Table 5, 'Overall'). Table 5 shows how the average summary scores $sum\text{-}score(C_{i,j}, S)$ are distributed across the 4 similarity types. It is noteworthy that in the case of the SS and ANO similarity types, scores given to the common summary C_{11} were considerably lower than the $sum\text{-}score(C_{10}, A)$ and $sum\text{-}score(C_{01}, B)$ scores, which is in accordance with our main hypothesis (see Section 5 for a discussion).

4.2 Testing our Main Hypothesis

Figure 2 already provides a first indication in support of our hypothesis, showing that pairs of stories that were of types $sim\text{-}type(A, B)$ that are considered more similar, were also associated with a better common summary C_{11}, as measured by the participants' average summary score $sum\text{-}score(C_{i,j}, S)$.

For the main part of our analysis, we ignored the labels given to the similarity and summary types by the experimenters, and we analyzed the relations between the triples of scores given by the participants for each triple of stories $(A,B,C_{i,j})$: i) the appropriateness of summary $C_{i,j}$ for story A ($sum\text{-}score(C_{i,j}, A)$), ii) the appropriateness of summary $C_{i,j}$ for story B ($sum\text{-}score(C_{i,j}, B)$), and iii) the similarity between stories A and B ($sim\text{-}score(A, B)$).

Figure 1 Participants' scorings matching the experimenters' labeling. **Left:** Average similarity scores $sim\text{-}score(A, B)$ given for each similarity type $sim\text{-}type(A, B)$. **Right:** Average summary scores $sum\text{-}score(C_{i,j}, S)$ given for each summary type $sum\text{-}type(C_{i,j}, S)$ with $S \in \{A, B\}$ and $i, j \in \{0, 1\}$. Error bars represent standard errors.

Table 5 Average summary scores $sum\text{-}score(C_{i,j}, S)$ for each Story $S \in \{A, B\}$ and for each summary $C_{i,j}$ with $i, j \in \{0, 1\}$, across the 4 similarity types $sim\text{-}type(A, B)$. Number of responses for each summary $C_{i,j}$ and each story S are given in parentheses.

Similarity type $sim\text{-}type(A, B)$	Average summary score $sum\text{-}score(C_{i,j}, S)$ for Story $S = A$, $S = B$ (number of responses)			
	C_{11}	C_{10}	C_{01}	C_{00}
LS	4.39, 4.27 (142)	4.29, 4.15 (34)	3.97, 4.24 (38)	2.53, 2.51 (70)
ANA	3.43, 3.32 (135)	3.43, 2.63 (35)	2.93, 3.61 (28)	2.21, 2.31 (72)
SS	1.99, 1.93 (136)	4.27, 1.45 (33)	1.22, 4.06 (36)	1.25, 1.06 (67)
ANO	2.04, 2.01 (138)	2.69, 1.34 (35)	1.34, 2.66 (35)	0.96, 0.91 (68)
Overall	2.96, 2.89 (551)	3.68, 2.42 (137)	2.43, 3.67 (137)	1.74, 1.70 (277)

Figure 2 Average summary scores $sum\text{-}score(C_{11}, A)$ and $sum\text{-}score(C_{11}, B)$ for the common summary C_{11} for each of the 4 similarity types $sim\text{-}type(A, B)$. Error bars represent standard errors.

■ **Table 6** Average similarity scores $sim\text{-}score(A, B)$ distributed across the possible summary score pairs $sum\text{-}score(C_{i,j}, A)$ and $sum\text{-}score(C_{i,j}, B)$, for all $C_{i,j}$ with $i, j \in \{0, 1\}$. Similarity scores above 2.50 are shown in bold and similarity scores above 3.00 are underlined.

$sum\text{-}score(C_{i,j}, B)$						
5	1.33	2.09	**2.71**	**<u>3.00</u>**	**<u>4.22</u>**	**<u>4.72</u>**
4	2.06	**2.55**	**<u>3.38</u>**	**<u>3.39</u>**	**<u>3.74</u>**	**<u>4.22</u>**
3	1.43	1.60	**2.75**	**2.93**	**<u>3.76</u>**	**<u>3.38</u>**
2	1.88	1.91	**2.98**	**2.81**	**2.78**	**<u>3.89</u>**
1	1.31	**2.52**	1.85	1.71	2.06	1.44
0	1.87	1.36	1.00	1.29	2.00	1.88
	0	*1*	*2*	*3*	*4*	*5*

$$sum\text{-}score(C_{i,j}, A)$$

Table 6 presents the average similarity scores $sim\text{-}score(A, B)$ given by participants, for each possible combination of scores $sum\text{-}score(C_{i,j}, A)$ and $sum\text{-}score(C_{i,j}, B)$ given for a particular summary $C_{i,j}$. Interpreting the table as a graph, each cell of the table corresponds to a point (x, y) in the graph, with coordinates the two summary scores $sum\text{-}score(C_{i,j}, A)$ and $sum\text{-}score(C_{i,j}, B)$ given for a summary $C_{i,j}$. For example, to determine the value placed in cell $(0,0)$ (i.e., in the lower left corner of the graph), we gathered all the triples $(A,B,C_{i,j})$ for all the summaries $C_{i,j}$ scored with 0 for both stories A and B and we computed the average similarity score $sim\text{-}score(A, B)$ between stories A and B for these triples. Similarly, the average similarity score $sim\text{-}score(A, B)$ for all the summaries $C_{i,j}$ scored with 1 for their appropriateness for story A and with 0 for story B is presented in the cell $(1, 0)$ of the table, and so on. Observe that for triples where summaries $C_{i,j}$ were rated with high summary scores for both stories A and B (upper right corner of the graph), the average $sim\text{-}score(A, B)$ between stories A and B was higher. Hence, the existence of an appropriate common summary for two stories suggests that the two stories are highly similar.

The data of Table 6 is also represented by the contour of Figure 3 (Left), where higher average similarity scores are represented with warmer colors than lower average similarity scores. Accordingly, the 3D curve of Figure 3 (Right) represents the two summary scores $(sum\text{-}score(C_{i,j}, A), sum\text{-}score(C_{i,j}, B))$ by a point (x, y) in the horizontal plane, and the average similarity score for each point (x, y) by the height z of the curve in the vertical axis. Taken together, Table 6 and the two graphs of Figure 3, confirm our main hypothesis: as the summary scores $sum\text{-}score(C_{i,j}, A)$ and $sum\text{-}score(C_{i,j}, B)$ increase, indicating that the summary $C_{i,j}$ becomes more and more appropriate for stories A and B, the average similarity score $sim\text{-}score(A, B)$ also increases.

Finally, we attempted to more systematically measure the appropriateness of a summary $C_{i,j}$ for *both* stories A and B simultaneously, with a single score. Using the two summary scores $sum\text{-}score(C_{i,j}, A)$ and $sum\text{-}score(C_{i,j}, B)$, we estimated the *appropriateness of summary $C_{i,j}$ as a common summary of A and B* with 4 different computational methods: i) the *sum* of scores, ii) the *product* of scores, iii) the *Euclidean (L_2) distance* between scores and the point $(5,5)$, subtracted from the maximum distance between two points (which is $5\sqrt{2}$), and iv) the *minimum* of the two scores. The above functions were selected so that the higher the values they produce, the more appropriate $C_{i,j}$ would be as a *common summary* of stories A and B. When graphing the relation between the average similarity score $sim\text{-}score(A, B)$ and the degree of $C_{i,j}$ as a common summary, as computed using each of the above methods, the relation was found to be directly proportional, and the two quantities extremely highly correlated, confirming our main hypothesis once again (Figure 4).

Figure 3 Average similarity scores $sim\text{-}score(A, B)$ distributed across the possible summary score pairs $sum\text{-}score(C_{i,j}, A)$ and $sum\text{-}score(C_{i,j}, B)$, for all $C_{i,j}$ with $i, j \in \{0, 1\}$. **Left:** Contour plot, where average similarity scores are represented by the colors of the spectrum; warmer colors (red) indicate higher similarity scores, while colder colors (violet) indicate lower similarity scores. **Right:** 3D plot, where average similarity scores are represented by the plot height (vertical axis).

5 Discussion

Taken together, the results of our study strongly confirmed our initial hypothesis: stories that are both summarized by a single summary are similar to each other, and the more appropriate this common summary is, the stronger the similarity of the stories becomes. The common summary C_{11} was meant to be an appropriate summary for both stories A and B, and we attempted to choose it to be so during the design of our experiment. However, summary scores given to C_{11}, regarding stories A, B of the SS and ANO similarity types, were relatively low, indicating that this summary was not appropriate enough for the SS and ANO pairs. Accordingly, it seems that we were unable to produce, even though we tried to, an appropriate enough C_{11} summary for the story pairs of low similarity. This is directly in line with our main hypothesis: dissimilar story pairs cannot have appropriate common summaries.

Moreover, the participants' ratings for similarity between the story pairs (A,B) were in accordance with relevant previous work on similarity judgments for stories [9]. Gentner, *et al.* designed several story pairs according to the similarity types of Table 1 and asked participants to rate, among others, the subjective similarity of these pairs on a scale rating from 1 to 5. According to their results, story pairs of the LS type were considered as more similar than those of the ANA type, and story pairs of the ANA type as more similar than those of the SS type. In their study they did not use any story pairs of the ANO type. However, given the definition of the Anomaly type as stories that share none of their predicates, we could plausibly assume that story pairs of the ANO type would be rated with the lowest degree of similarity among all types. Moreover, in Gentner *et al.* [9], the average similarity ratings for the ANA type were close enough to the ratings for the LS type, while ratings for the SS type were much lower. Our results on similarity judgments perfectly reproduced the pattern found in previous work, indicating that the ordering of the similarity types reflects the degree of relational overlap between the pairs.

However, in a previous attempt to empirically determine the factors that affect human

Figure 4 Average similarity score $sim\text{-}score(A,B)$ (y-axis) over the degree of appropriateness of $C_{i,j}$ as a common summary of A and B (x-axis), for each $i,j \in \{0,1\}$. In each case the x-axis value was calculated as a function F of the two summary scores $sum\text{-}score(C_{i,j},A)$ and $sum\text{-}score(C_{i,j},B)$. **I:** $F_I(C_{i,j},A,B) = sum\text{-}score(C_{i,j},A) + sum\text{-}score(C_{i,j},B)$ **II:** $F_{II}(C_{i,j},A,B) = sum\text{-}score(C_{i,j},A) \times sum\text{-}score(C_{i,j},B)$ **III:** $F_{III}(C_{i,j},A,B) = 5\sqrt{2} - d[(sum\text{-}score(C_{i,j},A), sum\text{-}score(C_{i,j},B)),(5,5)]$ **IV:** $F_{IV}(C_{i,j},A,B) = min[sum\text{-}score(C_{i,j},A), sum\text{-}score(C_{i,j},B)]$

judgments on story comparison, Fisseni and Löwe [6] concluded that structural factors are not the most important aspect for subjective similarity. This discrepancy with our and Gentner *et al.*'s results may be due to the different methodology used. Fisseni and Löwe asked participants to judge pairs of stories as same or different. In a second experiment, participants were additionally asked to justify their sameness ratings by stating as many differences between the stories as possible. This instruction, to explicitly state differences, might have affected their responses. People tend to find more differences for highly similar than for less similar pairs [8]. Since the story pairs used in that study were highly similar (variants of the same story), participants may have reported many differences, even though the stories were similar enough. Accordingly, this may have biased their sameness judgments.

Our empirically tested hypothesis comes from a previous work of Michael [19], as part of a logic-based theoretical framework attempting to computationally define aspects of story understanding. The formal definitions given therein offer qualitative metrics of the appropriateness of a summary and the degree of similarity between two stories. A credulous

common summary of two stories is defined as a summary that includes at least the most important parts of the stories; and the existence of a credulous common summary is taken to imply that the two stories are credulously similar in that they share at least their most important parts. The present empirical study, other than to empirically support the psychological validity of the hypothesis offered in that work, can also be seen to extend the logic-based framework by defining some quantitative metrics for the 'credulousness' of a common summary.

The standard distance metrics $L_1 = 10 - F_I$, $L_2 = 5\sqrt{2} - F_{III}$ and $L_\infty = 5 - F_{IV}$, based on the functions F_I, F_{III}, and F_{IV} as previously defined (Figure 4), represent the distance between the points $(sum\text{-}score(C_{i,j}, A),\ sum\text{-}score(C_{i,j}, B))$ and $(5,5)$. Considering that the point $(5,5)$ represents the common summary that is the most appropriate for both stories A and B (given that 5 is the maximum score for the score $sum\text{-}score(C_{i,j}, S)$), it follows that the lower the value of these metrics, the closer the points are to the 'most appropriate common' summary. Borrowing the notion of 'credulousness' of a summary [19], we could name these quantities as "metrics of the credulousness of a common summary".

Finally, we may say that our results reflect the two mapping principles of the SMT. The fact that participants judged story pairs of the ANA type as more similar than those of the SS type, indicates that people consider relational matches as more important than object matches and also, in accordance to the systematicity principle of the SMT, that higher-order relations count more than first-order relations for people's judgments.

6 Conclusion

Identifying similarities among stories is a central part of the process of making sense of stories, and building machines for the latter task will presumably require some solution to the former. In this work we have provided overwhelming psychological evidence that the more appropriate a given story is as a common summary of two other stories, the more similar the latter two stories are to each other. The validity of this hypothesis offers a sufficient condition to test for similarity, or more precisely, offers a way to lower bound the degree of similarity. The condition is not, however, necessary, since the failure of a candidate summary to be an appropriate common summary of two stories does not indicate lack of similarity between the two stories, since some other candidate summary could exist that would be appropriate. Devising a method to produce candidate summaries that would be the most specific common summaries of two stories would offer the missing link to establish the necessity of the condition as well. The role of expectations in stories [18, 19] would seem to be important to that end.

The present study was a first step towards the confirmation of our hypothesis for a certain sample of the possible types of stories. Further research could examine the applicability of this hypothesis to other genres of stories. In a different direction, we could analyze stories extracted automatically from online sources, in order to avoid manually selecting specific types of stories, and any bias this choice may bring to the empirical study.

It would be interesting to generalize our hypothesis beyond stories, and to examine whether similarity between two concepts is effectively equivalent to saying that the two concepts share a common abstraction which is appropriate for both of them. Such concepts could be short videos, simple images, or sound clips. We believe that the empirical methodology developed herein, and the type of analysis performed, could be applied equally well to such more general settings.

References

1 Jerome Bruner. The narrative construction of reality. *Critical Inquiry*, 18(1):1–21, 1991.

2 Nathanael Chambers. *Inducing Event Schemas and their Participants from Unlabeled Text.* PhD thesis, Stanford University, 2011.

3 David K. Elson. Detecting story analogies from annotations of time, action and agency. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 91–99, İstanbul, 2012.

4 David K. Elson. *Modelling Narrative Discourse.* PhD thesis, Columbia University, 2012.

5 Matthew P. Fay. Story comparison via simultaneous matching and alignment. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 100–104, İstanbul, 2012.

6 Bernhard Fisseni and Benedikt Löwe. Which dimensions of narratives are relevant for human judgments of story equivalence? In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 114–118, İstanbul, 2012.

7 Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, 1983.

8 Dedre Gentner and Arthur B. Markman. Structure mapping in analogy and similarity. *Americal Psychologist*, 52(1):45–56, 1997.

9 Dedre Gentner, Mary Jo Rattermann, and Kenneth D. Forbus. The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, 25:524–575, 1993.

10 Dedre Gentner and Linsey Smith. Analogical reasoning. In V. S. Ramachandran, editor, *Encyclopedia of Human Behavior*, pages 130–136, Oxford, UK, 2012. Elsevier.

11 Ulrike Hahn, Nick Chater, and Lucy B. Richardson. Similarity as transformation. *Cognition*, 87:1–32, 2003.

12 Ulrike Hahn, James Close, and Markus Graf. Transformation direction influences shape-similarity judgments. *Psychological Science*, 20(4):447–454, 2009.

13 Caryn E. Krakauer and Patrick H. Winston. Story retrieval and comparison using concept patterns. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 119–124, İstanbul, 2012.

14 Levi B. Larkey and Arthur B. Markman. Processes of similarity judgment. *Cognitive Science*, 29:1061–1076, 2005.

15 Wendy G. Lehnert. Plot units: A narrative summarization strategy. In Wendy G. Lehnert and Martin H. Ringle, editors, *Strategies for Natural Language Processing*, pages 375–414, Hillsdale, NJ, 1981. Erlbaum.

16 Wendy G. Lehnert. Plot units and narrative summarization. *Cognitive Science*, 5(4):293–331, 1981.

17 Benedikt Löwe. Comparing formal frameworks of narrative structures. In Mark Alan Finlayson, editor, *Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium*, number FS-10-04 in AAAI Technical Reports, pages 45–46. AAAI Press, 2010.

18 Loizos Michael. Similarity of narratives. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 105–113, İstanbul, 2012.

19 Loizos Michael. Story understanding... calculemus! In *Proceedings of the 11th International Symposium on Logical Formalizations of Commonsense Reasoning 2013, Ayia Napa, Cyprus*, 2013.

20 Robert M. Nosofsky. Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1):104–114, 1984.

21 Vladimir Yakovlevich Propp. *Morphology of the Folktale.* University of Texas Press, Austin, TX, 2nd edition, 1968. transl. L. Scott.

22   David E. Rumelhart. *Introduction to Human Information Processing.* John Wiley and Sons, New York, 1977.

23   Roger C. Schank and Robert P. Abelson. *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures.* Erlbaum, Oxford, England, 1977.

24   Roger C. Schank and Robert P. Abelson. Knowledge and memory: The real story. In Jr Robert S. Wyer, editor, *Knowledge and Memory: The Real Story*, pages 1–85, Hillsdale, NJ, 1995. Erlbaum.

25   Roger N. Shepard. Stimulus and Response Generalization: A Stochastic Model Relating Generalization to Distance in Psychological Space. *Psychometrika*, 22:325–345, 1957.

26   Roger N. Shepard. Representation of structure in similarity data: Problems and prospects. *Psychometrika*, 39:373–421, 1974.

27   Amos Tversky. Features of similarity. *Psychological Review*, 84:327–352, 1977.

Testing Reader Ethical Judgments over the Course of a Narrative*

Greg Lessard[1] and Michael Levison[2]

1   French Studies
     Queen's University
     Kingston, Ontario, Canada
     greg.lessard@queensu.ca
2   School of Computing
     Queen's University
     Kingston, Ontario, Canada
     levison@cs.queensu.ca

─── **Abstract** ───

We present a web-based environment – an *Ethics Workbench* – which allows a reader's ethical judgments to be solicited while reading a narrative. Preliminary results show generally consistent scores across subjects and test conditions, and suggest that it is possible to measure how individual readers respond to texts in terms of ethical judgments, how the linearity inherent in narrative plays a role in affecting ethical judgments, and how readers appear to synthesize judgments over the course of a text. Applications of the model include the empirical analysis of the ethical aspects of reading, the more detailed study of ethical issues, the potential for eliciting ethical discussions, and a means of dynamically planning texts to achieve maximum effect with respect to reader judgments.

1   Background

It would be difficult to find two broader fields than narrative and ethics. Each has given rise to an enormous range and depth of research, in literary studies, philosophy, ethnology, and so on. Their intersection has led to work on the evolutionary advantages of narrative as a model of ethical and other choices [3], a means of 'improving' readers on the ethical dimension [2], and as one of the principles which structures narrative [7]. Ethical phenomena in narrative have been studied in broader fields like film [8] and women's studies [4], applied to the analysis of personality traits such as empathy [6], and used in training of medical practitioners [5]. As a complement to these various approaches, it would be of value to have a means of determining readers' ethical judgments over the course of a narrative, to determine how these are influenced by narrative structures and content. We present here an environment designed to achieve these goals and we discuss some of our preliminary findings.

2   Testing reader reactions

To test reader reactions incrementally, we constructed a web page which presents a story one element at a time. At various points in the story, the reader is asked to respond to a certain number of ethical questions before moving on to the next story element. Reader judgments are expressed by moving a slider, where movement to the left indicates a progressively greater negative judgment, while movement to the right indicates a progressively more positive ethical judgment.[1] Each slider represents a linear scale, with 0 at the midpoint and -10 and +10 as endpoints. The slider is set initially at the midpoint of the scale. Once a reader judgment has been entered, the slider disappears, thereby preventing explicit comparison with previous judgments. The test framework is based on the Aesop's Fable *The Ant and the Grasshopper*. Although the protagonists are insects, the fable was designed to present a moral principle, and we have found that the strong human tendency to anthropomorphize animals leads to treating the two characters as if they were humans. Two test scenarios were developed, differing at only one point, as shown below in item (c), where the reason for the grasshopper's lack of work is presented either as age and illness, or as a desire to sing and mock the ant. The framework for both scenarios is shown below, where sections beginning with letters indicate items of the story as told, and bolded and italicized sections beginning with digits indicate requests for judgment on the part of the reader.

(a) Once upon a time there was an ant and a grasshopper. They both lived in a large field. All summer long, the ant worked hard to collect food for the winter.
(b) The grasshopper did no work and collected no food.
(1) **Evaluate the grasshopper's behaviour.**
(c) Either: He was old and ill. Or: He sang in the sun and mocked the ant for its laborious behaviour.
(2) **In light of this new information, evaluate the grasshopper's behaviour.**
(d) At the end of the summer, the ant had a lot of food stored away, but the grasshopper had none. When the first frost came, the grasshopper approached the ant and asked for food.
(e) The ant told the grasshopper that since he hadn't worked, he deserved no food.
(3) **How much do you agree with this decision?**
(f) The grasshopper slowly starved to death.
(4) **How appropriate is this outcome?**
(5) **In light of the whole story, evaluate the ant's behaviour.**

3   Some preliminary results

The scenarios were presented to 10 adult native speakers of English, 5 male and 5 female. All subjects were given as much time as required to read and to enter their judgments. Half the subjects were exposed to the *old and ill* scenario, and half to the *mocking* scenario. Assignment to a particular scenario was done at random. The object of the experiment was to consider four issues in particular: (a) Narrative is linear, and subsequent items in a test are invisible to a reader until they are reached. Given this, to what extent does exposure to new information lead to changes in ethical judgments? Questions 1 and 2 test this, by asking first for a judgment of the grasshopper's failure to work, without any background information, and

---

[1] See for example, Bard et al. [1] for discussion of this approach.

then asking again in light of further positive or negative information (ill-health or mocking). (b) Ethical decisions are presumably related to presented states of affairs. Question 3 tests whether, in the mind of the reader, the ant's decision not to share food is correlated with previous information. (c) Ethical decisions have consequences. In the version of the Aesop's Fable presented here, the ant's negative decision leads to the grasshopper's death. Question 4 measures the reader's reaction to this outcome. (d) By taking a decision which leads to the grasshopper's death, the ant becomes an ethical agent whose actions may be judged. Question 5 tests the reader's evaluation of the ant.

Figures 1 and 2 show the results of the tests.[2] The left hand axis represents the slider value at each point, the various lines labeled Subject1 to Subject10 represent the trajectory of judgments for each test subject, and the digits from 1 to 5 across the middle of the graph represent each of the five ethical evaluations expressed using sliders.

■ **Figure 1** Subject testing results for *old and ill* scenario.

3.1 Narrative order and inflection of ethical judgments

Examination of Figure 1 shows that in the *old and ill* scenario, the subsequent explanation of the grasshopper's inactivity causes all judgments to be inflected in a positive direction (see results for questions 1 and 2). On the other hand, as Figure 2 shows, 4 of 5 subjects inflect their evaluation of the grasshopper in a negative direction when learning of his singing in the sun and mocking the ant.[3] Taken together, this data provides at least *prima facie* evidence that reader judgments follow the ordered presentation of information in a narrative.

---

[2] Since this is preliminary work and the number of subjects is small, no statistical tests were done on the data presented here. The focus is rather on broad tendencies found in the data, as well as an indication of whether the environment used can work in practice.

[3] In the case of Subject 7 in the *mocking* scenario, the test subject noted that the evaluation chosen was based on the grasshopper's singing, which was seen as a good thing. This illustrates the importance of carefully linking ethical choices to well-defined aspects of the narrative.

Figure 2 Subject testing results for *mocking* scenario.

3.2   Ethical situations and ethical decisions

Both scenarios provide background and then ask the reader to pronounce on his or her agreement with the ant's decision. Analysis of Figures 1 and 2 shows that in the case of the old and ill grasshopper, all five subjects disagreed more or less strongly with the ant's decision, as shown by the negative values for Question 3. However, in the case of the mocking scenario, four of the five test subjects agreed with the ant's decision. This would suggest that previous information is being used to evaluate not just the players in this drama, but also the decisions being made.[4]

3.3   Ethical decisions and their consequences

The test environment spells out explicitly the fatal consequences for the grasshopper of the ant's ethical judgment. In the case of Figure 2, we can see a sharp divergence between three subjects who saw the mocking grasshopper's death as quite inappropriate (with scores of -6 to -10) and two subjects who saw the death as at least somewhat appropriate (with scores of +2 and +3). We find a similar divergence in Figure 1, where four subjects see the old and ill grasshopper's death as inappropriate, while one (Subject 1), sees the death as at least marginally appropriate, with a score of +1.[5] In general, though, the data presented appears to show that the gravity of the consequences of some act has a measurable effect on reader judgments.

---

[4] Note that once again, Subject 7 runs counter to the other four test subjects, perhaps in light of this subject's previous positive evaluation of the grasshopper. This is a case where more 'think-aloud' data would be of value.

[5] Note that this subject has also given the grasshopper the lowest initial score and that the curve of this subject's judgments never diminishes over the course of the scenario. The notion of the 'arc of ethical judgment' merits more attention.

3.4 From ethical paragon to ethical agent

The ant is at least partially responsible for the grasshopper's death from starvation. In light of this, it is possible to ask how the ant itself is judged. Consideration of the two Figures shows that in all cases but one (Subject 9, whose evaluation of the grasshopper's demise was positive but of the ant negative) subjects who considered the grasshopper's demise to be inappropriate judged the ant's behaviour negatively, while those who saw the grasshopper's death to be appropriate judged the ant's behaviour to be positive. Thus, in the *old and ill* scenario, Subject 1 is consistent in answering Questions 4 and 5, as is Subject 10 in the *mocking* scenario, by evaluating both the demise and the ant's decision as positive. This consistency would suggest that readers are capable of focusing on particular players in a scenario as they become involved with ethical decisions.

4 Applications of an Ethics Workbench

Clearly, the framework here is still very simple and the number of subjects tested small. Further research will apply the framework to other, human-based contexts, more test subjects, and richer narrative models. It will also be important to solicit initial ethical positions before texts are read. That being said, we believe that the framework presented here has potential applications in several areas, as: (i) a means of studying, in more detail than we have done here, the interplay of narrative and reader judgments in areas like the duration of ethical judgments over a narrative, and the role of repeated behaviours; (ii) a locus of discussion for dealing with ethical issues, for example in Philosophy or Sociology classes,[6] by instantiating preset principles in stories and examining outcomes, or by studying reader reactions to developing scenarios as a stimulus to ethical self-examination; (iii) a means of 'crowdsourcing' reader judgments, by the capture of evaluations across many readers, as in [9], or by 'shopping' political or even marketing narratives among target groups; (iv) to the extent that reader judgments are allowed to inflect evolving narratives, as one element of a dynamic narrative structure, a sort of 'choose your own narrative' with an ethical dimension.

──── **References** ────

1 Ellen Gurman Bard, Dan Robertson, and Antonella Sorace. Magnitude estimation of linguistic acceptability. *Language*, 72(1):32–68, 1996.

2 Wayne C. Booth. *The company we keep: an ethics of fiction*. University of California Press, Berkeley, 1988.

3 Brian Boyd. *On the origin of stories: evolution, cognition, and fiction*. Belknap Press of Harvard University Press, Cambridge, MA, 2009.

4 Carol Gilligan. *In a different voice: psychological theory and women's development*. Harvard University Press, Cambridge, MA, 1982.

5 Brian Hurwitz. Textual practices in crafting bioethics cases. *Bioethical Inquiry*, 9:395–401, 2012.

6 Susanne Keen. Narrative empathy. In Frederick Luis Aldama, editor, *Toward a cognitive theory of narrative acts*, pages 61–93, Austin, 2010. University of Texas Press.

7 James Phelan. Narrative judgments and the rhetorical theory of narrative: Ian McEwan's *Atonement*. In James Phelan and Peter J. Rabinowitz, editors, *A companion to narrative theory*, pages 322–336, Oxford, 2005. Blackwell.

───────

[6] We are indebted to Paul O'Marra for raising this point.

8   Jane Stadler. *Pulling focus: intersubjective experience, narrative film, and ethics*. Continuum, New York, 2008.

9   Reid Swanson and Arnav Jhala. A crowd-sourced collection of narratives for studying conflict. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 65–73, İstanbul, 2012.

Theoretical Issues in the Computational Modelling of Yorùbá Narratives

Olufemi D. Ninan and Odetunji A. Ọdẹ́jọbí

Department of Computer Science and Engineering
Ọbáfẹ́mi Awólọ́wọ̀ University
Ilé-Ifẹ̀, Nigeria
`{jninano,odejobi}@oauife.edu.ng`

── Abstract ───────────

Developing a coherent computational model for narratives across multiple cultures raises the question of the components and structure of a framework within which African narratives can be conceptualised and formalised. It is well known that narratives are influenced by cultural, linguistic, and cognitive factors. We identify and define entities, elements, and relations necessary for the adequate description of Yorùbá narratives. We also discuss these theoretical issues in the context of designing a formal framework that is amenable to computational modelling.

1 Introduction

This position paper addresses some of the important theoretical issues relating to the computational modelling of *Yorùbá* narratives. A major challenge we encountered in our ongoing research project [11] on the computational modelling of African folktales is the accurate formal description of the events composing the narratives. The main challenge relates to how to account for some concepts, spatial and temporal relations, and processes that are particular to African and *Yorùbá* culture. We think that the computational representations of narratives should correspond to the mental models formed by the people when communicating and comprehending events in the story. These mental models underlie the conception, process, and purpose of narratives embedded in the cognitive make-up of the people. It is well known that each African tribe has a unique culture and language and hence unique world-view. We think that the cosmology of a culture influences its world-view. Therefore narratives cannot be considered in isolation from the cultural environment from which they emanate. We need to address the question of how African narratives can be conceptualised and formalised in a computational model. Addressing this question subsumes other fundamental questions such as: is there a universal formal modelling framework for narratives? If the framework elements and relations are different, what features characterise this difference? Are these characteristics influenced by the cultural, linguistics, and cognitive factors? If the elements and symbolism of narrative are universal, casting narratives in a computational model should be constant across time and culture. We think that this is not the case as we consider narrative to be a creative work of art communicated by way of language and heavily influenced by culture. The thesis of this paper, therefore, is that an accurate computational model for narrative must necessarily account for formalisable aspects of the cosmology of the culture that produced

the narrative. Our motivation for this work is suggested in [4, 12, 13] and summed up in the following quotation from the symposium paper by Lakoff and Narayanan [10]:

> Narratives exploit the shared cognitive structure of human motivations, goals, emotions, actions, events, and outcomes. Computational models of narrative must therefore be capable of modelling these shared human understandings [10]

To argue our thesis, we explore the structure of narratives from the perspective of the *Yorùbá* cosmology.

2 Fundamentals of the Yorùbá cosmology of narratives

There are several fundamental concepts that are distinctive to the Yorùbá world view. They underpin the framework for understanding the dynamics of the Yorùbá narratives through time and space as well as some of its enduring philosophies. These concepts are expressed in words, images, signs, and actions used in narratives. We demonstrate our thesis through a critical analyses of Yorùbá folktales narratives.

2.1 Concept of being

The Yorùbá cosmology recognises three states of being: (i) the ancestors, (ii) the living and, (iii) the yet to be born (see Figure 1). A being can exist in two forms: the tangible (physical) and the intangible (spiritual) [1, 7]. This *three states dual existence* (TSDE) [6], concept is pervasive in Yorùbá narratives. The spiritual form of a living person or being, is considered to be the most powerful as it controls, coordinates, and determines everything that the physical form manifests. For example, the Yorùbá concept of *orí* (head) comprises *orí inún* (inside/intangible head) and *orí òde* (outside/tangible head). The *orí inún* is sacred and an object of worship while the *orí òde* is a tool for interacting with the physical world. This concept applies to other parts of the human body, particularly the eyes, legs, and hands.

Figure 1 Two world concept in Yorùbá narratives.

2.2 Origin

The origin of Yorùbá narratives are *communal*. Some narratives are associated with occupations: e.g., *Ìjálá* [14] are the chants of hunters; *Ìyèrè* are the chants of Ifá priest; *Oríkì* is a cultural praise poetry specific to an individual, a particular family, or linage. Even

in these cases, individual authorship is never claimed though an individual may introduce ideas, contents, and forms that account for the peculiar circumstance and situation into the narratives. This authorship model imposes a requirement for flexibility on the framework for the formalisation of concepts and structure in the design of a computational model for Yorùbá narratives.

2.3 Terminologies

Some amount of difficulty arises when a non-native reads a Yorùbá narrative text. The source of this difficulty is the use of culturally bounded concepts and terminologies which are evident in the use of language. A place to start this discussion therefore is to explain, as much as practically possible although perhaps not exactly, some of the concepts and terms necessary to comprehend the fundamentals of Yorùbá narratives. The meaning and use of words for days, months, seasons, numbers, proper names, possession, and relations, to name but a few are deeply rooted in the Yorùbá's cosmology [1, 2]. The meaning of the majority of these terms are often misconstrued with their English equivalents, particularly in translations. Some popular examples drawn from the home domain are listed in Table 1. The Yorùbá *ẹbí*, for example, does not have an exact equivalent English word. A word that is often used as equivalent is *family*. The Yorùbá *ẹbí*, unlike the English *family*, refers to the extended blood relatives and its members, most often, include the ancestors and those yet to be born. This use of the concept of *ẹbí* in a narrative will make sense only when interpreted within the context of the TSDE philosophy. Also, the meaning of most terms for expressing relations in Yorùbá narratives are only logical and permissible within the Yorùbá cosmology. The relation Ìyàwó (*wife*), for example, is used to refer to a woman married into an *ẹbí* (family), in contrast to the English concept of a woman legally married to a man. That is why it is not uncommon, and semantically correct, to read in a narrative that a woman refers to another woman, as *'my wife'*. A formal computational model of Yorùbá narratives must make it possible to express and adequately represent these concepts.

Table 1 Comparison of Concepts in Yorùbá and English.

Ser. No.	Term	English meaning	Yorùbá meaning
1.	Family	Nuclear and extended	*Ẹbí* applies to the extended family. Members include the ancestors and the un-born.
2.	Marriage	Union between two individuals male/female	Union between two extended families (*ẹbí*)
3.	Wife	A woman married to a man	A woman married into a family (*ẹbí*)
4.	Husband	A man married to a woman	A man assigned by the family (*ẹbí*) to take care of a woman married into the *ẹbí*
5.	Body parts	Only the physical or tangible	The physical (tangible) and non-physical (intangible). The concept of duality of being applies
6.	Child	A biological offspring	Any offspring in the family *(ẹbi)*
7.	Mother / father	Female/Male parent	Female/Male ancestors living or dead

2.3.1   Computational model

Within the context of the world view discussed above, we define *computational model of narrative* as the *creation* and *manipulation* of symbols with the aim to describe and communicate past or evolving events. The symbols in this case are labels that represent the entities in a narrative. These include those that describe the character and the props as well as all other elements or entities in the narrative. The manipulation of symbols is a process that specifies states and defines the events and state transitions in a narrative. In this case the plot of a narrative corresponds to the *algorithm* or *heuristics* of a computation process. This can be formally described within the context of reasoning about action based on ontology of the narrative domain [9].

2.3.2   Knowledge

The most appropriate definition of knowledge in the context of the world view expressed above is *"that which is acquired when information is applied"*. When a piece of information is subjected to practice, the experience resulting from that practice becomes knowledge. To know, therefore, is to perceive and personify a piece of information. This definition seems to differ from some Western definition of knowledge. of knowledge. In the Popperian cosmology [15], for example, knowledge is regarded as an object independent of knowing. Underlying that cosmology is a three world model. Popper called world-1 the physical world, world-2, the world of mind, and world-3, human knowledge expressed in its manifold forms. He opined that the growth of human knowledge could be said to be a function of the independent evolution of the three worlds. In that context, knowledge is considered an object that is independent of knowing. Structurally, this model contrasts with the four world model proposed by Ọdéjọbí [5] which is more suited to the concept of knowledge in Yorùbá narrative. What is clear here is that, the narrator has a mental model of something, real or imaginary, that he wishes to communicate. The aim of the narrator is to generate a transcript that facilitates the transfer of that mental model to the audience as accurately as possible. A challenge to this simple model of narrative knowledge is the counter-factual aspect of Yorùbá narratives, for example, situations where animals can be portrayed as exhibiting some characteristics of humans (e.g., marrying, talking, etc.), which are also, incidentally, found in many other cultures in the world. The representation of such counter-factual expressions and reasoning associated, side by side, with factual knowledge, is a challenge to current fact-based computational models of narrative.

2.3.3   Time in Narrative

In the Yorùbá cosmology, time is not defined in absolute terms but in relations to events. For example, people use phrases such as 'the last yam festival', 'the last moon' and 'when sun rises' in describing their experiences and planning their activities. Inclusive counting is used in date and time reckoning [16]. Terms such as *before*, *now*, *next day*, *later*, etc. are implicitly reckoned using inclusive counting. Interestingly, we are able to link this ordering to aspects of repeated motion in nature, such as the movement of the earth about its axis and around the sun. Logic can be used to relate the spatial and temporal dimensions into a coherent narrative, though the logic of sequence does not guarantee coherence of a narrative [8, 3].

3 Conclusion

A computational model of narrative must be sensitive to the peculiar logic and relations of the cultural domain in which it is operating. Therefore a single general-purpose (one-size-fit-all) mechanism is not sufficient. The character of the elements, relations, actions and events in a narrative are a function of purpose and world-view of the narrator. We also think that it is necessary to reach a balance between the language of computing, which is mathematical, formal and precise and the language of narratives which is linguistic, informal, and imprecise.

References

1 W. Abímbọlá. *An Exposition of Ifá Literary Corpus.* Oxford University Press, Ìbàdàn, Nigeria, 1976.

2 K. A. Appiah. *In my Father's House, Africa in the Philosophy of Culture.* New York, Oxford, 1992.

3 H-J. Backe. Narrative rules? Story logic and the structures of games. *Literary and Linguistic Computing*, 27(3):243–260, 2012.

4 Ben-Amos Dan. Analytical categories and ethnic genres. *Genre*, 2(3):275–301, 1969.

5 O. A. Ọdẹ́jọbí. *Introduction to the Principles of Artificial Intelligence.* DovePower Technologies, Ọ̀yọ́, Nigeria, 1997.

6 O. A. Ọdẹ́jọbí. Creation, existence and nature: A perspective from constraint and computing. Seminar paper, Cork Constraint Computation Center, 2012.

7 S. M. Ọpẹ́ọlá. A way of applying science education to interpret literary corpus. *Odu*, 33:149–162, 1988.

8 S.-Y. Iwasaki. A Cognitive Grammar account of time motion 'metaphors': A view from Japanese. *Cognitive Linguistics*, 20(2):341–366, 2009.

9 A. Kakas and R. Miller. A simple declarative language for describing narratives with actions. *Journal of Logic Programming*, 31(1):157–200, 1997.

10 G. Lakoff and S. Narayanan. Towards a computational model of narratives. In Mark Alan Finlayson, editor, *Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium*, number FS-10-04 in AAAI Technical Reports, pages 21–28. Association for the Advancement of Artificial Intelligence, 2010.

11 D. O. Ninan and O. A. Ọdẹ́jọbí. Towards a digital resource for African folktales. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 75–81, Ìstanbul, 2012.

12 'T. Ògúnpọlá. Classification of Ìjẹ̀bú Yorùbá prose narratives. In *2nd Annual Congress, Nigerian Folklore Society*, pages 1–28, The University of Ìlọrin, Nigeria, 1986.

13 'T. Ògúnpọlá. Traditional prose narratives as culture indicator: The Yorùbá example. In *17th Annual Congress of The West African Linguistic Society*, pages 1–15, The University of Ìbàdàn, Nigeria, 1986.

14 S. Olorode. *Ìjálá.* PhD thesis, Obafemi Awolowo University, Ile-Ife, Nigeria, 1986.

15 Karl Popper. Three worlds. Tanner lectures on Human Values, delivered at the University of Michigan, 1977/78.

16 C. Zaslavskey. *Africa Counts: Number and Pattern in African Cultures.* Lawrence Hill Books, Chicago, third edition, 1973.

Constructing Spatial Representations from Narratives and Non-Narrative Descriptions: Evidence from 7-year-olds*

Angela Nyhout and Daniela K. O'Neill

Center for Child Studies
University of Waterloo
Waterloo, Ontario, Canada
`aknyhout@uwaterloo.ca, doneill@uwaterloo.ca`

──── **Abstract** ────

Although narratives often contain detailed descriptions of space and setting and readers frequently report vividly imagining these story worlds, evidence for the construction of spatial representations during narrative processing is currently mixed. In the present study, we investigated 7 year old children's ability to construct spatial representations of narrative spaces and compared this to the ability to construct representations from non-narrative descriptions. We hypothesized that performance would be better in the narrative condition, where children have the opportunity to construct a multi-dimensional situation model built around the character's motivations and actions. Children listened to either a narrative that included a character traveling between 5 locations in her neighbourhood or a description of the same 5-location neighbourhood. Those in the narrative condition significantly outperformed those in the description condition in constructing the layout of the neighbourhood locations. Moreover, regression analyses revealed that whereas performance on the narrative version was predicted by narrative comprehension ability, performance on the description version was predicted by working memory ability. These results suggest the possibility that building spatial representations from narratives and non-narratives may engage different cognitive processes.

1 Introduction

The subjective experience of readers and listeners of narratives is often one of being transported into the narrative world, vicariously participating in the unravelling events [7]. Individuals may "feel" hot sand beneath their feet as a story protagonist walks along a beach or "see" the destruction caused by a tornado to a character's house. To the experiencer, narrative processing, with the imagination, perspective-taking, and emotional engagement that it encourages and induces, may seem qualitatively different from non-narrative processing. Common, subjective experience suggests that narrative processing is a feat of the imagination, often including visuospatial components [6, 12, 27]. Space and setting (the *where*) are the

---

components of narrative that are most often associated with these visuospatial representations [9]. Yet, experimental evidence for representations of space in narratives is currently mixed. There is little doubt that most individuals are *able to* form spatial representations of narratives, but whether they do so spontaneously is another matter [9, 14, 29].

Representing narrative spaces may be important for several reasons, both outside and within the narrative world. Creating an accurate representation of space in narrative may be important for navigation in the real world. Narratives told by foraging peoples place great emphasis on "the lay of the land, travel routes, or orienteering-knowledge critical to undertaking extended hunting, trading, or visiting trips, which are an important part of forager life" [20, p. 243]. Additionally, accurately representing a narrative's space may be important for understanding events within the narrative. In some cases, understanding a causal sequence and making inferences may hinge on building an accurate spatial representation [9]. Although constructing a spatial representation may not be necessary for survival or comprehension in many cases, one often finds that storytellers include descriptions of spaces and settings of the story world to add detail and colour to their narratives.

1.1 The situation model

Underlying the imaginative experience that many readers and listeners of narratives report is the construction of a situation model. During narrative processing, individuals create representations not only of word order and meaning, but also representations of the situation the text or spoken language is about [11, 24]. These representations of the situations described by the text are known as *mental models* [11] or *situation models* [24, 29]. The situations described by sentences are retained in memory and used to make judgments [1, 8, 9]). Situation models are multidimensional representations, created by combining the content of the text with prior knowledge, and may include temporal, spatial, causal, person and object, and intentional information [11, 29].

There is evidence to suggest that adults construct situation models along most of the major dimensions during narrative processing [28, 30], but, as mentioned above, experimental evidence for the construction of spatial situation models is currently mixed.

1.2 Children's spatial situation models

To our knowledge, there has been no investigation of children's spatial situation model construction during narrative processing. Studies investigating other aspects of children's situation models have found that children track characters' physical [19, 26], mental [5, 17], and spatiotemporal perspectives [5]. Although they have not directly assessed children's ability to construct spatial situation models, these studies provide strong evidence that children spontaneously track characters' perspectives and movements. It may be reasonable to expect that they should also represent the space in which the characters are perceiving and moving. Indeed, Bruner has argued that "the inseparability of character, setting, and action must be deeply rooted in the nature of narrative thought. It is only with difficulty that we can conceive of each of them in isolation" [2, p. 39]. In other words, it is potentially difficult to construct a situation model that lacks one of the major dimensions; a character who acts without goals, an occurrence without apparent cause, or an event devoid of setting. Perhaps we should not conceive of the dimensions of situation models as independent of one another, but, rather, as deeply intertwined.

Outside the domain of narrative, Uttal, Fisher, and Taylor [23] compared eight- and 10-year-old children's and adults' ability to create representations of space from descriptions

to representations created from maps. Participants heard a description of a six-room building or saw a map of the same space and were then asked to assemble the space using six cards. Eight-year-old children who heard the description had difficulty with the task; their situation models seemed to be tied to the sequential order in which locations were mentioned.

Although this study suggests that children have difficulty constructing spatial representations from language when no visual information is available, it tells us little about how children may be able to create spatial situation models during *narrative* comprehension. Perhaps representations created during narrative processing are qualitatively different from representations created from descriptions, or perhaps this ability is similarly limited during narrative processing.

The process of constructing a spatial representation from a narrative may be fundamentally different than that from a description. In the former, one is following a character through space, whereas in the latter, one must conceive of the space from a characterless perspective. The provision of character, actions, motivations, and time may mean the construction of the spatial situation model is more character-driven than in the case of a non-narrative description, in which a multidimensional situation model cannot be constructed. Thus, when spatial information is presented in the form of a narrative, the system that builds spatial situation models may be engaged more readily than in the case of a non-narrative description. The construction of spatial representations from non-narrative descriptions may be more of a working memory process—something akin to memorizing a grocery list.

1.3   Outline of experiment

The experiment described here aimed to discover what children's spatial situation models of narratives look like—the amount and type of detail they include, and how they may be different from spatial representations of non-narratives. We also included measures to attempt to uncover the abilities related to children's construction of situation models, such as language comprehension, general language ability, spatial ability, and working memory. If the construction of spatial situation models is indeed a different process depending on whether spatial information is presented in the form of a narrative or a non-narrative description, one may expect different abilities to be recruited, and, thus, success on each task may be associated with strengths in different areas.

The present experiment compared seven-year-old children's abilities to construct spatial representations of narratives and non-narrative descriptions. Seven-year-olds were chosen, because it is at this age that children begin regularly encountering narratives that they must mentally construct, without the support of any visuals. Canadian children, the sample that participated in the present study, typically begin reading short novels at this age and often hear stories read aloud in class. Additionally, we felt the task may be too complex for younger children (cf., e.g, [23]).

2   Method

2.1   Participants

Participants were 38 7-year-old children ($M = 7.55$ years, range $= 7.17$ to 8.0, $SD = 2.88$ months; 20 girls). All children were recruited through a laboratory database and were in Canadian second grade.

2.2 Task procedures

Spatial situation model (SSM) task. Children were randomly assigned to one of two conditions: *narrative* or *description* and heard one of one of two corresponding pre-recorded passages about a character's neighbourhood. Children in the *narrative* condition heard about a child (Molly or Max) who bakes cookies and delivers them to four locations in the neighbourhood. The relative position of the locations comes about through the character's movement through space; for example, the character is described as walking "over the bridge to the library that's across the river from her house." Children in the *description* condition heard a description of the same four locations without the presence of a character moving between them. The relative position of the locations was explicitly stated; for example, the library is described as being "across the river from Molly's house, over the bridge". The passages in both conditions were designed to be as similar as possible, with the critical difference between the two being the presence of a goal-driven character moving through space. The narrative passage also included a three-sentence introduction that presented the character's motivations for visiting the locations in the neighbourhood.

After having listened to the passage twice, participants in both conditions were presented with a box with the following three-dimensional model pieces placed randomly within it: house, fire station, veterinarian's office, library, toy store, road, river, and bridge, and were asked to build Molly's neighbourhood.

Coding. Participation was video recorded for later analysis. Two coders, the second blind to participant condition and the purpose of the study, provided a code for each participant based upon a screen capture image provided. The coding scheme used required participants to represent meaningful relations between locations in the neighbourhood. Participants received a score ranging from 0 to 5 based upon their placement of the five locations (the character's house and the four locations she visited.)

Narrative comprehension. Two stories were chosen from the Neale Analysis of Reading Ability Test [16], a standardized tool designed to assess children's reading accuracy and comprehension. Although the tool is designed to be a reading test, children listened to the stories, because the SSM task involved listening, rather than reading. After listening to each story on headphones, children are asked a series of comprehension questions.

Listening comprehension. The Listening Comprehension subtest from the Woodcock-Johnson Tests of Achievement [25], a measure of language comprehension, was administered. Participants listen to sentences and short passages of increasing difficulty and are asked to provide a word to complete the passage. Appropriate completion depends on having processed and comprehended the passage as a whole.

Picture vocabulary. Children completed the Picture Vocabulary subtest from the Woodcock-Johnson Tests of Achievement [25], as a measure of general language ability. In this expressive vocabulary test, participants are asked to provide a label for pictures of increasing difficulty.

Sentence Span. A sentence span test, a test of verbal working memory, adapted from the widely-used reading span test [4] by Swanson, Cochrane, and Ewers [21], was administered. In this task, participants are presented with sets of unrelated sentences on a screen and are asked to remember the last word from each sentence. To ensure participants are paying attention to sentences as a whole, they are asked a factual comprehension question about one

■ **Table 1** Correlations between performance on narrative or description SSM versions and other measures.

	Narrative Comp.	Listening Comp.	Picture Vocab.	Mental Rotation	Sentence Span (WM)
Narrative SSM	.63**	.35	.18	.01	−.23
Description SSM	.27	.54*	.31	.08	−.53*

*) indicates correlation is significant at .05 level
**) indicates correlation is significant at .01 level.

of the sentences before being cued to recall the words. Participants only receive credit for recalling words on sets for which they have answered the comprehension question correctly.

Mental Rotation. A mental rotation test [13] was included as a measure of children's spatial ability. On this test, children are required to choose from four, candidate whole shapes the shape two pieces would make if put together. Items require mental translation and/or rotation of the pieces to arrive at the correct answer.

3 Results

3.1 Spatial situation model task

No gender differences were found, so results for both genders were analyzed together. Children in the narrative condition ($M = 3.44$, $SE = .25$) significantly outperformed those in the description condition ($M = 2.75$, $SE = .23$), $t(36) = 2.08$, $p = .045$.

3.2 Spatial situation model task performance and its predictors

Because it was hypothesized that different processes may underlie the construction of spatial representations depending on whether the information is presented in the form of a narrative or a description, the data were divided up by condition. Performance on the narrative version of the SSM task was significantly correlated with narrative comprehension scores ($r = .64$, $p = .006$). Performance on the description version of the SSM task was significantly correlated with listening comprehension scores ($r = .54$, $p = .014$) and sentence span (working memory) scores ($r = .53$, $p = .017$). See Table 1.

Multiple regression analysis was used to investigate whether certain abilities predicted performance on the narrative and description versions of the SSM task. Children's performance in the narrative condition was best predicted by narrative comprehension ($\beta = .72$, $p = .001$) and sentence span (working memory) ($\beta = -.49$, $p = .013$), explaining 51.2% of variance in narrative SSM task scores, Adjusted $R^2 = .516$, $F(2, 17) = 10.07$, $p = .002$. Children's performance in the description condition was best predicted by sentence span alone ($\beta = .53$, $p = .017$), explaining 23.8% of variance in description SSM task scores, Adjusted $R^2 = .238$, $F(1, 19) = 6.94$, $p = .017$.

4 Discussion

Children in the present study created more accurate external models of the neighbourhood in the narrative condition than in the description condition. Additionally, preliminary regression

analyses suggest that different cognitive processes *may* be recruited to perform the narrative and description versions of the SSM task. Performance in the narrative condition was predicted by narrative comprehension scores and was negatively predicted by sentence span, whereas performance in the description condition was predicted by sentence span alone.

Note that these results were not necessarily predictable from the outset. Research in adult education has yielded mixed results when comparing undergraduates' performance on tests of material encountered in either a narrative or non-narrative (expository) text, with some studies showing a narrative advantage [10] and others showing no difference between genres [3]. Our study is the first, to our knowledge, to compare spatial representations constructed from narratives and non-narratives and to look at correlates of these abilities.

What is the reason for the observed advantage on the narrative task? There are three key ways in which the narrative and description versions differed. First, in the narrative, participants were presented with a character with specific goals (i.e., to deliver cookies) that motivated her to travel to the various locations. This may have given participants the opportunity to construct a multidimensional situation model that supported the construction of their spatial situation models. Perhaps including goals, characters, and actions through space scaffolds the construction of a spatial representation. In other words, reading or listening to a narrative may engage a different set of cognitive processes than a non-narrative [2].

Second, the narrative invited participants to take a perspective *within* the narrative, whereas the description may have encouraged participants to take more of a bird's eye view of the space. This distinction would be similar to that between route and survey perspectives, respectively. If participants are inclined to step into characters' shoes, as suggested by previous studies (e.g., [5, 15, 19]), perhaps they take on something of a route perspective. However, previous studies have demonstrated that adults' spatial representations are more or less the same whether they are derived from survey or route descriptions [18, 22]. It is possible that the same pattern will not hold for children. Follow-up studies will investigate the underlying reasons for the advantage on the narrative task my manipulating factors such as the opportunity to take a perspective within the space.

Third, it may simply be the case that participants found the narrative more interesting, which served to maintain their attention. This explanation would suggest that the unique characteristics of the narrative, such as the opportunity to construct a multidimensional situation model or take a character's perspective, were not driving the effect, but rather that the narrative yielded superior performance because it was more engaging. Although this explanation cannot be entirely ruled out, there are two potential problems with it. If it were simply a matter of participants in one condition devoting more attention to the task than those in the other, one may have expected to see effects of passage length. The description version was substantially shorter than the narrative version. Additionally, interest is an inherently subjective matter. Indeed, children in both conditions reported enjoying the passages and the accompanying activity quite frequently.

The results of the correlational and regression analyses are most in line with the first interpretation. It is intriguing that performance on the narrative version of the task was most strongly associated with narrative comprehension, whereas performance on the description version was most associated with working memory. This lends support to an explanation that suggests that different cognitive processes support construction of spatial representations based on the two types of language (narrative versus description). Successful construction of spatial representations from narratives may depend on the ability to build situation model representations, whereas successful construction from descriptions may depend more on holding a series of propositions in working memory.

The possibility that children in the description group held verbatim representations of the sentences in memory could explain why children with stronger verbal working memory abilities had an edge in the task. The finding that narrative comprehension was a significant predictor of performance on the narrative version, but not the description version could be interpreted in a few ways. Children with strong narrative comprehension skills may demonstrate such strength because they are better at constructing spatial (and other types of) situation models. That is, the ability to create detailed and accurate situation models may bolster children's comprehension. Or, children may require a certain level of competence in their comprehension abilities to be able to process the sentences they have heard, before they begin to construct a situation model. However, the narrative comprehension measure, but not the listening comprehension measure predicted children's performance on the narrative SSM task, suggesting that there was something unique about narrative comprehension abilities involved in task success. Of course, the children who performed well on both the narrative comprehension measure and the narrative SSM task may just have been those who enjoy stories more. However, there remains the intriguing possibility that the effect is due to another reason; children who are better at visualizing and creating spatial situation models may be better comprehenders because of it.

4.1   Conclusions

The findings of the present study lend support to the idea that a special mode of narrative thought exists distinctly from non-narrative thought [2]. When presented with the exact same spatial information in narrative or non-narrative formats, participants had differential success. Furthermore, performance in each condition was associated with different abilities. These findings suggest that constructing spatial representations from narrative is, on average, easier than constructing representations from descriptions, but also raise intriguing questions about why this may be the case.

─── **References** ───────────────────────────────────

1   John D. Bransford, J. Richard Barclay, and Jeffery J. Franks. Sentence memory: A constructive versus interpretive approach. *Cognitive Psychology*, 3(2):193–209, 1972.

2   Jerome Bruner. Actual minds, possible worlds. *Cambridge, MA: Harvard University*, 1986.

3   Lawrence J. Cunningham and M. D. Gall. The effects of expository and narrative prose on student achievement and attitudes toward textbooks. *Journal of Experimental Education*, 58:165–175, 1990.

4   Meredyth Daneman and Patricia A. Carpenter. Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior*, 19(4):450–466, 1980.

5   Agnieszka M. Fecica and Daniela K. O'Neill. A step at a time: Preliterate children's simulation of narrative movement during story comprehension. *Cognition*, 116(3):368–381, 2010.

6   Rebecca Fincher-Kiefer. Perceptual components of situation models. *Memory & Cognition*, 29(2):336–343, 2001.

7   Richard J. Gerrig. *Experiencing narrative worlds: On the psychological activities of reading.* Yale University Press, 1993.

8   Arthur M. Glenberg, Marion Meyer, and Karen Lindem. Mental models contribute to foregrounding during text comprehension. *Journal of Memory and Language*, 26(1):69–83, 1987.

9   Georg Jahn. Three turtles in danger: Spontaneous construction of causally relevant spatial situation models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(5):969–987, 2004.

10   Adrian S. Janit, Georgina S. Hammock, and Deborah S. Richardson. The power of fiction: Reading stories in abnormal psychology. *International Journal for the Scholarship of Teaching and Learning*, 5(1):1–14, 2011.

11   Philip N. Johnson-Laird. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press, 1983.

12   William Langston, Douglas C. Kramer, and Arthur M. Glenberg. The representation of space in mental models derived from text. *Memory & Cognition*, 26(2):247–262, 1998.

13   Susan C. Levine, Janellen Huttenlocher, Amy Taylor, and Adela Langrock. Early sex differences in spatial skill. *Developmental Psychology*, 35(4):940–948, 1999.

14   Daniel G. Morrow. Spatial models created from text. In Herre van Oostendorp and Rolf A. Zwaan, editors, *Naturalistic Text Comprehension*, pages 57–78, Norwood, NJ, 1994. Ablex.

15   Daniel G. Morrow, Steven L. Greenspan, and Gordon H. Bower. Accessibility and situation models in narrative comprehension. *Journal of Memory and Language*, 26(2):165–187, 1987.

16   Marie D. Neale. *Neale analysis of reading ability: Manual*. ERIC, 1999.

17   Daniela K. O'Neill and Rebecca M. Shultis. The emergence of the ability to track a character's mental perspective in narrative. *Developmental Psychology*, 43(4):1032–1038, 2007.

18   Walter Perrig and Walter Kintsch. Propositional and situational representations of text. *Journal of Memory and Language*, 24(5):503–518, 1985.

19   Jaime Rall and Paul L. Harris. In Cinderella's slippers? Story comprehension from the protagonist's point of view. *Developmental Psychology*, 36(2):202–208, 2000.

20   Michelle Scalise Sugiyama. Food, foragers, and folklore: The role of narrative in human subsistence. *Evolution and Human Behavior*, 22(4):221–240, 2001.

21   H. Lee Swanson, Kathryn F. Cochran, and Cynthia A. Ewers. Working memory in skilled and less skilled readers. *Journal of Abnormal Child Psychology*, 17(2):145–156, 1989.

22   Holly A. Taylor and Barbara Tversky. Spatial mental models derived from survey and route descriptions. *Journal of Memory and Language*, 31(2):261–292, 1992.

23   David H. Uttal, Joan A. Fisher, and Holly A. Taylor. Words and maps: developmental changes in mental models of spatial information acquired from descriptions and depictions. *Developmental Science*, 9(2):221–235, 2006.

24   Teun Adrianus Van Dijk and Walter Kintsch. *Strategies of discourse comprehension*. Academic Press New York, 1983.

25   Richard W. Woodcock, Kevin S. McGrew, and Nancy Mather. *Woodcock-Johnson® III Tests of Achievement*. Itasca, IL: Riverside Publishing, 2001.

26   Fenja Ziegler, Peter Mitchell, and Gregory Currie. How does narrative cue children's perspective taking? *Developmental Psychology*, 41(1):115–123, 2005.

27   Rolf A. Zwaan. Embodied cognition, perceptual symbols, and situation models. *Discourse Processes*, 28:81–88, 1999.

28   Rolf A. Zwaan, Joseph P. Magliano, and Arthur C. Graesser. Dimensions of situation model construction in narrative comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(2):386–397, 1995.

29   Rolf A. Zwaan and Gabriel A. Radvansky. Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2):162–185, 1998.

30   Rolf A. Zwaan, Gabriel A. Radvansky, Amy E. Hilliard, and Jacqueline M. Curiel. Constructing multidimensional situation models during reading. *Scientific Studies of Reading*, 2(3):199–220, 1998.

Linking Motif Sequences with Tale Types by Machine Learning

Nir Ofek[1], Sándor Darányi[2], and Lior Rokach[1]

1   Department of Information Systems Engineering
     Ben-Gurion University of the Negev
     Beer-Sheva, Israel
     {nirofek,liorrk}@bgu.ac.il
2   Swedish School of Library and Information Science
     University of Borås
     Borås, Sweden
     sandor.daranyi@hb.se

―――― Abstract ――――

Abstract units of narrative content called motifs constitute sequences, also known as tale types. However whereas the dependency of tale types on the constituent motifs is clear, the strength of their bond has not been measured this far. Based on the observation that differences between such motif sequences are reminiscent of nucleotide and chromosome mutations in genetics, i.e., constitute "narrative DNA", we used sequence mining methods from bioinformatics to learn more about the nature of tale types as a corpus. 94% of the Aarne-Thompson-Uther catalogue (2249 tale types in 7050 variants) was listed as individual motif strings based on the Thompson Motif Index, and scanned for similar subsequences. Next, using machine learning algorithms, we built and evaluated a classifier which predicts the tale type of a new motif sequence. Our findings indicate that, due to the size of the available samples, the classification model was best able to predict magic tales, novelles and jokes.

1   Introduction

Digital humanities and the emerging field of cultural analytics implement powerful multidisciplinary metaphors and methods to process texts in unprecedented ways. One of the new concepts is the reference to "narrative genomics" or "narrative DNA" – more and more authors point out similarities between sequences of genetic material building up living material, and those of literary units constituting "memetic", i.e., cultural products whose transmission can be traced by means of population genetics [29].

The idea that canonical sequences of content indicators constitute higher order content units has been pervading biology in the 20th century, and then slowly spilled over to other domains, prominently linguistics. Namely, strictly regulated strings of nucleotides constitute genes whereas canonical strings of genes amount to chromosomes. As a parallel, first the notion of indexing languages as sentence-like sequences of classification tags was born [27], then disciplinary sublanguages as content indicator chains were proposed [15], and finally,

treebanks suggested to manifest the linguistic genome [1], albeit at the cost of giving up canonical sequences for a more loose concept of syntax; that is, at this point in development anything can be considered a genome as long as the expression is sequential and well-formed, i.e., grammatical. As a latest development in this respect, recently Jockers claimed to study the 19th century literary genome of English novels, having extended the genetic metaphor to corpus linguistics to come up with new findings for cultural analytics [18].

Our current endeavor below relates to this latter tradition in the building, although following its own line of thought when considering tale types as canonical motif sequences [5, 6]. By motifs we mean abstracted, generic content tags which summarize segments of the plot. For a more detailed discussion of motifs and related considerations, see, e.g., [4].

As a next step, in this more technical approach to apply methodology from bioinformatics to problems of the literary genome, here we introduce machine learning to reveal the probabilistic scaffolding of tale types in terms of motif content. Whereas the idea is simple – if motifs are condensed expressions of multiple sentence content, then tale types "sum up" motif sequences to yield broader topics –, this first attempt still bears all the hallmarks of a dry run and therefore comes with a caveat.

This paper is organized as follows. Section 2 explains background considerations with special emphasis on formulaity and metadata. Section 3 discusses the research problem. Section 4 offers a brief introduction to sequence mining, with Section 5 outlining the methodology used in the experiment. Section 6 reports the results, whereas Section 7 sums up our conclusions with suggestions for future research. The appendix gives insight in the structure of the tale type catalog and the motif index used in the experiment.

2 Background considerations

2.1 Formulaity as a means of storyline preservation

It has been known for almost a hundred years that the oral communication of folklore texts often applies *formulaity* to help the singer remember his text [24, 25, 22]. Filed under different names, structural and formal investigations of tales [30, 26, 17] and myths, indeed mythologies, have proposed the same approach [21, 23]. Less known is the fact that linguistic evidence points in the same direction: as exemplified by a now famous study in immunology, scientific sublanguages, characteristic of subject areas, may use a formulaic arrangement of content elements in a sequential fashion for the presentation of experiments, results, and their discussion [15]. Formulae as storytelling aids abound in oral literature on all levels and in all genres; consult e.g., [22, 16] for various formulae in the genre of oral epics and [28] for the genre of fairy tale.

Several kinds of formulaity exist, ranging from short canonical phrases such as the *epitheton ornans* in Homeric epics, to longer ones used in orally improvised poetry, including canonical sequences of content elements and leading to story grammars [20, 11] or narrative algebra [12, 13, 14]. We will focus on such sequences only.

To recall, according to the oral-formulaic theory developed by Milman Parry [24, 25] and Albert Lord [22], stock phrases could enable poets to improvise verse called orally improvised poetry. In oral composition, the story itself has no definitive text, but consists of innumerable variants, each improvised by the teller in the act of telling the tale from a mental stockpile of verbal formulas, thematic constructs, and narrative incidents. This improvisation is for the most part subconscious so that texts orally composed will differ substantially from day to day and from teller to teller. The key idea of the theory is that poets have a store of formulas (a formula being 'an expression which is regularly used, under the same metrical conditions,

to express a particular essential idea' [22]), and that by linking these in conventionalized ways, they can rapidly compose verse.

Such linking, however, seems to be pertinent to storytelling in prose as well. The following example displays a chain of motifs which characterize a particular tale type about supernatural adversaries:

> 300 *The Dragon-Slayer*. A youth acquires (e.g., by exchange) three wonderful dogs [B421, B312.2]. He comes to a town where people are mourning and learns that once a year a (seven-headed) dragon [B11.2.3.1] demands a virgin as a sacrifice [B11.10, S262]. In the current year, the king's daughter has been chosen to be sacrificed, and the king offers her as a prize to her rescuer [T68.1]. The youth goes to the appointed place. While waiting to fight with the dragon, he falls into a magic sleep [D1975], during which the princess twists a ring (ribbons) into his hair; only one of her falling tears can awaken him [D1978. 2].
>
> Together with his dogs, the youth overcomes the dragon [B11.11, B524.1.1, R111.1.3]. He strikes off the dragon's heads and cuts out the tongues (keeps the teeth) [H105.1]. The youth promises the princess to come back in one year (three years) and goes off.
>
> An impostor (e.g., the coachman) takes the dragon's heads, forces the princess to name him as her rescuer [K1933], and claims her as his reward [K1932]. The princess asks her father to delay the wedding. Just as the princess is about to marry the impostor, the dragon-slayer returns. He sends his dogs to get some food from the king's table and is summoned to the wedding party [H151.2]. There the dragon-slayer proves he was the rescuer by showing the dragon's tongues (teeth) [H83, H105.1]. The impostor is condemned to death, and the dragon-slayer marries the princess [32].

Square brackets refer to forkings in the plot where alternative motifs can result in valid tale variants (Figure 1).

What matters for our argumentation is that as much as a certain sequence of specific Proppian functions amounts to a fairy tale plot [26], it takes a certain linking of consecutive motifs to constitute a specific tale type. Extracting chains of symbolic content from text in the above sense is the formulaic representation of sentences as proposed by Harris *et al.* [15], bridging the gap between scientific sublanguages and so far unidentified agglomerations of sentences amounting to sequentially linked functions, motifs etc.

2.2   Metadata in folktale research

The case we want to test our working hypothesis on, outlined in Section 5, is the Aarne-Thompson-Uther Tale Type Catalog (ATU), a classification and bibliography of international folk tales [32]. In the ATU, tale types are defined as canonical motif sequences such that motif string A constitutes Type X, string B stands for Type Y, etc. Also, it is important to note that types were not conceived in the void, rather they extract the essential characteristic features of a body of tales from all corners of the world, i.e., they are quasi-formal expressions of typical narrative content, mapped from many to one.

ATU is an alphanumerical, basically decimal classification scheme describing tale types in seven major chapters (animal tales, tales of magic, religious tales, realistic tales [novelle], tales of the stupid ogre [giant, devil], anecdotes and jokes, and formula tales), with an extensive Appendix discussing discontinued types, changes in previous type numbers, new types, geographical and ethnic terms, a register of motifs exemplified in tale types, bibliography and abbreviations, additional references and a subject index.

ATU tale type 300: The Dragon-Slayer.
[B421 B312.2] B11.2.3.1 [B11.10 S262] T68.1 D1975 D1978.2 [B11.11 B524.1.1 R111.1.3] H105.1 K1933 K1932 H151.2 [H83 H105.1]
Sequence variants

B421	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	B11.11	H105.1	K1933	K1932	H151.2	H83
B421	B11.2.3.1	S262	T68.1	D1975	D1978.2	B11.11	H105.1	K1933	K1932	H151.2	H83
B421	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	B524.1.1	H105.1	K1933	K1932	H151.2	H83
B421	B11.2.3.1	S262	T68.1	D1975	D1978.2	B524.1.1	H105.1	K1933	K1932	H151.2	H83
B421	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	R111.1.3	H105.1	K1933	K1932	H151.2	H83
B421	B11.2.3.1	S262	T68.1	D1975	D1978.2	R111.1.3	H105.1	K1933	K1932	H151.2	H83
B421	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	B11.11	H105.1	K1933	K1932	H151.2	H105.1
B421	B11.2.3.1	S262	T68.1	D1975	D1978.2	B11.11	H105.1	K1933	K1932	H151.2	H105.1
B421	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	B524.1.1	H105.1	K1933	K1932	H151.2	H105.1
B421	B11.2.3.1	S262	T68.1	D1975	D1978.2	B524.1.1	H105.1	K1933	K1932	H151.2	H105.1
B421	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	R111.1.3	H105.1	K1933	K1932	H151.2	H105.1
B421	B11.2.3.1	S262	T68.1	D1975	D1978.2	R111.1.3	H105.1	K1933	K1932	H151.2	H105.1
B312.2	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	B11.11	H105.1	K1933	K1932	H151.2	H83
B312.2	B11.2.3.1	S262	T68.1	D1975	D1978.2	B11.11	H105.1	K1933	K1932	H151.2	H83
B312.2	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	B524.1.1	H105.1	K1933	K1932	H151.2	H83
B312.2	B11.2.3.1	S262	T68.1	D1975	D1978.2	B524.1.1	H105.1	K1933	K1932	H151.2	H83
B312.2	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	R111.1.3	H105.1	K1933	K1932	H151.2	H83
B312.2	B11.2.3.1	S262	T68.1	D1975	D1978.2	R111.1.3	H105.1	K1933	K1932	H151.2	H83
B312.2	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	B11.11	H105.1	K1933	K1932	H151.2	H105.1
B312.2	B11.2.3.1	S262	T68.1	D1975	D1978.2	B11.11	H105.1	K1933	K1932	H151.2	H105.1
B312.2	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	B524.1.1	H105.1	K1933	K1932	H151.2	H105.1
B312.2	B11.2.3.1	S262	T68.1	D1975	D1978.2	B524.1.1	H105.1	K1933	K1932	H151.2	H105.1
B312.2	B11.2.3.1	B11.10	T68.1	D1975	D1978.2	R111.1.3	H105.1	K1933	K1932	H151.2	H105.1
B312.2	B11.2.3.1	S262	T68.1	D1975	D1978.2	R111.1.3	H105.1	K1933	K1932	H151.2	H105.1

Figure 1 *300 The Dragon Slayer* as a motif chain and its equally valid story variants.

The numbering of the types runs from 1 to 2399. Individual type descriptions uniformly come with a number, a title, an abstract-like plot mostly tagged with motifs, known combinations with other types, technical remarks, and references to the most important literature on the type plus its variants in different cultures. At the same time, as the inclusion of some 250 new types in the Appendix indicates, tale typology is a comprehensive and large-scale field of study, but also unfinished business: not all motifs in the *Motif Index* [30] were used to tag the types, difficulties of the definition of a motif imposed limitations on its usability in ATU, and narrative genre related considerations related to classification in general had to be observed.[1]

To turn to Thompson's *Motif-Index*, it offers worldwide coverage of folk narrative. As Alan Dundes suggested, in spite of its shortcomings, "It must be said at the outset that the six-volume *Motif-Index of Folk-Literature* and the Aarne-Thompson tale type index constitute two of the most valuable tools in the professional folklorist's arsenal of aids for analysis. This is so regardless of any legitimate criticisms of these two remarkable indices, the use of which serves to distinguish scholarly studies of folk narrative from those carried out by a host of amateurs and dilettantes. The identification of folk narratives through motif and/or tale type numbers has become an international *sine qua non* among bona fide folklorists. For this reason, the academic folklore community has reason to remain eternally grateful to Antti Aarne (1867–1925) and Stith Thompson (1885–1976) who twice revised Aarne's original 1910 *Verzeichnis der Märchentypen*—in 1928 and in 1961—and who compiled two editions of the *MotifIndex* (1922–1936; 1955–1958)" [9].

In appendices A and B we give an overview of the structure in ATU and the structure of a sample class of motifs.

---

[1] Hans-Jörg Uther, personal communication.

3   Research problem

Both transmitted genetic content and transmitted text content undergo variation over time. Genetic variation is called mutation and affects, e.g., nucleotides, amino acids, genes etc. Text variation does not have a specific name. Narrative elements that can vary include motifs (i.e., abstracted, generic content tags which summarize segments of the plot). Motif chains are the "backbones" of tale types, clusters of multilingual texts with related content. As motif insertion, deletion, and crossover were demonstrated to exist in tale types [5], types of mutation known from genetics apparently also occur in storytelling.

With the above observations about formulaity in oral tradition in mind, and to use the terminology of Dawkins [7], given the phenomenon of text variation in folklore, the existence of tale motifs and tale types on a global scale is universal evidence for semantic content resisting erosion, i.e., meme loss. This stability of memetic products invites the study of the relationship between two forms of memes, tag content vs. type content as a classification problem, the relationship between the features of a class and their sum total reflected in a set of documents being a major research issue well beyond folklore research. Therefore to ask about the interplay between genres like animal tales, and the content of motifs which build up such tales so that the result ends up in that genre, justifies one's curiosity. Put another way, this time we were interested in the correlation between two respective semantic fields [31], one described by tale types, the other by thematic motif groups and subgroups.

The stability of content sequences is documented in different corners of text research, lately for example by (Danescu-Niculescu-Mizil *et al.* [3]) who studied the memorability of phrases. In this particular field, Darányi and Forró [5] have shown that motifs are not the ultimate level of tale content available for indexing. In a sample of 219 tale types over 1202 motifs (ATU 300-745A, "Tales of magic" segment), their semiautomated analysis found granularity in ATU on two more levels, in the pattern of motif co-occurrences and in collocated motif co-occurrences, both apparently having been stable enough to resist text variation. On the other hand, Karsdorp *et al.* [19] have indicated that tale types in ATU show reasonably unique motif sequences whose subsequences are hardly ever repeated over different types.

With these considerations in mind, next we briefly introduce the data mining methodology we decided to apply to the problem.

4   Sequence mining by machine learning

Sequential pattern mining is a prevalent data mining approach [8]. The input of the learning process is a set of class labeled sequences (here tale types), which are used to train a model to predict the label of any unlabeled sequence. The learning process for classification uses the information of subsequences derived from the original sequences to discriminate class types. That is performed mainly by calculating a discrimination ratio based on statistics.

Sequence data include sequences of DNA, protein, customer purchase history, web surfing history, and more. Ferreira and Azevedo [10] used sequence mining in conjunction with a machine learning algorithm to classify protein sequences.

5   Methods

We considered motifs as entries in an indexing vocabulary and tale types as the document vectors in a corpus indexed by them, the latter being sparse motif strings which at the same time constitute "sentences", i.e., are predicates about type-specific tale content.

Figure 2 Tale type sequence lengths (in motifs).

We anticipated thematic dependencies between the 7 narrative genres defined as tale types in ATU (i.e., animal tales, tales of magic, religious tales, realistic tales [novelle], tales of the stupid ogre [giant, devil], anecdotes and jokes, and formula tales) vs. the 23 major motif groups in the *Motif Index* (mythological motifs, animals, taboo, magic, the dead, marvels, ogres, tests, the wise and the foolish, deceptions, reversal of fortune, ordaining the future, chance and fate, society, rewards and punishments, captives and fugitives, unnatural cruelty, sex, the nature of life, religion, traits of character, humor, and miscellaneous). The research question was, how do motifs from the above 23 groups constitute sequences resulting in those 7 genres? We assumed that by exploring the dependency structure of motifs vs. tale types, one can unveil the underlying probabilistic underpinnings of storyline construction.

5.1   Material

We used 94% of the complete ATU for this first experiment, i.e., out of the 2399 types we worked with 2249. The remaining 6% were left out from preprocessing because of their non-standard motif notation in the types, e.g., also containing running text in square brackets. Figure 2 shows the frequency of tale type length in terms of number of motifs in the string. Figure 3 displays the number of tale types and subtypes per genre.

It is an open question if due to text erosion or because of still being in a nascent stage, but many of the tale types consist of a single motif only. This undermines the very notion of tale type as a motif sequence [9]. In ATU, one-motif narratives are typical for anecdotes and jokes, formula tales and animal tales, whereas they are least characteristic for tales of magic, with the other genres statistically placed between them. Contrary to Karsdorp *et al.* [19], we feel that tale type is a bicomponential concept, having to satisfy both a formal and a topical constraint, and where the formal aspect, i.e., being a string, is met to a minimum only, topicality still prevails and accounts for the existence of genres grouping short texts; furthermore nothing prevents one from concatenating them in order to generate new narrative types with a higher dose of adventure than in anecdotes etc.

Figure 3 Number of tale types and subtypes per genre.

5.2   Preprocessing

In order to employ sequence mining in a conjunction with machine learning, a preprocessing stage is required. In tale types as motif strings, forkings in the plot may occur as alternative motifs which can be used as filler in particular loci in the plot (see Figure 1). To remove this obstacle, different respective sequences were treated as subtypes (motif string variants) of the same type with a renumbered identifier constructed from the type number and the variant number. This increased the number of 2249 published types to 7050 types and subtypes.

To reflect their similarities, every motif is encoded by a set of unique characters in a certain format in the *Motif Index*. For example, since motifs *"Blindness miraculously cured"* and *"Cripple marvelously cured"* share similar content, they are represented by a similar code having the same prefix letter and a similar number: F952 and F953, respectively. Not to overfit the learning process, we had to generalize motifs in tale types as their sequences and represent them in a less granular way. Decimal motif numbers were gradually truncated, from full notation to class tags only. Due to this process, each level of granularity reduction has fewer number of sequences; however, we do not regard them as new sets of tales, instead, each truncated motif sequence is still considered as another representation of its original sequence. In more details, for every motif, we employed two types of generalization: (1) On a first level we considered only the prefix letter and the integer number to the left of the dot symbol. By doing so, we aggregated similar tales to their father node in the *Motif Index*. In a similar manner, e.g., the first level representation of the two motifs B143.0.3 *"Owl as prophetic bird"* and B143.0.4 *"Raven as prophetic bird"* is generalized into motif B143 *"Prophetic bird"*. (2) Prefix generalization – every motif became represented by its prefix letter only; e.g., the above two motifs both were represented by the letter 'B'. As a result, each original sequence is represented by three sequences: the original, truncation to first level, and truncation to prefix. The principle of this process is displayed in Figure 4.

5.3   Constructing a dictionary of sequences

In machine learning, every instance (tale) is represented by a set of features. We hypothesized that features that are based on sequence frequencies are beneficial for the training process;

[K1555.2 K1554.1] K443.1

K1555.2 K443.1
K1554.1 K443.1
K1555 K443
K1554 K443
K1555.2
K443.1
K1554.1
K443.1
K1555
K443
K1554
K443
K K
K
k
K K
K
k

Generate all permutations

K1555.2 K443.1
K1554.1 K443.1

Generalizing

Original: K1555.2 K443.1
              K1554.1 K443.1
First level: K1555 K443
              K1554 K443
Prefix: K K
              K K

Extracting all sequences

Figure 4 Preprocessing of tale 1358A and extraction of its motif sequences.

therefore, in our work the set of features was extracted from a dictionary of sequence frequencies.

We extracted subsequences from every motif sequence (tale type) and stored their frequencies according to the classes they represented. Since motif sequences are relatively short, we chose to store all sequences of size 1 to 4. Thereby, we avoided storing relatively long sequences that seldom occur and are not likely to add any useful information to our analysis. This is done for each of the three types of sequence representation i.e., the original, and the two generalization rounds. The right side of Figure 4 details subsequence extraction. An example of dictionary entries is given in Table 1.

It is important to note that our dataset contains permutations of motif sequences, as explained in the pre-processing section. As such, they are dependent on their original motif sequence. For example the two permutations in Figure 4 both contain motif K443.1 in their

Table 1 Frequency of motif sequences per tale genre in the dictionary (excerpt).

sequence	tale type	frequency
B184.1 D961 B435.1 H1242	magic	0.0053
L161	magic	0.0498
L161	novelle	0.00182
J H	magic	0.00027
J H	jokes	0.0029
J H	novelle	0.6577

Table 2 Tale types in our corpus, after generating all permutations.

Tale type class	Number of instances
animal	349
formula	68
jokes	858
magic	4668
novelle	688
ogre	133
religious	286

second position. Given that, in our evaluation we ensure that permutations (subtypes) of any original tale type should be assigned to only one set – test or train. If a tale was sampled for the training set, then all of its permutations also belonged to the same set. Only the training set was used to generate the dictionary. Thus, the dictionary does not contain information based on tales from the test set, to avoid dependent tales being used for training.

5.4 Experiment design

After preprocessing and constructing the dictionary, we wish to train a classification model by using sequence discrimination ratios based on statistics.

The original dataset contains seven classes of tale types. Table 1 displays the number of instances for each.

The learning process of any machine learning algorithm requires to be provided with a sufficient number of sampled instances which represent the population of each class. Thus, an effective learning process can be employed. However, in our dataset, in some classes of tale types there are only few dozens of observed instances. Not only having insufficient number of training instances of some classes of tale types, the dataset is also imbalanced.

In a classification problem, class imbalance occurs when there are more examples of a certain class than of any other, on a large scale. For a variety of reasons, imbalanced datasets pose difficulties for induction algorithms [2]. The most obvious problem is that standard machine learning techniques are overwhelmed by the majority class and ignore the minority class. This problem is reflected in the phrase: like a needle in a haystack. Much more than the needle is being small, the problem is the fact that the needle is obscured by a huge number of strands of hay. Therefore, a classifier can achieve high accuracy by always predicting the majority class, particularly if the majority class constitutes most of the dataset, as for the 'magic' class type. Some class labels have a relatively low number of instances, sometimes down to a ratio of 69 less times than in other classes. Therefore, we experimented with those classes of types (i.e., genres) that have a more balanced number of instances.

In our first experiment we tried to discriminate between tales of two type classes, 'jokes' and 'novelle' which have a similar number of instances. The next experiment contains the 'magic' tale types, since this type has the largest number instances that allow an effective learning process. In addition, the mentioned two tale types ('jokes' and 'novelle') were selected as the closest in our dataset to the 'magic' type in terms of number of training instances.

In our classification task, each instance is a sequence of motifs. The goal was to train a model that can be used to predict the class of any unlabeled instance. The first step was the

Table 3 The dictionary as a lookup table of sequences and frequency-ratio, for any given class type.

sequence	tale type	frequency ratio
B184.1 D961 B435.1 H1242	magic	6.89
L161	magic	63.13
L161	novelle	0.041
J H	magic	0.0009
J H	jokes	0.0339
J H	novelle	911.9

extraction of features by which the instances could be represented. To take into consideration the order of motifs, but at the same time also to avoid a rigid structure of motif sequence, we segmented every sequence into several subsequences as detailed in the previous section. We used the subsequences dictionary in the following way. We calculated for every entry its likelihood in each class in contrast with all other classes. This ratio was calculated by dividing the frequency of the subsequences in a specific class by its frequency in all other classes. The calculated frequency ratio was stored in the dictionary, which now functions as a frequency-ratio lookup table (Table 3). We expected that a class whose value was high for a certain subsequence is more likely to be the tale type of a motif sequence that contained this subsequence.

In order to construct the dictionary, in the next step, for every instance we extracted all of its subsequences of size 1 to 4 as a pool of subsequences (see right side of Figure 4). By extracting subsequences of size 1 to 4, we took into consideration the order of the motif sequence, to some extent, while focusing on relatively short subsequences that would more likely to occur in motif sequences, and to avoid overfitting.

We computed two types of features for each class type. First, for all the subsequences from the instance's pool, we attached their frequency-ratio from the lookup table. Then, we sorted them and got the ratio of the highest score. This is the first feature type, which is called 'top 1'. The second is an accumulation of the top three ratios, denoted as 'top 3'. This was repeated for the original motif sequence, and for both generalization rounds of the motif identifiers, i.e., beyond the original, truncation to first level, and truncation to prefix. The total number of features is given by: 2 (*top 1 and top 3 ratios*) × 3 (*generalization levels*) × *class types* (Figure 4). By using abstracted top ratios features, we try to avoid overfitting, as could be the case in a bag-of-motifs feature space approach.

We believe that features that are strongly related to the actual (true) class type of the tale will have higher values than the same features for other class types.

In each experiment we split the dataset into two sets: 80% of the examples were used for training and 20% for testing. The dictionary and its subsequences statistics were also constructed only according to the training instances.

We trained a classification model based on the training set to predict the actual class of each tale. We evaluated several types of machine learning classification algorithms that we find adequate for that task. We chose to display results for the Bayes Network classifier since it yielded the best results, and for a decision tree as a comparison and for illustrating its interpretable model. To train the models and perform the experiments, the WEKA machine learning program suite was used [33].

■ **Table 4** A set of 18 features calculated for each instance.

	Top 1 ratio			Top 3 ratios		
Class / Generalization	Magic	Jokes	Novelle	Magic	Jokes	Novelle
Original	88.5	0	106.60	88.5	0	106.6
First level	118.5	0	136.6	118.5	0	136.6
Prefix	1208	0.11	985	1208	0.11	985

6 Results

The best results are given by Bayes Network classifier. Table 5 details the performance.

In the next experiment we added the 'magic' class instances as well, and the task is to discriminate among the three class types. The results are given by Table 3. The best performance is given by using our approach with the Bayes Network algorithm which outperforms the decision tree algorithm as it yielded better result in more tale types, and across all measurements. We compared our methodology with a baseline. In the baseline, we used a bag-of-words approach, i.e., each tale is represented by a feature vector which is its set of motifs. The best results for the baseline were given after generalizing the motifs to their first-level and by using a decision tree classifier. Our approach outperforms the baseline by all measurements. That is since our approach uses abstracted features, and the baseline uses a set of nearly 2000 features (motifs) that might cause an overfitting. Since the 'magic' class has a substantially higher number of instances, and in order to show performance on each tale type separately, we evaluated each class separately and not the weighted average of the measurements. We analyze the error of the triplet classification experiment. On average, the normalized error rate by motif sequence length is 4.7%, taking into consideration only prominent lengths. Tales of lengths one, four, six and twelve were the most difficult to classify, and resulted in 7%–9% classification errors. Jokes and magic tales were confused

■ **Table 5** Results for binary class experiment. The Bayes Network classifier outperformed the decision tree by F-measure and AUC for both tale types. Best results for each tale type are in bold.

Classifier	Class	Precision	Recall	F-measure	AUC
Bayes Network	novelle	**0.912**	0.601	**0.725**	**0.867**
	jokes	0.749	**0.953**	**0.839**	**0.867**
Decision Tree	novelle	0.463	**0.964**	0.626	0.537
	jokes	**0.783**	0.105	0.185	0.537

Table 6 Results for trinary class experiment. The Bayes Network classifier is found to be superior. Best results for each tale type are in bold.

classifier	class	Precision	Recall	F-measure	AUC
Bayes Network	magic	**0.97**	0.796	0.875	**0.935**
	novelle	0.223	**0.767**	**0.345**	**0.799**
	jokes	**0.844**	**0.409**	**0.551**	0.864
Decision Tree	magic	0.803	**0.983**	**0.884**	0.913
	novelle	**0.6**	0.175	0.271	0.614
	jokes	0.55	0.069	0.123	**0.921**
Decision Tree	magic	0.883	0.056	0.105	0.515
	novelle	0.0	0.0	0.0	0.506
*baseline	jokes	0.136	0.981	0.239	0.518

to be novelles, at almost all error cases; novelles confused to be magic for 70% of its errors. That is for the Bayes Network classifier. However, for the decision tree there is not a same tendency, therefore we can not state that tales are mostly confused to be novelles.

Since the decision tree model is easy to be described and is interpretable, we will explain its structure. The generated tree is a directed graph that consists of a root node (a starting point), internal nodes (nodes that are pointed at and point to other nodes) and leaves (ending points). During the classification process, the classified item "travels" from the root to one of the leaves, where a classification decision is made. Figure 5 illustrates a sub-tree of the generated binary decision tree model. If the 'top 1' subsequence ratio for class 'jokes' in the original pool of subsequences is greater than 4.621245 and the 'top 1' ratio of for class 'magic' in the first level pool of subsequences is not greater than 5.174637 and the 'top 1' ratio for class 'novelle' in the prefix pool of subsequences is not greater than 2.724305, then the instance is of class 'jokes'. The support for this decision is 22/2, based on the training instances.

7   Conclusion and future research

Considering the existence of "narrative DNA", we used sequence mining methods used in bioinformatics to learn more about the nature of tale types as a corpus. 94% of the Aarne-Thompson-Uther catalogue (2249 tale types in 7050 variants) was analyzed as individual motif strings based on the *Motif Index* and scanned for similar subsequences. Next, using a machine learning classification algorithm, we built and evaluated a classifier which predicts the tale type of a new motif sequence. Our findings indicate that the probabilistic underpinnings of tale types by motif co-occurrences are robust enough to develop the classification model which, on this instance, was able to predict motif strings characterizing magic tales, novelles and jokes. We plan to continue this work and combine our framework with sequence transformation analysis to learn more about the DNA-like nature of narrative content.

...

original_top1_class_jokes > 4.621245

yes   no

firstLevel_top1_class_magic <= 5.174637   ...

yes   no

prefix_top1_class_novelle <= 2.724305   ...

yes   no

jokes (22/2)   ...
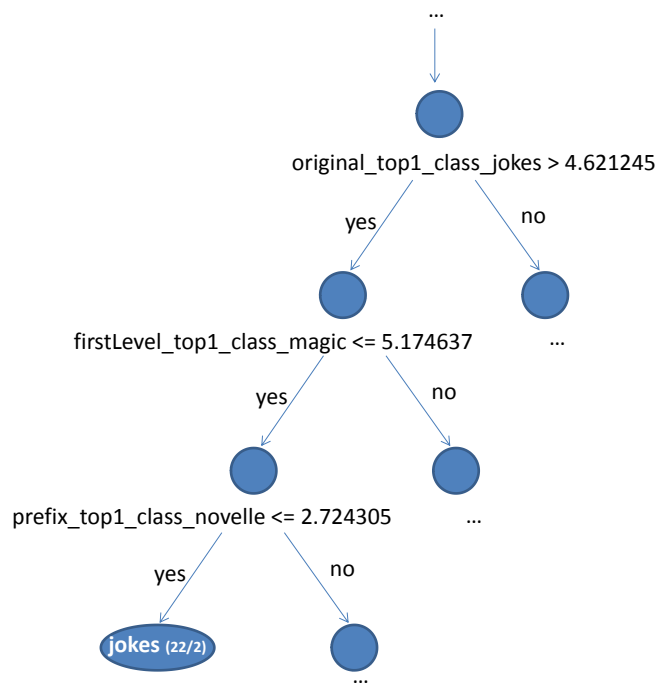
Figure 5 A sub-tree of a decision tree model, from the root note to one of its decision leaves.

References

1 M. Berti. The Ancient Greek and Latin dependency treebanks. blog post, `http://www.monicaberti.it/2010/10/the-ancient-greek-and-latin-dependency-treebanks/`, 2010.

2 N. V. Chawla, N. Japkowicz, and A. Kołcz. Editorial: special issue on learning from imbalanced data sets. *ACM SIGKDD Explorations Newsletter*, 6(1):1–6, 2004.

3 C. Danescu-Niculescu-Mizil, J. Cheng, J. Kleinberg, and L. Lee. You had me at hello: How phrasing affects memorability. Preprint, 2012.

4 S. Darányi. Examples of formulaity in narratives and scientific communication. manuscript, University of Szeged, Hungary, 2010.

5 S. Darányi and L. Forró. Detecting multiple motif co-occurrences in the Aarne-Thompson-Uther tale type catalog: A preliminary survey. *Anales de Documentación*, 15(1), 2012.

6 S. Darányi, P. Wittek, and L. Forro. Toward sequencing "narrative DNA": Tale types, motif strings and memetic pathways. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 2–10, İstanbul, 2012.

7 R. Dawkins. *The Selfish Gene*. Oxford University Press, Oxford, 1976.

8 G. Dong. *Sequence data mining*. Number 33 in Advances in Database Systems. Springer, 2009.

9 A. Dundes. The motif-index and the tale type index: A critique. *Journal of Folklore Research*, 34(3):195–202, 1997.

10 P. Ferreira and P. Azevedo. Protein sequence classification through relevant sequence mining and Bayes classifiers. In Carlos Bento, Amílcar Cardoso, and Gaël Dias, editors, *Progress in Artificial Intelligence, 12th Portuguese Conference on Artificial Intelligence,*

EPIA 2005, Covilhã, Portugal, December 5–8, 2005, Proceedings, number 3808 in Lecture Notes in Computer Science, pages 236–247. Springer, 2005.

11 A. Garnham. What is wrong with story grammars. *Cognition*, 15:145–154, 1983.

12 M. Griffin. An expanded, narrative algebra for mythic spacetime. *Journal of Literary Semantics*, 30:71–82, 2001.

13 M. Griffin. More features of the mythic spacetime algebra. *Journal of Literary Semantics*, 32:49–72, 2003.

14 M. Griffin. Mythic algebra uses: Metaphor, logic, and the semiotic sign. *Semiotica*, 158–1/4:309–318, 2006.

15 Z. S. Harris, M. Gottfried, T. Ryckman, P. Mattick, A. Daladier, T. N. Harris, and S. Harris. *The form of information in science: analysis of an immunology sublanguage.* Kluwer, Dordrecht, 1989.

16 H. Jason. *Motif, Type and Genre. A Manual for Compilation of Indices and A Bibliography of Indices and Indexing.* Academia Scientiarum Fennica, Helsinki, 2000.

17 H. Jason and D. Segal, editors. *Patterns in oral literature.* Mouton, The Hague, 1977.

18 M. L. Jockers. *Macroanalysis: Digital methods and literary history.* University of Illinois Press, Champaign, Il., 2013.

19 Folgert Karsdorp, Peter Van Kranenburg, Theo Meder, Dolf Trieschnigg, and Antal Van den Bosch. In search of an appropriate abstraction level for motif annotations. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 22–26, İstanbul, 2012.

20 G. P. Lakoff. Structural complexity in fairy tales. *The Study of Man*, I:128–190, 1972.

21 C. Lévi-Strauss. *Mythologiques I–IV.* Plon, Paris, France, 1964–1971.

22 A. Lord. *The singer of tales.* Harvard University Press, Cambridge, 1960.

23 P. Maranda. *The double twist: from ethnography to morphodynamics.* University of Toronto Press, Toronto, 2001.

24 M. Parry. Studies in the epic technique of oral verse-making. I: Homer and Homeric style. *Harvard Studies in Classical Philology*, 41:73–143, 1930.

25 M. Parry. Studies in the epic technique of oral verse-making. II: The Homeric language as the language of an oral poetry. *Harvard Studies in Classical Philology*, 43:1–50, 1930.

26 Vladimir Yakovlevich Propp. *Morphology of the Folktale.* University of Texas Press, Austin, TX, 2nd edition, 1968. transl. L. Scott.

27 S. R Ranganathan. *The Colon Classification Vol. 4.* Rutgers Graduate School of Library Service, New Brunswick, NJ, 1965.

28 N. Roshianu. *Traditsionnuie formuly skazki (Traditional formulae of the fairy tale).* Moscow, 1974.

29 R. M. Ross, S. J. Greenhill, and Q. D. Atkinson. Population structure and cultural geography of a folktale in Europe. *Proceedings of the Royal Society B*, 280(1756):1471–2954, 2013.

30 S. Thompson. *Motif-index of Folk-Literature: a Classification of Narrative Elements in Folktales, Ballads, Myths, Fables, Medieval Romances, Exempla, Fabliaux, Jest-Books, and Local Legends. 6 volumes.* Indiana University Press, Bloomington, 2nd edition, 1955–1958.

31 J. Trier. Das sprachliche Feld. *Neue Jahrbücher für Wissenschaft und Jugendbildung*, 10:428–449, 1934.

32 H. J. Uther. *The types of international folktales: A classification and bibliography based on the system of Antti Aarne and Stith Thompson.* Academia Scientiarum Fennica, Helsinki, Finland, 2004.

33 I. H. Witten and E. Frank. *Data Mining: Practical machine learning tools and techniques.* Morgan Kaufmann, 2005.

A   Appendix: The Types of International Folktales (from [32])

1. ANIMAL TALES (1–299)
Wild Animals 1–99
    The Clever Fox (Other Animal) 1–69
    Other Wild Animals 70–99
Wild Animals and Domestic Animals 100–149
Wild Animals and Humans 150–199
Domestic Animals 200–219
Other Animals and Objects 220–299
    Birds 220–249
    Fish 250–253

2. TALES OF MAGIC (300–749)
Supernatural Adversaries 300–399
Supernatural or Enchanted Wife (Husband) or Other
Relative 400–459
    Wife 400–424
    Husband 425–449
    Brother or Sister 450–459
Supernatural Tasks 460–499
Supernatural Helpers 500–559
Magic Objects 560–649
Supernatural Power or Knowledge 650–699
Other Tales of the Supernatural 700–749

3. RELIGIOUS TALES (750–849)
God Rewards and Punishes 750–779
The Truth Comes to Light 780–799
Heaven 800–809
The Devil 810–826
Other Religious Tales 827–849

4. REALISTIC TALES (NOVELLE) (850–999)
The Man Marries the Princess 850–869
The Woman Marries the Prince 870–879
Proofs of Fidelity and Innocence 880–899
The Obstinate Wife Learns to Obey 900–909
Good Precepts 910–919
Clever Acts and Words 920–929
Tales of Fate 930–949
Robbers and Murderers 950–969
Other Realistic Tales 970–999

5. TALES OF THE STUPID OGRE (GIANT, DEVIL) (1000–1199)
Labor Contract 1000–1029
Partnership between Man and Ogre 1030–1059
Contest between Man and Ogre 1060–1114
Man Kills (Injures) Ogre 1115–1144
Ogre Frightened by Man 1145–1154
Man Outwits the Devil 1155–1169
Souls Saved from the Devil 1170–1199

6. ANECDOTES AND JOKES (1200–1999)
Stories about a Fool 1200–1349

Stories about Married Couples 1350–1439
                The Foolish Wife and her Husband 1380–1404
                The Foolish Husband and his Wife 1405–1429
                The Foolish Couple 1430–1439
Stories about a Woman 1440–1524
                Looking for a Wife 1450–1474
                Jokes about Old Maids 1475–1499
                Other Stories about Women 1500–1524
Stories about a Man 1525–1724
                The Clever Man 1525–1639
                Lucky Accidents 1640–1674
                The Stupid Man 1675–1724
Jokes about Clergymen and Religious Figures 1725–1849
                The Clergyman Is Tricked 1725–1774
                Clergyman and Sexton 1775–1799
                Other Jokes about Religious Figures 1800–1849
Anecdotes about Other Groups of People 1850–1874
Tall Tales 1875–1999

7. FORMULA TALES (2000–2399)
Cumulative Tales 2000–2100
Catch Tales 2200–2299
Other Formula Tales 2300–2399

B   Appendix: Excerpt from Thompson's Motif Index [30]

B0–B99.Mythical animals
        B10. Mythical beasts and hybrids
        B20. Beast-men
        B30. Mythical birds
        B40. Bird-beasts
        B50. Bird-men
        B60. Mythical fish
        B70. Fish-beasts
        B80. Fish-men
        B90. Other mythical animals
B100–B199. Magic animals
   B100–B119. Treasure animals
        B100. Treasure animals–general
        B110. Treasure-producing parts of animals
   B120–B169. Animals with magic wisdom
        B120. Wise animals
        B130. Truth-telling animals
        B140. Prophetic animals
        B150. Oracular animals
        B160. Wisdom-giving animals
   B170–B189. Other magic animals
        B170. Magic birds, fish, reptiles, etc.
        B180. Magic quadrupeds
        B190. Magic animals: miscellaneous motifs
B200–B299. Animals with human traits
        B210. Speaking animals
        B220. Animal kingdom (community)

B230. Parliament of animals
B240. King of animals
B250. Religious animals
B260. Animal warfare
B270. Animals in legal relations
B280. Animal weddings
B290. Other animals with human traits
B300–B599. Friendly animals
B300–B349. Helpful animals–general
B310. Acquisition of helpful animal
B320. Reward of helpful animal
B330. Death of helpful animal
B340. Treatment of helpful animal—miscellaneous
B350–B399. Grateful animals
B360. Animals grateful for rescue from peril of death
B370. Animals grateful to captor for release
B380. Animals grateful for relief from pain
B390. Animals grateful for other kind acts
B400–B499. Kinds of helpful animals
B400–B449. Helpful beasts
B400. Helpful domestic beasts
B430. Helpful wild beasts
B450. Helpful birds
B470. Helpful fish
B480. Helpful insects
B490. Other helpful animals
B500–B599. Services of helpful animals
B500. Magic power from animal
B510. Healing by animal
B520. Animals save person's life
B530. Animals nourish men
B540. Animal rescuer or retriever
B550. Animals carry men
B560. Animals advise men
B570. Animals serve men
B580. Animals help men to wealth and greatness
B590. Miscellaneous services of helpful animals
B600–B699. Marriage of person to animal
B610. Animal paramour
B620. Animal suitor
B630. Offspring of marriage to animal
B640. Marriage to person in animal form
B650. Marriage to animal in human form
B700–B799. Fanciful traits of animals
B710. Fanciful origin of animals
B720–B749. Fanciful physical qualities of animals
B720. Fanciful bodily members of animals
B730. Fanciful color, smell, etc. of animals
B740. Fanciful marvelous strength of animals
B750. Fanciful habits of animals
B770. Other fanciful traits of animals
B800–B899. Miscellaneous animal motifs
B870. Giant animals

Character Networks for Narrative Generation: Structural Balance Theory and the Emergence of Proto-Narratives

Graham Alexander Sack

English and Comparative Literature Department
Columbia University
New York, NY, USA
`gas2117@columbia.edu`

──── **Abstract** ────

This paper models narrative as a complex adaptive system in which the temporal sequence of events constituting a story emerges out of cascading local interactions between nodes in a social network. The approach is not intended as a general theory of narrative, but rather as a particular generative mechanism relevant to several academic communities: (1) literary critics and narrative theorists interested in new models for narrative analysis, (2) artificial intelligence researchers and video game designers interested in new mechanisms for narrative generation, and (3) complex systems theorists interested in novel applications of agent-based modeling and network theory. The paper is divided into two parts. The first part offers examples of research by literary critics on the relationship between social networks of fictional characters and the structure of long-form narratives, particularly novels. The second part provides an example of schematic story generation based on a simulation of the structural balance network model. I will argue that if literary critics can better understand sophisticated narratives by extracting networks from them, then narrative intelligence researchers can benefit by inverting the process, that is, by generating narratives from networks.

1 Introduction

Throughout this paper, I will make extensive use of concepts from social network analysis and structural balance theory. The basic unit of analysis in structural balance theory is the *triad*, defined as a triangular configuration of friendship and enmity ties between three mutually connected nodes. Some triad configurations are socially unstable and, when embedded in networks with many interdependencies, may trigger cascading social events. These cascades, I will argue below, can be treated as a proto-narrative—the skeleton of a story from which complex social dramas may be constructed.

My approach to narrative is loosely inspired by a variety of sources.

In *Deceit, Desire, and the Novel* (1961), literary critic René Girard argues that a defining feature of the novel as a modern story-telling form is the way characters' desires are embedded in and mediated by indirect social relations. Taking Cervantes' *Don Quixote* as the prototype for the modern novel, Girard asserts that Quixote's desire is "triangular":

The straight line is present in the desire of Don Quixote, but it is not essential. The mediator is there, above that line, radiating toward both the subject and the object. The spatial metaphor that expresses this triple relationship is obviously the triangle. [10]

Girard argues that in novels such as *Don Quixote*, the subject desires the object not because of its inherent qualities, but because some third character, the mediator, also desires the object. The mediator may be a role model who the subject intentionally imitates, or he may be a character against whom the subject competes for the desired object, as with a rival in polite society, romance, or commerce. In either case, the dyadic relationship between subject and object cannot be understood without reference to the mediator. While I will not use the concept of "triangular desire" directly, I will adapt Girard's general interpretive framework by drawing a parallel between the idea that triads of characters are the basic unit of narrative analysis and structural balance theory's assumption that triads of nodes are the basic unit of network analysis.

In *Identity and Control: How Social Formations Emerge* (2008), network theorist Harrison White posits that "social networks emerge only as ties mesh with stories" [25]. White suggests that social ties ought to be thought of multi-dimensionally in terms of "netdoms": "'dom' from domain of topics and 'net' from network relations." When two social agents encounter one another, they struggle for recognition and status by "switch[ing] from netdom to netdom, finding footings in different networks in differing domain contexts." Two co-workers, for example, might initially relate to one another professionally, but then switch between political, religious, or even romantic domains as the social tie between them evolves. The more "netdom switchings" occur, the more complex and nuanced the relationship between the identities becomes. Over time, these switchings settle down into a stable tie that is comprehended by the participants via a "story". White summarizes the process as follows:

A story is a tie placed in context. Stories structure switchings into accounts with a beginning, middle, and end; so story-making frames social time... These relations are characterized by stories told in and about them with meanings drawn from the switchings between netdoms... A network can be traced as similar stories appear across a spread of dyads. [25]

White's notion of "netdom switchings" provides another potential link between structural balance theory and narrative. As the unstable triads in a structural balance network evolve towards stability, the edges connecting each pair of nodes undergo a simplified version of domain switching, oscillating between friendship and enmity. Following White's logic, the more frequently a link in a triad switches, the more complex and nuanced the relationship between the associated node-characters becomes, gradually forming a story-tie.

I instantiate the structural balance model in NetLogo, an interactive development environment for agent based modeling, and run the simulation forward in time under different parameter configurations, producing a range of possible proto-narrative event sequences that vary in length and outcome. My use of simulation to generate proto-narratives is influenced by the work of computational social scientists Joshua Epstein and Robert Axtell. In *Growing*

Artificial Societies (1996), the authors describe how their now canonical "Sugarscape" model, originally constructed to study the emergence of wealth inequality, can generate "proto-histories"—schematic social and cultural histories in which individuals agglomerate and form tribes, battle, and trade. Although they do not use the term, Epstein and Axtell argue, in effect, that each run of the Sugarscape model is analogous to an historical narrative.[1] Although simulations such as Sugarscape are now widely used in the social sciences, practitioners rarely if ever discuss the idea that the temporal progression of a simulation can be treated as a narrative. Epstein and Axtell emphasize moreover that they "grow this history 'from the bottom up."' [3] That is, their proto-historical narratives are complex adaptive systems (CAS) displaying the property of emergence.[2] Structural balance models provide a similarly CAS-based approach to narrative generation. Unstable triads update based on local stability rules, yet produce a narrative chain of events tracing the formation of global network structures (see the discussion of "social mitosis" below).

2 From Narratives to Networks

Over the past several years, literary critics have begun researching the relationship between social networks and narrative structure, including several efforts to extract character networks from literary works [8, 20, 22]. The guiding principle behind literary network analysis is that narratives are not merely depictions of individual experience in language but are also artificial societies whose imaginary social forms can be quantified and analyzed. What such analyses reveal is that narrative structure, such as plot, genre, and characterization, is intimately related to network structure.

Included below are examples drawn from my own research on literary networks. Figure 1 shows the sociograms[3] for several canonical European novels: Cervantes's *Don Quixote de la Mancha* (1605), Charles Dickens's *David Copperfield* (1850), and Virginia Woolf's *Mrs. Dalloway* (1925). Noticeable contrasts between these networks reflect key differences in literary conventions across historical periods and genres.

Don Quixote, widely regarded as the first modern novel, possesses an episodic plot structure derived in part from the picaresque story-forms popular during the Spanish Golden Age. It consists of a series of adventures that all feature the iconic knight and his squire but which are otherwise minimally connected in terms of plot. *Don Quixote*, moreover, operates as a frame narrative, encompassing many interpolated stories-within-the-story—such as Cardenio's autobiography or the pastoral poems describing Marcela. The sociogram reflects this disconnected plot structure. The network is centered on a main axis connecting Don Quixote and Sancho, which is embedded in a diffuse web of characters that extends outwards in several layers. For a relatively small network, it has a high diameter, indicative of the fact that it is effectively a network of networks: when Don Quixote meets another character, such

---

[1] In a chapter subtitled "The Emergence of History", Epstein and Axtell write, "The basic aim of this chapter is to "grow" a very simple caricature of history—a "proto-history" if you will... The social story is as follows: In the beginning, there is a small population of agents..." [3]

[2] "Emergent narrative" is an overloaded term. Researchers in interactive narrative use the phrase, along with "character based narrative", to refer to minimally plotted stories generated spontaneously through live user interaction (see Aylett 1999 [4]). My meaning is drawn from complex systems theory and refers to the emergence of systemic properties from local interactions. A connection could be drawn—one might argue that human interactions with non-player characters in virtual reality environments exhibit CAS-like behavior—but the terms are not trivially synonymous.

[3] Social networks are generated for each novel by running a 10-word window through the text and counting the number of times each pair of character names co-occurs. Edges are drawn for pairs with greater than 3 co-occurrences and are weighted by frequency.

(a) Don Quixote (Cervantes)
- Nodes: 30
- Graph Density: 0.26
- Clustering Coeff: 0.72
- Avg. Degree: 7.67
- Stdev. Degree: 5.85
- Stdev / Avg. Degree   0.76
- Avg. Shortest Path: 1.89
- Diameter: 4

(b) David Copperfield (Charles Dickens)
- Nodes: 83
- Graph Density: 0.07
- Clustering Coeff: 0.80
- Avg. Degree: 5.81
- Stdev. Degree: 9.05
- Stdev / Avg. Degree   1.56
- Avg. Shortest Path: 2.03
- Diameter: 4

(c) Mrs. Dalloway (Virginia Woolf)
- Nodes: 24
- Density: 0.54
- Clustering: 0.79
- Avg. Degree: 12.33
- Stdev. Degree: 5.02
- Stdev / Avg Degree   0.41
- Avg. Shortest Path: 1.49
- Diameter: 3

Figure 1 Character Network Sociograms for three illustrative novels.

as Cardenio, often this new acquaintance temporarily assumes the role of narrator, relating a micro-narrative with its own stand-alone character network not linked to the original action. We can see this most clearly in the presence of several peripheral cliques, particularly the one between Anselmo, Lothario, Camilla, and Leonela, whose tale appears in a found manuscript read aloud by Quixote's priest. Moreover, the network contains a low proportion of "strong ties": aside from the heavily weighted edge between the knight and his squire, most edges are thin and light, indicating brief, glancing interactions with secondary characters who provide color and variety in particular episodes but who rarely recur or interact significantly with one another.

David Copperfield is one of the preeminent examples of the 19th Century *Bildungsroman*, depicting the education and development of its eponymous protagonist as he finds his way in the world, seeking benefactors, a career, and a marriage partner. Like many serialized mid-Victorian novels, it features an expansive cast: the network consists of 83 nodes, with a high proportion of isolates and low graph density (7%). The network is highly centralized with an obvious star-shape, reflecting an egocentric focus on its protagonist, who is the hub for virtually all character interactions. This accounts for its very high standard deviation in node degree, indicative of inequality in connectedness and social importance. The prevalence of strong ties is noticeably greater than in *Don Quixote*, illustrative of the 19th Century *Bildungsroman's* concern with complex relationships developed over a long duration, rather than the brief encounters with strangers common in the Renaissance-era picaresque.

The character network for *Mrs. Dalloway* contrasts noticeably with the others. While mid-Victorian novels often featured sprawling casts, the network for *Mrs. Dalloway*, a canonical work of high British modernism, is delimited. The focus is on psychological depth rather than sociological breadth. The sociogram consists of a single large component with no isolates and a very high graph density (54%) and clustering coefficient (79%), reflecting embedded relationships between characters with many common social ties. Unlike *David Copperfield*, the narrative is not singularly focused on its title character: point of view shifts approximately every ten pages. The network, correspondingly, does not have a pronounced center: it is bifurcated into two cliques—one concentrated around Clarissa Dalloway and the other around her narrative alter-ego, Septimus. Moreover, the network has a low standard deviation in node degree, indicating that character interaction is broadly and evenly distributed across the ensemble. Lastly, the diagram exhibits a high proportion of strong ties. The overall picture is that of a tightly knit social world focused on the intimate relationships between a small set of equally significant characters.

As these brief examples suggest, there is a close association between narrative structure and network structure. Authorial decisions related to linear vs. episodic plot or the balance of focus between protagonist and ensemble are visible in network properties such as centralization, graph density, diameter, clustering, and prevalence of strong vs. weak ties. But if literary critics can better understand sophisticated narratives by extracting networks from them, perhaps narrative intelligence researchers can benefit by inverting the process, that is, by generating narratives from networks.

3 From Networks to Narratives

In the remainder of this paper, I offer a simple example of how social networks may be used to generate narratives. While the networks in the preceding section summarize character interactions, the networks that follow produce interactions. Descriptive and generative networks may at first appear quite different, but I will relate them by showing how structural balance networks produce event sequences that can then be converted back into descriptive sociograms analogous to those shown above. The direct connection will be established towards the end of this section.

3.1 Background: Structural Balance Model

The model I will describe is based on ideas from structural balance theory, also known as social balance theory. SBT was originated in the mid-1940s by Fritz Heider, who studied patterns of belief coherence in individual psychology [11]. In the mid-1950s, Cartwright and Harary generalized Heider's theory of coherence and applied it to social relations, representing

Figure 2 (a) Local triad stability rules. (b) Emergent globally stable network patterns.

stable and unstable configurations with basic graph theory [6]. SBT has since become a sub-branch of social network theory.

Consider a set of nodes representing, for example, people or countries. Each node may be joined to each other node by an edge, which represents their relationship. If two nodes are joined, they are either (1) friends or (2) enemies. The fundamental unit of analysis in SBT is a triad of three mutually linked nodes. A triad is considered unstable if there is social pressure to change one of the relationship links. It is considered stable if there is no social pressure to change.

Let $(+)$ represent friendship and $(-)$ represent enmity. There are several possible configurations:

1. $(+)(+)(+)$: If all 3 nodes are friends / allies, the triad is considered stable.
2. $(-)(-)(-)$: If all 3 nodes are enemies, the triad is unstable, since two nodes have an incentive to ally against the third (thereby becoming friends with each other).
3. $(+)(+)(-)$ or $(+)(-)(+)$ or $(-)(+)(+)$: If one node is friends with two that are enemies with one another, it will be pressured to pick a side, and therefore the triad is unstable.
4. $(+)(-)(-)$ or $(-)(+)(-)$ or $(-)(-)(+)$: If two nodes are friends with each other and both are enemies against a third, the triad is stable.

The rule for stability can be summarized as follows: a triad is stable if the multiplicative product of the signs is positive [6]. The stability of the various triads conforms to the following simplified social principles: (1) my friend's friend is my friend; (2) my friend's enemy is my enemy; (3) my enemy's friend is my enemy; (4) my enemy's enemy is my friend.

Note that changing any single link in an unstable triad will make it stable: $(-)(-)(-)$ becomes any cyclical permutation of $(-)(-)(+)$; $(+)(+)(-)$ becomes either $(+)(+)(+)$ or any cyclical permutation of $(-)(-)(+)$. Likewise, changing any link in a stable triad will make it unstable: $(+)(+)(+)$ becomes any cyclical permutation of $(-)(+)(+)$; $(-)(-)(+)$ becomes either $(-)(-)(-)$ or any cyclical permutation $(-)(+)(+)$. Local triad stability rules are illustrated in figure 2, part (a).

The stability of three mutually connected nodes is easy enough to evaluate, but the complexity increases as nodes are added to create larger graphs with many interdependent
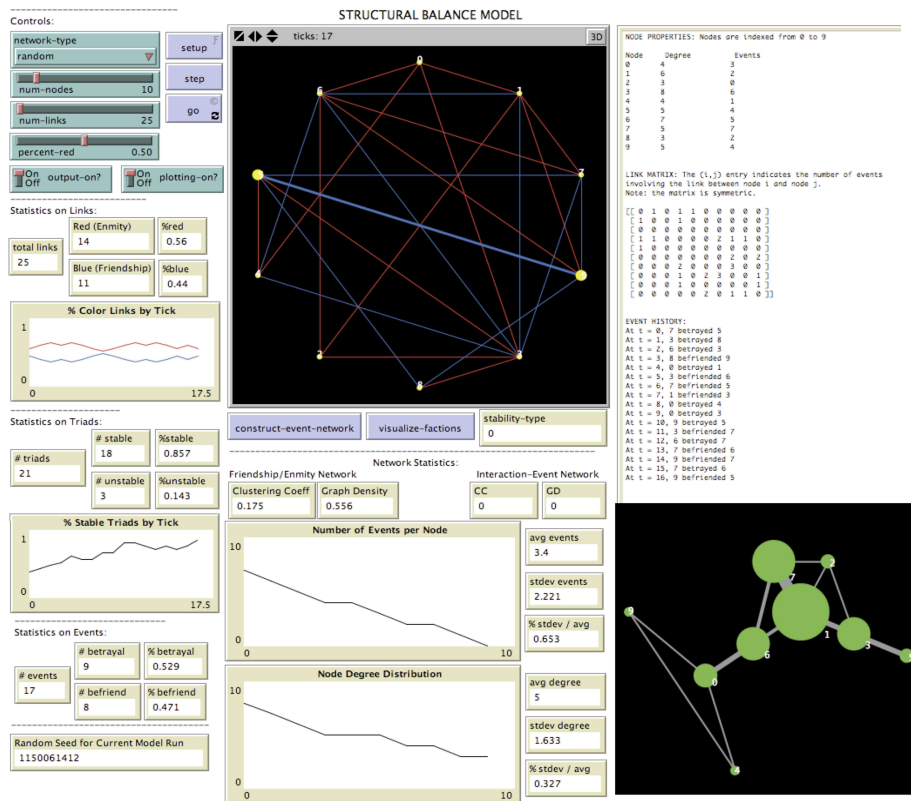
Figure 3 Synopsis of Model Run: (i) an unstable SBT network is constructed with user-specified number of nodes, number of links, and % red (enmity) vs. blue (friendship) links; (ii) as the model runs, link colors change generating an event history; (iii) the model halts when the percentage of stable-triads is 100%; (iv) a character-interaction network is generated from events between nodes.

triads. Nevertheless, global patterns emerge from the local interactions (see figure 2, part (b)). One such pattern is social mitosis: it can be proven that there are only two ways for a complete graph, i.e., one with no missing edges, to be structurally balanced: (1) everyone is friends (universal harmony); or (2) there are two factions of friends with total enmity between them (bi-polar factions).[4] For an incomplete graph, two more outcomes are possible: (3) nodes divide into three or more groups with total enmity between them (multi-polar factions); (4) some nodes are enemies but no factions form (mixed outcome).[5] While the general properties of the equilibrium state of any graph are deterministic, the dynamic process by which that graph reaches an equilibrium is not. This is what makes it interesting and useful as a narrative generation mechanism.

---

[4] For a simple proof of this theorem, see [7, chapter 5].

[5] A complete graph with n nodes, has $_nC_2 = n(n-1)/2$ edges, each of which can be in 2 states, $(+)$ or $(-)$; thus, there are $2^{n(n-1)/2}$ states for the network. It can be shown that $2^n - 1$ of these are stable outcomes. This corresponds to the number of ways to divide a group with n members into two factions of size m and $(n-m)$. The non-polarized solution ("universal harmony") is simply the trivial solution where $m = 0$. An incomplete graph has more stable configurations since multi-polar and mixed outcomes are permitted.

Avg. Event History Length vs. Percent-Red Links & #-Links

Avg. Event History Length vs. Network-Type & #-Links

* Data shown is based on 88,200 runs of the model – Averages are over 200 runs per input parameters configuration

■ **Figure 4** The length of the "event history", measured by the mean ticks to reach global stability, is affected by the SBT network's topology and number of links and the percentage of friendship vs. enmity ties.

3.2 Model Implementation

The next several pages describe a version of the structural balance model that I have implemented in NetLogo, an IDE for agent-based modeling; see figure 3 for the synopsis of a model run. This is not the first computer simulation of structural balance dynamics (see [12, 13, 24] for alternative implementations). The crucial difference is that my focus is not on structural balance in its own right, but rather on motivating a series of observations about the narrative generating potential of social-network-based simulations. Towards this end, my emphasis is on the proto-narratives generated by the model's dynamics, which are captured by the event history, node history, and relationship-link history discussed below.

At set-up, the user specifies the number of nodes, the number of links, the percentage of links that will be enmity, indicated in red, as opposed to friendship, indicated in blue, and the network's degree distribution, which may be uniform, random, or maximally unequal. With each time step, the model's algorithm checks whether there are any unstable triads. If so, the algorithm randomly selects one of the unstable triads and randomly changes one of its links from red to blue or from blue to red. Changing a link's color stabilizes the selected triad, but may inadvertently destabilize other triads. The model continues stepping forward in time until all triads in the graph have been made stable.[6]

As it runs, the model generates several types of output. First, it produces global network statistics, such as graph density and clustering coefficient. Since clustering coefficient measures

---

[6] As noted in footnote 5, an n-node complete graph has $2^n - 1$ stable outcomes. Because the simple algorithm implemented here proceeds through the random selection of unstable triads and the random flipping of links, there is a non-zero probability of transition between any two graph states. Therefore, the algortithm will halt in finite time. In alternative implementations, this is not necessarily so. For example, if we require that links only be flipped if doing so will immediately increase the number of stable triads, then it is possible for the graph to reach a "jammed state." For a discussion of such results, see [2].
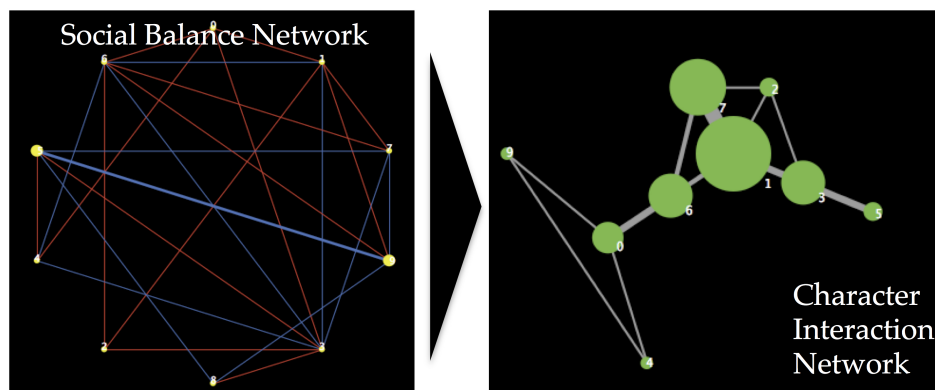
■ **Figure 5** The structural balance network generates character interaction events via link-switchings. The number of events between each pair of characters is represented in a separate character-interaction network.

the prevalence of complete triads, the greater its value, the more complex the structural balancing problem and the more time-steps on average required to reach stability. Second, the model tracks link and triad statistics. These include: (1) number and percentage of friendship vs. enmity ties and (2) number and percentage of stable vs. unstable triads. The program halts when the percentage of stable triads equals 100%. These metrics merely provide basic information about the state of the network. Of greater relevance for narrative generation are the outputs involving events, nodes, and relationships.

An "event" is defined as a change in link color. There are two types of events: (1) befriending: when a red link changes to blue, meaning that the two end-nodes have changed from enemies to friends; (2) betrayal: when a blue link changes to red, meaning that the two end-nodes have changed from friends to enemies. One event occurs each time step until the network reaches global stability. Each event is logged in the event history and is listed as "At $t = T$, node X befriended / betrayed node Y." As the model runs, it produces a simple proto-narrative, represented by the list of events that has occurred up to the current time step. This proto-narrative is akin to the "proto-histories" that Robert Axtell generates with the Sugarscape model.

The network topology constitutes a rudimentary setting representing social space rather than physical space. Like the setting of a novel, the geometric configuration of nodes and links defines the environment in which character interaction-events will unfold. Different settings engender different event sequences. The degree distribution, representing how equal vs. unequal the initial allocation of social ties is, is a key topological feature. The model has three settings: (1) uniform degree distribution, (2) random, and (3) maximum degree inequality. Maximum inequality networks are more centralized and have higher clustering coefficients given the same number of links: because the structural balancing problem is more complex, the event history is longer on average for these networks.

Nodes may be thought of as rudimentary characters. Like characters, they have basic descriptive attributes including their degree, initial number of friends and enemies, and location in the network's topology. Nodes undergo a simple version of "character development". As the model runs, each node accumulates a personal history consisting of the link-switching events in which it has been involved. Levels of development vary. Some character-nodes are important to the proto-narrative and are involved in many events, while others are marginal. As the model runs, it draws two distributions, both shown in figure 3: (1) the

Figure 6 Example run showing percentage of stable triads over time.

degree distribution of the nodes at set-up, and (2) the time-evolving event distribution, showing how skewed the history of the model has been towards particular node-characters. These distributions are imperfectly correlated. High degree nodes, called "hubs", generally figure in more events, however, it is theoretically possible to have hubs that are embedded in only stable triads: such nodes are central to the network, but peripheral to its narrative of development.

Link-switchings constitute rudimentary character interactions and are designated as either "befriending" or "betrayal". Just as each node in the structural balance network has a history, so does each relationship-link. Some relationships are active and tumultuous, with many oscillations between friendship and enmity, while others are uneventful. The history of each relationship is stored in an adjacency matrix where the (i, j) entry tracks the number of events between node-characters i and j. It can be visualized as a separate character-interaction network. Figure 5 shows an example. It is important to distinguish between the social balance network, which generates character interactions, and the character interaction network, which summarizes them after the fact. Consistent with White's concept of "netdom switchings", as links in the structural balance network switch between friendship and enmity domains, they build up story-ties in the character interaction network. The more domain switchings in the SBT model, the stronger the tie formed. The character interaction network is comparable to the novelistic networks shown in the first half of this paper, which also summarize character interactions. The SBT and character interaction networks may have significant topological differences: in the example above, the character interaction network is more centralized and has a lower clustering coefficient and graph density than the structural balance network that generated it.

■ **Figure 7** (a) Example run showing percentage of friendship (blue) vs. enmity (red) ties over time. (b) Initial and final SBT network.

Lastly, the model's progression from instability to stability provides both a rudimentary narrative arc and sense of closure. To paraphrase Aristotle's *Poetics*, a narrative begins when an initially stable situation is disturbed by an inciting incident. The resulting disequilibrium constitutes the central problem of the narrative, towards which all actions aim according to Aristotle's principle of unity. The narrative ends when the problem is resolved, providing a sense of closure through the establishment of a new equilibrium.

The SBT simulation conforms to the Aristotelian structure, telling the simple but meaningful story of how a community journeys from an initially unstable configuration to a stable configuration. It possesses a clear beginning, middle, and end and achieves narrative closure through the establishment of a structurally balanced equilibrium. Each event either advances the proto-narrative towards resolution by increasing the percentage of stable triads or constitutes a reversal by decreasing the percentage of stable triads.

The SBT dynamics adhere most closely to an Aristotelian paradigm when we begin the simulation from a state only slightly perturbed from equilibrium by, for example, introducing a single friendship link between two enemy factions or a single enmity link into a network of universal friendship. These initial conditions have obvious analogues in classical narrative. The former roughly corresponds to the inciting incident in a drama such as *Romeo & Juliet*, in which members of rival factions fall in love, while the latter roughly corresponds to the inciting incident in a tragedy such as *Julius Caesar*, in which an act of betrayal between friends tears a community apart. Once disrupted, there are two ways to restore stability: (1) the disturbing link can be extinguished or forced back into conformity as in *Romeo & Juliet*; or (2) the old social structure can be completely unraveled and a new equilibrium established, as in *Julius Caesar*. The latter case is pictured in figure 7. At $t = 0$, the network consists
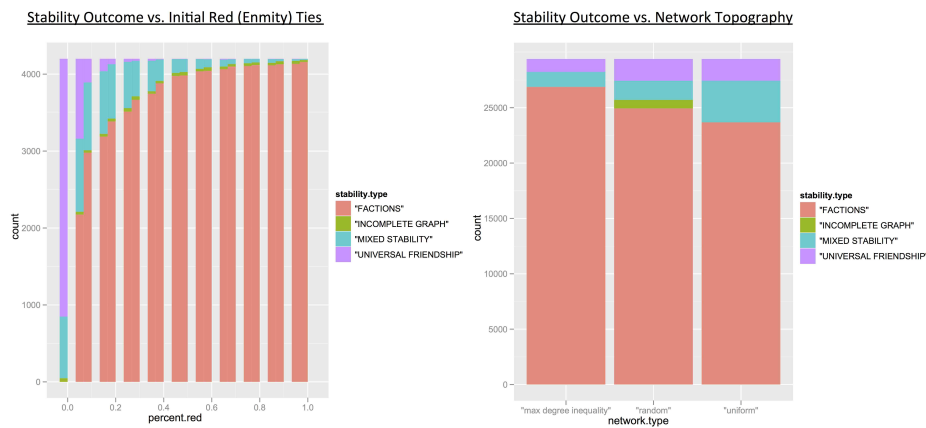
Figure 8 Story ending, represented by stability outcome, is affected by network topology and percentage of friendship vs. enmity ties.

of only friendship ties except for one enmity link. Universal harmony unravels in the first phase of the run—what we might call "act one"—evidenced by the rapid fall in friendship links and rise in enmity links. This phase ends when the percentage of friendship and enmity ties is equalized, indicating the original social structure has completely collapsed. A middle phase of instability and flux follows. Finally, a new equilibrium structure begins to form consisting of two polarized factions. The model terminates when social mitosis is complete. The sequence of events is reminiscent of Harrison White's assertion that "a network can be traced as similar stories appear across a spread of dyads. The "story" that spreads in this case is one of mounting enmity—an initial act of betrayal ripples outward, cleaving the network.

While the model will eventually find an ending represented by a stable outcome, when and what type of ending are indeterminate. There are three possibilities for an incomplete graph: (i) universal harmony, or what we might consider a "happy ending"; (2) polarized factions, or what we might consider an "unhappy ending"; (3) or a mixed outcome. The indeterminacy of the ending provides a rudimentary version of narrative suspense. Figure 8 shows the effect of different input parameters on the ending of the proto-narrative.

4   Conclusion

Lest the "proto-narrative" produced by this structural balance simulation strike us as overly simplistic, it is worth observing that the acts of betraying and befriending and the reconfiguration of social allegiances are the core events of many classical and contemporary social dramas, ranging from French court novels such as *Le Princesse de Clèves* and *Les Liaisons dangereuses* to contemporary soap operas such as *Gossip Girl*. In a recent paper entitled "Facebook for Vikings", folklorist Timothy Tangherlini argues that the plot structures of Scandinavian story cycles can be understood in terms of shifting alliances and enmities consistent with SBT:

> In a great deal of saga scholarship there is an understandable emphasis on understanding enmity, with friendship acting as a powerful counter force (Byock 1982; Miller 1983 and 1990). Network analysis allows one to consider friendly interactions and antagonistic relationships both as individual features of the saga narrative and in concert with each other... Perhaps one of the most complicated aspects of social

interaction considered in the sagas is the selection of friends and its inverse, the selection of enemies. [23]

Tangherlini analyzes the famous "Höfuðlausn" or head ransom episode in *Egil's Saga* in terms of interdependent triads between four characters—Egil, Eirik, Arinbjorn, and Gunnhild. All friendship and enmity ties are determined prior to the episode except for the relationship between Arinbjorn and Gunnhild. Tangherlini argues that the dramatic arc of the head ransom episode consists in the determination of the Arinbjorn-Gunnhild relationship and the reconfiguring of the Egil-Eirik relationship based on the stability requirements of SBT.[7] Tangherlini's approach is confirmed by a recent study of 'mythological networks,' in which MacCarron and Kenna find that several ancient and medieval epics, including *Beowulf* and *The Iliad*, obey structural balance rules [15].

Studies such as Tangherlini's demonstrate that structural balance theory has value as a descriptive model of socially complex narratives, but it has perhaps even greater potential as a generative mechanism that could be used for new story creation.

Finding realistic but tractable story generation mechanisms is an ongoing challenge for artificial intelligence researchers working on narrative. Most story generators rely on either (a) corpora of pre-existing stories (e.g., MEXICA [21]), or (b) story grammars. Researchers in this area are now attempting to expand the suite of generative mechanisms to include games as well as crowdsourcing [14]. In a recent paper, Pablo Gervás argues for the value of chess as a narrative generation mechanism:

> Chess provides a finite set of characters (pieces), a schematical representation of space (the board), and time (progressive turns), and a very restricted set of possible actions. Yet it also allows very elementary interpretations of game situations in terms of human concepts such as danger, threat, conflict, death, survival, victory or defeat, which can be seen as interesting building blocks for story construction. [9]

There is a related body of work that makes use of sports game statistics to generate simple narratives akin to newspaper articles [1]. Like chess and sports, dynamic network models such as structural balance could provide a narrative generation mechanism to complement corpus analysis. They possess the added advantage of supporting social behaviors that are more complex than direct competition or physical conflict.

Most promisingly, there is growing emphasis in the interactive narrative community on story-worlds that incorporate complex and nuanced social dynamics. Notable examples include *The Sims* and Michael Mateas' experimental game *Façade* [16]. UC Santa Cruz's *Prom Week*—in which players manipulate the social interactions between high schoolers in the week leading up to prom—is perhaps the most sophisticated example of an emerging

---

[7] Tangherlini's study also provides a useful set of benchmarks for calibrating the SBT model. In *Egil's Saga*, Tangherlini counts a total of 316 friendship interactions and 158 enmity interactions, of which 61 are "lethal" and 97 are "non-lethal" (a roughly 2:1 ratio of friendship to enmity events). This total of 474 interactions is spread across 200 "saga actants." The ratio of events to characters that Tangherlini finds is low compared to that generated by the SBT model I have outlined. This is likely due to several factors. First, the graph for *Egil's Saga* is highly incomplete, with very low graph density. As shown in figure 4, the length of the "event history" is very sensitive to the number of links in the network: relatively complete graphs take exponentially longer to resolve than relatively incomplete graphs. Second, and perhaps more importantly, there is a key difference that should be drawn between story-events and discourse-events. Our model generates story-events, not all of which need ultimately be rendered narratable in the final text. The events in the discourse of *Egil's Saga* have been selected from underlying story-events that are deemed worthy of narration because, for example, they involve specific characters of interest, such as Egil.

category of "social games" that incorporate sociological and psychological models into AI to 'gamify' social experience. *Prom Week* makes explicit use of social networks and claims to have more than five thousand social rules guiding character behavior, including "if someone is mean to me, then someone else does something mean to them, I'm more likely to want to date that person". [17] This enemy-of-my-enemy-is-my-friend rule could be integrated easily with structural balance dynamics. As "social games" gain traction in interactive narrative, generative network models such as the one I have outlined on structural balance will have increasing value.

References

1   Nicholas D. Allen, John R. Templon, Patrick Summerhays McNally, Larry Birnbaum, and Kristian Hammond. StatsMonkey: A data-driven sports narrative writer. In Mark Alan Finlayson, editor, *Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium*, number FS-10-04 in AAAI Technical Reports, pages 2–3. AAAI Press, 2010.

2   T. Antal, P. L. Krapivsky, and S. Redner. Dynamics of social balance of networks. *Physical Review E*, 72:036121, 2005.

3   Robert Axtell and Joshua Epstein. *Growing Artificial Societies*. The Brookings Institution, 1996.

4   Ruth Aylett. Narrative in virtual environments – towards emergent narrative. In M. Mateas and P. Sengers, editors, *Narrative Intelligence: Papers from the AAAI Fall Symposium*, number FS-99-01 in AAAI Technical Reports, pages 83–86, 1999.

5   Jesse L. Byock. *Feud in the Icelandic saga*. University of California Press, 1982.

6   Dorwin Cartwright and Frank Harary. Structural balance: A generalization of Heider's theory. *Psychological Review*, 63(5), 1956.

7   D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, 2010.

8   David K. Elson, Nichoas Dames, and Kathleen McKeown. Extracting social networks from literary fiction. In Jan Hajič, Sandra Carberry, Stephen Clark, and Joakim Nivre, editors, *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 138–147, Stroudsburg, PA, 2010. Association for Computational Linguistics.

9   Pablo Gervás. From fleece of fact to narrative yarns. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 125–134, İstanbul, 2012.

10   René Girard. *Deceit, Desire, and the Novel: Self and Other in Literary Structure*. Johns Hopkins Press, Baltimore, 1961.

11   Fritz Heider. Attitudes and cognitive organization. *Journal of Psychology*, 21:107–112, 1946.

12   Norman P. Hummon and Patrick Doreian. Some dynamics of social balance processes. *Social Networks*, 25:17–49, 2003.

13   K. Kulakowski, P. Gawronski, and P. Gronek. The Heider balance – a continuous approach. *International Journal of Modern Physics C*, 16:707–716, 2005.

14   Boyang Li, Stephen Lee-Urban, Darren Scott Appling, and Mark O. Riedl. Automatically learning to tell stories about social situations from the crowd. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 142–151, İstanbul, 2012.

15   Pádraig MacCarron and Ralph Kenna. Universal properties of mythological networks. *EuroPhysics Letters*, 99:28002, 2012.

16   Michael Mateas and Andrew Stern. Façade: An experiment in building a fully-realized interactive drama. Presented at the Game Developers Conference, Game Design track, San Jose, California, 2003.

17 Josh McCoy, Mike Treanor, Ben Samuel, Brandon Tearse, Michael Mateas, and Noah Wardrip-Fruin. Authoring game-based interactive narrative using social games and *Comme il Faut*. Presented at the 4th International Conference & Festival of the Electronic Literature Organization, 2010.

18 William Ian Miller. Choosing the avenger: some aspects of the bloodfeud in medieval iceland and england. *Law and History Review*, 1:159–204, 1983.

19 William Ian Miller. *Bloodtaking and peacemaking feud, law, and society in Saga Iceland*. University of Chicago Press, 1990.

20 Franco Moretti. Network theory, plot analysis. *New Left Review*, 68:80–102, 2011.

21 Rafael Pérez y Pérez and Mike Sharples. MEXICA. *Journal of Experimental and Theoretical Artificial Intelligence*, 13:121–139, 2006.

22 Graham Sack. Quantifying imaginary social forms: Character networks in the 19th century British novel. Master's thesis, Columbia University, 2011.

23 Timothy Tangherlini. Facebook for Vikings: Social network analysis and Egils Saga. Presented at the Society for the Advancement of Scandinavian Study Conference, 2011.

24 Zhigang Wang and Warren Thorngate. Sentiment and social mitosis: Implications of Heider's Balance Theory. *Journal of Artificial Societies and Social Simulation*, 6(3), 2003.

25 Harrison White. *Identity and Control: How Social Formations Emerge*. Princeton University Press, New Jersey, 2008.

A Data-Driven Approach for Classification of Subjectivity in Personal Narratives*

Kenji Sagae[1], Andrew S. Gordon[1], Morteza Dehghani[1],
Mike Metke[1], Jackie S. Kim[1], Sarah I. Gimbel[2], Christine Tipper[2],
Jonas Kaplan[2], and Mary Helen Immordino-Yang[2]

1   Institute for Creative Technologies
     University of Southern California
     Los Angeles, CA, USA
     {sagae,gordon,morteza,mmetke,skim}@ict.usc.edu
2   Brain and Creativity Institute
     University of Southern California
     Los Angeles, CA, USA
     {sgimbel,tipper,jtkaplan,immordin}@usc.edu

Abstract

Personal narratives typically involve a narrator who participates in a sequence of events in the past. The narrator is therefore present at two narrative levels: (1) the extradiegetic level, where the act of narration takes place, with the narrator addressing an audience directly; and (2) the diegetic level, where the events in the story take place, with the narrator as a participant (usually the protagonist). Although story understanding is commonly associated with semantics of the diegetic level (i.e., understanding the events that take place within the story), personal narratives may also contain important information at the extradiegetic level that frames the narrated events and is crucial for capturing the narrator's intent. We present a data-driven modeling approach that learns to identify subjective passages that express mental and emotional states of the narrator, placing them at either the diegetic or extradiegetic level. We describe an experiment where we used narratives from personal weblog posts to measure the effectiveness of our approach across various topics in this narrative genre.

1   Introduction

Beyond communicating a simple description of a sequence of connected events, personal narratives are often crafted to evoke emotions or sway opinions by delivering a story through the point-of-view of one of its participants. Fully understanding a personal narrative therefore requires more than an accurate representation of the semantics of the story; the storyteller's intent is often expressed through subjective statements that may be used to frame specific events in ways that influence their interpretation by the audience. Similarly to how a soundtrack can set a specific mood in film to heighten the emotional impact of the sights

---

and sounds of a story, skilled rhetoric can serve to enhance the impact of events depicted in writing.

Subjective language, which expresses opinions, emotions, thoughts, preferences, and other mental states of the narrator, is crucial for delivery of the intended interpretation of a personal story. Despite significant efforts in research on identification of subjective language and detection of sentiment polarity for a handful of language genres, existing approaches fall short of the requirements for modeling subjectivity in personal narratives. Homodiegetic narrative, where the narrator is also a character in the story, presents an interesting challenge: subjective language may refer to mental or emotional states of the narrator as the storyteller, or of the narrator as a participant in the story. One way to characterize this distinction is to place specific instances of subjective language as referring to one of two *narrative levels*: the extradiegetic level, where the narration takes place, or the diegetic level, where the events of the story take place. In addition to its importance in interpreting narrative discourse, where it is important to distinguish emotional states occurring within the story from those that apply to the storyteller during the act of narration, automatic classification of diegetic level and subjectivity of narrative segments can be beneficial in a variety of practical applications involving narrative data. For example, when searching a large corpus of narratives by specific activities, such as visits to the zoo or protest rallies, it may be desirable to focus on text at the diegetic level by appropriately weighing query terms. Isolating events at the diegetic level would also be desirable in automatic induction of schemas and acquisition of commonsense knowledge from narratives. On the other hand, when the information need targets the intellectual or emotional impact of an experience, without specific constraints on the activity described, focus should be on subjective statements and on passages referring to the extradiegetic level.

We present a data-driven modeling approach for identification of subjective language in each of these two narrative levels, showing how text classification techniques, combined with human annotation, can learn to classify subjectivity in personal narratives. Much of motivation in our work is shared by the line of research on computational approaches to subjectivity in narrative due to Wiebe [38], while our view of subjectivity is that defined by later work by Wiebe and colleagues (e.g., [37]). Unlike Wiebe's original investigation of subjectivity in third person narratives, we deal here with first person narratives, a genre that we describe in more detail in section 2.1. Additionally, while Wiebe's analysis focused on characterization of subjectivity in terms of linguistic elements, our approach focuses instead on the application of machine learning and text classification techniques to the task of identification of subjectivity, following recent research that we discuss in section 2.2. We conclude section 2 with a brief discussion about subjectivity and narrative levels. In section 3 we describe the narrative data used in our investigation, taken from personal Internet weblogs and selected from topics where we expect to find examples of subjective language referring to both the diegetic and extradiegetic levels. In sections 4 we present our computational modeling approach and experiments, in two parts. The first part of our approach involved the design of an annotation scheme for subjectivity and diegetic levels for first person narratives, and manual annotation of 40 narratives (section 4.1). The second part involved learning multiclass classification models from the annotated corpus (section 4.2). We present and discuss our results in section 5, and finally we present our conclusions and briefly discuss future work in section 6.

2   Background

2.1   Personal Narratives

The genre of the personal narrative is broadly defined as the non-fiction stories that people share with each other about their own life experiences. This genre of discourse includes the stories told among family members while reviewing old photographs [4], the accounts shared among coworkers in office environments [6], the testimonials of people in interviews [10], and the reflections of daily experiences of people written to private diaries [35]. In this research our focus is the written forms of personal narrative (text documents), which are more amiable to automated analysis than other forms. Specifically, we develop and evaluate our methods on personal narratives extracted from Internet weblogs.

The phenomenal rise of personal weblogs has afforded new opportunities to collect and study electronic texts of personal narratives on a large scale. While blogging is popularly associated with high-profile celebrities and political commentators, the typical weblog takes the form of a personal journal, read by a small number of friends and family [25]. As with the adoption of other forms of electronic communication, personal narratives in weblogs take on several new characteristics in adapting to a social media environment that is increasingly public and interconnected. Eisenlauer and Hoffman [7] argue that the on-going technological development of weblog software has led to an increase of collaborative narration, moving the form further toward Ochs and Capps [26] conception of the hypernarrative, where discourse is best understood as a conversation among multiple participants. Langellier and Peterson [18] characterize this collaborative narration as a form of public performance, creating a productive paradox between the insincerity needed to craft a good story and the sincerity of the blogger as a character in the narrated events.

This productive paradox seen in weblog storytelling helps distinguish personal narrative from other narrative forms. As in all narrative, personal narrative consists of descriptions of multiple events that are causally related, but requires further that the narrative perspective is the author's own. The expectation of the reader is that the narration reflects the truthful interpretation of events actually experienced by the author, but the truth of the narration is constrained by the demands of good storytelling.

2.2   Analysis of Subjective Language

Because personal narratives feature a storyteller who is also a character in the story, it is common for the events of the story to be framed by the narrator in terms of opinions, emotions, preferences, and other commentary that influences the reader's interpretation of the events. While it is possible for a narrator to be objective in recounting first-hand participation in a story, our analysis is focused on personal narratives that are framed by subjective language employed by the narrator, and more precisely on computational models for identification of subjectivity in personal narratives.

Although there has been remarkable interest in analysis of sentiment and subjectivity in text in the past decade, the bulk of the research has been focused on a few language genres, with the most prominent example being reviews of movies, products, restaurants, etc. (e.g., [27, 3, 33]). Reviews are attractive as the target of sentiment analysis, as they are abundantly available online, they restrict language processing tasks to well-defined domains, and they necessarily express opinions that can often be binned into negative or positive categories relatively easily. In analysis of reviews, it is common to frame the task as sentiment polarity classification ("thumbs up" vs "thumbs down"), often aided by a preprocessing step that

identifies *subjective* language, which Pang and Lee [28] define simply as opinion-oriented language. Another language genre where subjectivity and sentiment analysis has been studied extensively is news, where the identification of subjectivity is itself the target of analysis, rather than binary classification of sentiment polarity. In their work on subjectivity analysis, Wiebe et al. [37] take a broader view of subjective language, which they define as the expression of *private states* [29], which includes emotions, opinions, evaluations and speculations. A third major area of application of sentiment and subjectivity analysis, which has been growing rapidly, is user-generated content, including Twitter, discussion boards, political weblogs, and YouTube video reviews (e.g., [1, 23, 24]).

Although far from exhaustive, the list of language genres mentioned above serves to illustrate how the goals of subjectivity analysis can vary widely when different types of content are considered. For example, in reviews it is more important to determine whether statements are positive or negative, while in news there is a greater focus on separating opinion from fact. Even though goals and even definitions may vary, the most common types of application are related to fulfilling information needs or estimating public interest and opinion regarding specific issues, products, etc. In the case of narrative, however, analysis of subjectivity and sentiment can play a different type of role. Correctly assessing the mental and emotional state of the narrator is crucial for understanding the intent of a narrative beyond the facts and events of the story; narratives are often crafted with the explicit goal to have an emotional impact on the reader, sometimes more so than they are to convey a specific sequence of events. In contrast to the main role of subjectivity in reviews or editorial pieces, subjective language in narrative goes far beyond opinions. The expression of emotions, thoughts, preferences and other mental and emotional states is of primary importance.

In our work, we adopt Wiebe et al.'s notion of *subjective language* as the linguistic expression of *private states* (including opinions, evaluations, emotions, speculations and other mental processes), which are experienced but are not open to external observation or verification by others. Our main focus is on private states of the narrator, since we are dealing with personal narratives, which express the narrator's point of view. While it can be tempting to define subjective language as the statement of opinions, in contrast to objective statement of facts, this would be an imprecise definition. For example, while the text segment *I know her name* may be considered a statement of fact by the narrator, it is a case of subjective language. The key issue here is not whether a statement is true or made with certainty or privileged knowledge, or even whether it can be considered a fact, and rather whether it expresses a private state and not something that can be observed or measured objectively and externally. For example, while *I felt sick* is a subjective statement, since it cannot be observed externally, the statement *I had a 102-degree fever* is objective. Similarly, *it was hot yesterday* is subjective (the narrator's opinion), while *it was 95 degrees yesterday* is objective. It is not important at this point to distinguish whether a statement such as *he was sad* is subjective because it expresses a private state of a third person, or because it expresses the narrator's opinion or evaluation of a third person, since in either case the statement is subjective. On the other hand, *he said he was sad* is an objective statement, since it describes an event that can be observed externally (namely, the act of saying).

2.3 Private States and Narrative Levels

The genre of the personal narrative is particularly interesting from the point-of-view of analysis of subjectivity in that the narrator experiences emotions and holds opinions both within the story, as a character along with other story participants, and also outside of story. Accordingly, the narrator may employ subjective language that applies to at least

two different narrative levels. Consider, for example, a narrative that recounts events that include the narrator being afraid of a puppy and disliking dogs as a child, but also expresses the now adult narrator's current embarrassment of this long abandoned fear and current fondness for dogs. The universe of the story, where the narrator is a child, is sometimes referred to as the diegetic (or intradiegetic) level, and the act of narration is performed at the extradiegetic level, where in this case the narrator is an adult addressing the reader. This example includes expression of several private states experienced by the narrator: as a character in the story (i.e., at the diegetic level), the narrator experiences fear at a specific moment, and holds a negative preference for dogs; in contrast, the narrator expresses the private states of embarrassment and positive preference for dogs at the time of storytelling (i.e., at the extradiegetic level).

Although in our discussion we adopt the terms proposed by Genette [9] to speak of diegetic levels in narrative, we do so only to determine whether private states are either internal or external to the universe of the story, leaving aside the more complex issues of matching private states to more levels in embedded narratives. In other words, instead of performing a complete analysis of diegetic levels, we make only a binary distinction between the extradiegetic level and all other (intradiegetic) levels, with no distinction made in the levels of embedded narratives. An alternative way to characterize what we refer to as private states at the diegetic and extradiegetic levels is to use the notion of time points due to Reichenbach [30]: the narrator might refer to private states at speech time (at the time of narration), or at the event or reference time. However, our main concern is not necessarily one of time; the distinction we make in the present work is between private states experienced by the narrator as a character in the story, and private states experienced by the narrator as the storyteller. This distinction reflects the narrator's exclusive advantage in framing the story to influence the audience's interpretation and reaction. The impact of diegetic and extradiegetic material can be understood intuitively by considering the soundtrack in a movie. When watching a movie, we observe events taking place and a story unfolding, which may evoke emotion. External to the universe where the story takes place, however, we may also hear music (e.g., romantic music for a romantic scene, or fast-paced music for a car chase), which sets a specific mood and serves to evoke or amplify emotional reactions. This music is at the extradiegetic level: it is audible to the audience only, and does not exist for the characters in the story.

3   Data

In developing and evaluating a data-driven approach to our classification task, we required a corpus of personal stories containing substantial amounts of subjective statements describing private states belonging to either the diegetic or extradiegetic level, meaning that the narrator experiences the private state either within the universe of the story, or outside, at the time of the act of narration, respectively. Although weblog content is abundant and readily available, selecting and annotating random weblog posts would be inefficient. Gordon and Swanson [13] estimated that only 4.8% of randomly sampled non-spam English-language weblog posts can be characterized as personal stories, defined by them as non-fiction narrative discourse describing a specific series of events in the past, spanning minutes, hours, or days, where the storyteller or close associate is a participant. Even within this small subset of posts, our expectation is that the balance between description of private states held at the diegetic and extradiegetic levels will vary widely. For example, the play-by-play narrative of a baseball game might focus entirely on the diegetic universe, with descriptions of excitement, happiness

or apprehension applying only to the diegetic level, while an account of cherished childhood memories might move the narrator to describe an emotional state triggered by the events in the diegetic universe but experienced during narration, at the extradiegetic level.

To curate a well-balanced corpus for analysis, we focused our efforts on finding weblog posts about situations and activities that would lend themselves to a mix of these two types of expression of subjectivity, or private states. We specifically targeted narratives of socially-questionable behavior, e.g., stories of stealing, quitting a team, giving a child up for adoption, or getting into a physical fight. We expected that bloggers who shared personal narratives about socially-questionable behavior would feel the need to be descriptive of the events that occurred, including opinions, thoughts and emotions held at the time, and to provide some rationale or justification for their behavior, leading to expression of their current feelings about the past events of the narrative. Collectively, we brainstormed a list of such situations that could potentially be found in public weblogs (Figure 1).

To conduct these situation-specific searches, we used the technologies and methodologies described by Gordon et al. [14], which were used by them to find hundreds of personal narratives in weblogs related to health emergencies. The approach begins with the automatic filtering of personal narratives from streams of weblog posts, applying supervised story classifier to three years of non-spam English-language weblog posts provided by a weblog aggregator (Spinn3r.com). The filtered story collection (over 20 million posts) was then indexed using a text retrieval engine (Apache Lucene), which could be queried with a large array of weighted terms. Initial queries were authored for each of the socially-questionable behavior following Gordon and Swanson [12], where paragraph-sized fictional prototypes were used to retrieve similar instances. Retrieved posts were then annotated as to their relevance to the query, and this feedback was used to further refine the query and weight query terms using the Rocchio relevance feedback algorithm [32]. We identified 460 posts containing narratives of socially-questionable behavior using this approach, from which we selected 40 posts that we judged to be most compelling as personal narratives. These 40 posts include 22 of the 26 topics in Figure 1, with no topic appearing in more than three stories. Topics appearing in multiple posts include lying, divorce, protest rallies, breaking the law, and quitting a team, abortion, disobeying a superior, murdering someone, getting into a physical fight, killing an animal, prostitution and physically punishing a child. The four topics that are not represented in our selection of posts are: stealing, taking an unfair advantage, putting your own interests above others, and neglecting to care for children.

4 Approach

Following previous data-driven efforts on subjectivity and sentiment analysis, exemplified by the work of Wiebe et al. [37] and Pang and Lee [27], we use a machine learning approach typically associated with text classification. While Pang and Lee leverage "found-data" to train a classifier for subjectivity detection in reviews without the need for manual annotation, our approach has in common with Wiebe et al.'s that it involves the definition of an annotation scheme, training of human annotators, and manual annotation of training data for subjectivity classification. As an initial attempt to address subjectivity modeling in personal narratives, we use a simple discriminative text classification approach, with a relatively small training dataset consisting of 40 narratives. These narratives were found originally in personal weblogs, and for reasons unrelated to the work described here were edited for length prior to

▬ Participating in a protest rally	▬ Quitting a team
▬ Quitting a job	▬ Cheating on a test
▬ Telling a lie	▬ Killing an animal
▬ Getting into a physical fight	▬ Prostitution
▬ Converting from one religion to another	▬ Taking an unfair advantage
▬ Having an abortion	▬ Physically punishing a child
▬ Stealing something	▬ Violating a religious practice
▬ Crossing a picket line	▬ Breaking the law
▬ Changing one's political party	▬ Murdering someone
▬ Changing the country of your citizenship	▬ Disobeying a superior
▬ Making a large personal sacrifice	▬ Putting a child up for adoption
▬ Cheating in a romantic relationship	▬ Putting your own interest above others
▬ Getting a divorce	▬ Neglecting to care for children

■ **Figure 1** List of topics involving socially questionable behavior for personal narrative selection.

annotation[1]. In most cases, editing consisted largely of removing sentences from the original weblog post, with occasional addition of a few words to restore coherence to the edited text. The edited narratives retained the vocabulary and much of the narrative structure of the original posts, and contained 160 to 185 tokens (words and punctuation) each. The average length for these narratives is 169 tokens.

4.1 Text Annotation

After selection of narratives from personal weblogs, the first step towards creation of the dataset necessary for training a data-driven model and subsequent empirical validation of our overall approach was the definition of an annotation scheme. The annotation scheme described here and used in our experiments is the product of iterative refinement involving a computational linguist and two annotators. The annotators, whose backgrounds are in Linguistics and Psychology, first acquired familiarity with basic concepts in narratology and computational analysis of narrative by reading the background chapter of Indejeet Mani's book *Computational Modeling of Narrative* [19]. They then annotated a practice set of about 30 narratives, individually but in frequent consultation. This process resulted in refinements to the annotation scheme and guidelines for dealing with borderline cases, and was followed by annotation of the 40 narratives in our dataset by each of the two annotators individually.

The annotation task consists of tagging segments in the narrative with one of six labels, described below. Text segments were determined automatically and consist of one or more clauses. The use of clauses as the granularity for annotation was motivated by concerns both principled and practical in nature. Perhaps the easiest segmentation strategy for identification of subjective passages in narrative is to consider each sentence as a target for labeling. Sentences, however, are clearly too coarse grained, since a single sentence may express an unbounded number of objective and subjective statements through coordination.

---

[1] These narratives were used as stimuli in a separate series of experiments that examine the emotional reactions of readers of personal narratives. These experiments required that human subjects read each narrative in 36 seconds, and for this reason each of the 40 narratives was shortened from its original weblog post version. In future work we plan to investigate the relationship between subjectivity in narrative, as annotated in the work describe here, and emotional impact on readers.

A more suitable strategy is to define segments in the spirit of the Elementary Discourse Units (EDUs) in Rhetorical Structure Theory [20], as applied to entire texts by Marcu [22]. Instead of addressing the challenges of adapting full EDU segmentation to the needs of our task, we opted to use a simplified segmentation scheme inspired by EDUs, taking clauses as the target of annotation, with the application of rules and simple heuristics to prevent segmentation of certain types of subordinate clauses that tend not to be relevant to our annotation. For example, the non-finite subordinate clauses in *he told her not to go* and *I like going to the movies* are not split into segments separate from their matrix clauses. Other examples of subordinated language that results in multi-clause segments include *going to the movies is what I like to do on weekends* (one segment with four clauses) and *he said he would return* (one segment with two clauses). Our segmentation approach is based on identification of syntactic patterns in parse trees produced automatically by the Stanford parser [16], and largely follows the EDU segmentation approach described by Tofiloski et al. [36], but without the full set of rules and lexical patterns necessary for complete EDU segmentation according to the RST guidelines. Because our segments are sometimes too fine-grained, and because narratives sometimes include passages that do not fall within one of the categories defined by our scheme, the annotation scheme provides the option of tagging specific clauses as Other/None (see below).

The tags in our annotation scheme are:

Story Event Denotes a clause that corresponds to an event in the story.

Story Private State Denotes a clause that corresponds to an expression of a private state of the narrator that applies within the diegetic level.

Story Other Denotes other material that applies mainly to the story, such as descriptions, direct quotes, etc.

Subjective Statement Denotes an expression of a private state at the time of narration (at the extradiegetic level), rather than within the story.

Objective Statement Denotes an expression of fact that applies at the time of narration (at the extradiegetic level), rather than within the story.

Other/None Used for tagging clauses that do not fall within one of the categories above.

The annotation scheme is intended to distinguish discourse segments along two dimensions: (1) subjective vs. objective language; and (2) language that refers to the diegetic vs. the extradiegetic level. Although the notions of *emotion* and *sentiment* are certainly important aspects of narratives that are relevant to our overall goals, we focused our efforts on the related notion of *subjectivity* as the expression of private states. This simplifies labeling of cases such as reported speech and reported emotions. For example, in *he said he was sad*, we do not treat *he was sad* as an independent segment, since it is subordinated language. The single segment is labeled as a Story Event, reflecting the *saying* event, even though it involves reporting of a private state. However, in *I knew he was sad*, there is a single segment and it is labeled as a Story Private State, not because of the emotion reported, but because *knowing* is a private state.

The manual annotation process was done with the aid of the Story Workbench [8], a flexible environment for linguistic annotation, customized specifically for our annotation task. Figure 2 shows a narrative being annotated using the Story Workbench. The top middle section shows the text of the narrative, and the top right shows the segments to be annotated. Segmentation is performed automatically using clause boundaries determined by the Stanford parser [16] integrated in the Story Workbench. From the point-of-view of the annotators, this segmentation and population of the top right area with segments happen automatically
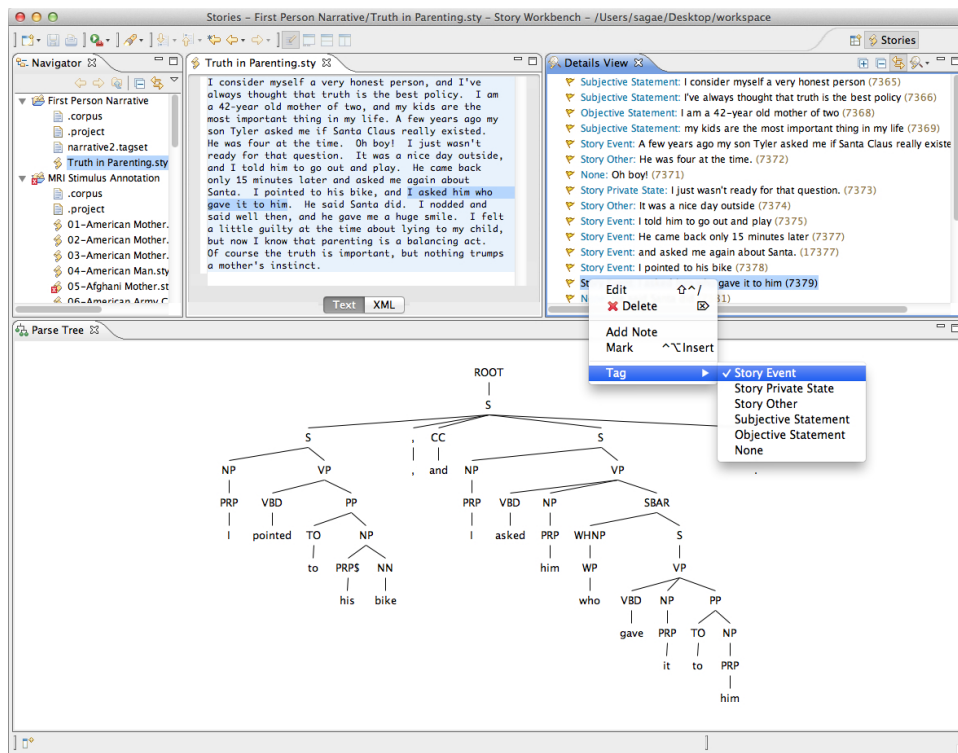
Figure 2 The Story Workbench tool for linguistic annotation [8] customized for annotation of private states in their narrative levels. The top three sections of the interface show a list of narratives to annotate, the text of a narrative, and a list of segments to be labeled. The bottom area shows an automatic syntactic analysis of the sentence being annotated.

and nearly instantly once text is loaded or entered in the text area. Annotators simply go down the list of segments, choosing one of the six labels for each of the segments.

Raw pairwise agreement over 571 segments (40 stories, with 12 to 18 segments per story) on the six-way labeling task was 84%. We measured chance-corrected agreement using Krippendorf's α and obtained a value of 73%. To produce the final annotations, cases where the annotators disagreed were discussed and a final label was chosen. The most frequent label is the first in the list above, Story Event, which accounts for 34% of the segments. Story Private State, the second most frequent label, accounts for 29% of the segments. Table 1 shows several characteristics of the final annotated narrative corpus used in our experiment. Appendix A shows an example of how a narrative is annotated according to our annotation scheme[2].

4.2   Classification of Subjective Language and Narrative Level

We now turn to automatic classification of narrative segments according to our annotation scheme. As in the manual annotation process, segmentation is performed as a pre-processing step using the clause boundaries produced by the Stanford parser. The next step is then

---

[2] Because of issues regarding the expectation of privacy from bloggers [15] and the nature of the material in our narratives, we do not use examples taken from our corpus.

Table 1 Characteristics of the annotated corpus of narratives.

Corpus attribute	Value	
Number of narratives	40	
Average number of tokens per narrative	169	
Average number of segments per narrative	14	
Frequency of Story Event label	194	(34%)
Frequency of Story Private State label	166	(29%)
Frequency of Story Other label	57	(10%)
Frequency of Subjective Statement label	85	(15%)
Frequency of Objective Statement label	63	(11%)
Frequency of None label	5	(1%)

to tag each resulting segment with one of the six categories in the annotation scheme. We first approached this step as a straightforward text classification task at the segment level, treating each segment as independent. We use multiclass classification with Maximum Entropy models [2][3], and for each segment we extract features of the following types:

- Bag of words (unigram features, or w_i);
- Part-of-speech tags for each word in the segment, as assigned by the Stanford parser (t_i);
- Word bigrams (w_i, w_{i+1});
- Part-of-speech tag bigrams (t_i, t_{i+1});
- Part-of-speech tag trigrams (t_i, t_{i+1}, t_{i+2});
- Word/part-of-speech tag bigrams (w_i, t_{i+1} ; t_i, w_{i+1});
- Word/part-of-speech tag trigrams (w_i, t_{i+1}, t_{i+2} ; t_i, w_{i+1}, t_{i+2} ; t_i, t_{i+1}, w_{i+2});

These features are intended to capture the bag-of-words representation widely used in text classification, augmented with n-grams to provide some context, and backing off to part-of-speech tags to reduce the sparsity of n-grams. Since our annotation scheme includes two separate dimensions of the narrative segments, the classification task could be performed in two steps (subjectivity detection, and narrative level classification), or a single step of six-way classification. In both cases the same set of feature types is used. Missing entirely from our classification approach is any notion that each segment is in fact not independent from its context. A possible extension to our current model is to add dynamic features for neighboring segments, making the overall model a conditional random field [17] that optimizes the entire set of segment labels for the entire narrative jointly. This is left as future work.

5 Results and Discussion

Because single-step six-way classification and the two-step classification approaches discussed in the previous section produced very similar results, we focus our discussion on the simpler approach of single-step classification. To perform an empirical evaluation of this approach, we performed a "leave one narrative out" cross-validation using our annotated set of 40

---

[3] Our classifier was implemented with Yoshimasa Tsuruoka's Maximum Entropy library available at `http://www.logos.t.u-tokyo.ac.jp/~tsuruoka/maxent/`

■ **Table 2** Precision and recall values for each of the categories in the annotation scheme, obtained through cross-validation of our 40-narrative dataset containing 571 segments.

Label	Precision	Recall
Story Event	0.56	0.87
Story Private state	0.69	0.57
Story Other	0.44	0.27
Subjective Statement	0.49	0.42
Objective Statement	0.58	0.24
Other	1.0	0.2

narratives: to label each of the 40 narratives, we used the remaining 39 to train a classification model. The overall accuracy using the six labels was 58%, a substantial improvement over a simple majority baseline (34%). While far below the level of human annotation in this task, our results are encouraging given our simple text classification approach. A more informative evaluation is to consider the precision and recall of each category individually. Precision for a category c is defined as the number of correct assignments of the c label divided by the total number of times the classifier assigned the c label to a segment. Recall for a category c is defined as the number of correct assignments of the c label, divided by the total number of instances of the c label in our answer key, or manual annotation. Intuitively, precision corresponds to how often the classifier is correct when it assigns a certain label, and recall corresponds to what portion of the items with a certain label the classifier can find. By labeling *every* segment as c, we would obtain perfect recall, but poor precision. Conversely, by assigning the c label very conservatively and only in cases of very high confidence, we could obtain high precision, at the cost of low recall. Table 2 show the precision and recall values for each of the categories in our annotation scheme.

The imbalance of high recall and lower precision for Story Event reflects that our classifier tends to prefer the assignment of the Story Event label over other labels. In particular, a substantial number of segments that should be labeled Story Private State or Objective Statement are labeled by the classifier as Story Event. In one case, the error appears to be in the more general dimension of subjectivity, and in the other, the error is related to distinguishing between narrative levels. This is also reflected in higher precision than recall in identifying the Story Private State and Objective Statement categories. The confusion between Story Event and Story Private State reflects that, even though the model often correctly identifies that the segment is referring to the diegetic level (which is likely due to part-of-speech features that reflect verb tense), it is less accurate in distinguishing between events and private states. In those cases, the error is in subjectivity classification. The confusion between Story Event and Objective Statement, conversely, shows that the classifier sometimes distinguishes subjective segments correctly, but fails to assign them the correct diegetic level. Not surprisingly the Story Private State category is also often confused with the Subjective Statement category (segments corresponding to expressions of private states of the narrator as the storyteller, outside of the story). This highlights the challenge of classifying correctly along both dimensions in our annotation scheme, which is necessary for analysis of subjectivity specifically at the extradiegetic and diegetic levels. Our results for identification of subjective statements that apply to the extradiegetic level are more balanced: we correctly identify almost half of the narrator's expressions of private states, with a relatively low rate of false alarm at about 50%. This is a particularly important category, since it corresponds to the narrator's reflections about the events in the story.

Table 3 Accuracy results for our main classification task using our six-category scheme (Section 4.1), and for two binary classification tasks, each focusing only on subjectivity or diegetic level. Accuracy of a majority baseline classifier is also shown for comparison.

Classification task	Majority baseline accuracy	Accuracy
Six-way classification	34%	58%
Binary subjectivity classification	56%	78%
Binary diegetic level classification	69%	81%

When we consider each dimension separately, we observe substantially higher accuracy, corresponding to easier classification tasks. On the binary task of identifying segments that refer to the diegetic level vs. to the extradiegetic level, which we evaluate simply by grouping the the first three labels of our scheme into one category (diegetic), and the remaining three labels into another category (extradiegetic), we obtain 81% accuracy. For comparison, a majority baseline for this task would assign the diegetic label to all segments and obtain 69% accuracy, since segments that refer to the extradiegetic level are substantially outnumbered by segments referring to the diegetic level. In the binary subjectivity classification task (grouping the Story Private State and Subjective Statement categories into one subjective category), where segments are simply classified as subjective or not, as is common in natural language processing, our approach does well, with 78% accuracy, compared to 56% for a majority baseline. These results, summarized in Table 3, highlight that the combined task of finding subjective language within the appropriate narrative level is predictably more challenging than either subjectivity classification or narrative level classification in isolation.

6 Conclusion and Future Work

We have described a methodology for analysis of subjective language in narrative that involves manual annotation to produce training material that can be used to learn computational models for automatic identification of subjectivity at the diegetic and extradiegetic levels. Although our classification accuracy still needs improvement, it shows promise given the small number of narratives in our training data, and it highlights some of the challenges in this type of classification. Our next step is to annotate a larger set of personal narratives to generate a larger training set and separate development and test sets using unedited text from weblogs. We believe a larger training set will improve the accuracy of our simple classification framework, and that further accuracy improvements may be obtained by going beyond our current framework where each segment is classified independently. In future work, we plan to abandon the assumption that segments are independent, and apply a structured classification approach (e.g., conditional random fields [17]). Additional annotated data will be important for exploring the use of features beyond unigrams and part-of-speech tags (e.g., features extracted from syntactic trees) using development data. In addition, although our current set of 40 narratives similar in length allows us to see how well our classification approach performs across a variety of topics, we plan to confirm that our models generalize to personal narratives from weblogs in their original forms.

With our current text classification model, subjectivity classification accuracy (78%) is at a level where automatic identification of subjective language in personal narratives could be of practical use. For example, Riloff et al. [31] have shown that subjectivity classification at this level of accuracy is useful for improving the precision of information extraction systems. Similarly, our approach to the classification of the aggregated diegetic and extradiegetic

categories performs well enough (81% accuracy) for potential use in a range of other natural language processing technologies. In many cases it would be desirable to filter out passages that refer to the diegetic or extradiegetic level in order to improve performance or precision. For example, information retrieval system that support searches for narratives of specific activities, such as protest rallies or automobile crashes, may garner improvements by indexing only the diegetic material in the document collection. Where relevance feedback is used to refine queries [14], diegetic material could be weighted more heavily when selecting and weighting query terms. Gains should also be expected in language processing systems that aim to generalize over events described in narratives, as in schema induction [5] and commonsense knowledge extraction [21, 11]. Similarly, some systems may benefit by ignoring the events of narratives, particularly where the emotional or intellectual impact of an experience defines the retrieval criteria [34].

References

1   Apoorv Agarwal, Boyi Xie, Ilia Vovsha, Owen Rambow, and Rebecca Passonneau. Sentiment analysis of Twitter data. In *Proceedings of the Workshop on Languages in Social Media*, pages 30–38, Stroudsburg, PA, 2011. Association for Computational Linguistics.

2   Adam L. Berger, Vincent J. Della Pietra, and Stephen A. Della Pietra. A maximum entropy approach to natural language processing. *Computational Linguistics*, 22(1):39–71, 1996.

3   John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, Bollywood, boom boxes and blenders: Domain adaptation for sentiment classification. In John A. Carroll, Antal van den Bosch, and Annie Zaenen, editors, *ACL 2007, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, June 23–30, 2007, Prague, Czech Republic*, pages 440–447, 2007.

4   Richard Chalfen. *Snapshot Version of Life*. Bowling Green State University Popular Press, Bowling Green, OH, 1987.

5   Nathanael Chambers and Dan Jurafsky. Unsupervised learning of narrative schemas and their participants. In Keh-Yih Su, Jian Su, and Janyce Wiebe, editors, *ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2–7 August 2009, Singapore*, pages 297–305, Singapore, 2009.

6   Stephanie J. Coopman and Katherine B. Meidlinger. Interpersonal stories told by a Catholic parish staff. *American Communication Journal*, 1(3), 1988.

7   Volker Eisenlauer. Once upon a blog. . . Storytelling in weblogs. In Christian R. Hoffmann, editor, *Narrative Revisited: Telling a story in the new age of media*, pages 79–108. John Benjamins Publishing Company, 2010.

8   Mark Alan Finlayson. The Story Workbench: An extensible semi-automatic text annotation tool. In *Intelligent Narrative Technologies IV, Papers from the 2011 AIIDE Workshop, Stanford, California, USA, October 10–11, 2011*, pages 21–24, Stanford, CA, 2011. AAAI Press.

9   Gerard Genette. *Narrative Discourse: An Essay in Method*. Cornell University Press, 1980.

10   Andrew S. Gordon. The fictionalization of lessons learned. *IEEE Multimedia*, 12(4):12–14, 2005.

11   Andrew S. Gordon, Cosmin A. Bejan, and Kenji Sagae. Commonsense causal reasoning using millions of personal stories. In Wolfram Burgard and Dan Roth, editors, *Twenty-Fifth Conference on Artificial Intelligence*, San Francisco, CA, 2011. Association for the Advancement of Artificial Intelligence.

12   Andrew S. Gordon and Reid Swanson. Storyupgrade: Finding stories in internet weblogs. In Eytan Adar, Matthew Hurst, Tim Finin, Natalie S. Glance, Nicolas Nicolov, and Belle L. Tseng, editors, *Proceedings of the Second International Conference on Weblogs and Social Media, ICWSM 2008, Seattle, Washington, USA, March 30 – April 2, 2008*, pages 188–189, 2008.

13   Andrew S. Gordon and Reid Swanson. Identifying personal stories in millions of weblog entries. In Ian Soboroff and Akshay Java, editors, *Data Challenge Workshop, Papers from the 2009 ICWSM Workshop*, number WS-09-01 in AAAI Technical Reports, pages 16–23. Association for the Advancement of Artificial Intelligence, 2009.

14   Andrew S. Gordon, Christopher Wienberg, and Sara O. Sood. Different strokes of different folks: Searching for health narratives in weblogs. In *2012 International Conference on Privacy, Security, Risk and Trust, PASSAT 2012, and 2012 International Confernece on Social Computing, SocialCom 2012, Amsterdam, Netherlands, September 3–5, 2012*, pages 490–495. IEEE Computer Society, 2012.

15   Erika Check Hayden. Guidance issued for US internet research. *Nature News*, 496:411, 2013.

16   Dan Klein and Christopher D. Manning. Fast exact inference with a factored model for natural language parsing. In Suzanna Becker, Sebastian Thrun, and Klaus Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 3–10, Cambridge, MA, 2003. MIT Press.

17   John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In Carla E. Brodley and Andrea Pohoreckyj Danyluk, editors, *Proceedings of the Eighteenth International Conference on Machine Learning*, pages 282–289, San Francisco, CA, 2001. Morgan Kaufmann Publishers Inc.

18   Kristin Langellier and Eric E. Peterson. *Storytelling in Daily Life.* Temple University Press, Philadelphia, PA, 2004.

19   Inderjeet Mani. *Computational Modeling of Narrative.* Number 5 in Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2012.

20   William C. Mann and Sandra A. Thompson. Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8:243–281, 1988.

21   Mehdi Manshadi, Reid Swanson, and Andrew S. Gordon. Learning a probabilistic model of event sequences from internet weblog stories. In David Wilson and H. Chad Lane, editors, *Proceedings of the Twenty-First International Florida Artificial Intelligence Research Society Conference, May 15–17, 2008, Coconut Grove, Florida, USA*, pages 159–164, Coconut Grove, FL, 2008.

22   Daniel Marcu. A decision-based approach to rhetorical parsing. In Robert Dale and Kenneth Ward Church, editors, *27th Annual Meeting of the Association for Computational Linguistics, Univeristy of Maryland, College Park, Maryland, USA, 20–26 June 1999*, pages 365–372. Association for Computational Linguistics, 1999.

23   Prem Melville, Wojciech Gryc, and Richard D. Lawrence. Sentiment analysis of blogs by combining lexical knowledge with text classification. In John F. Elder IV, Françoise Fogelman-Soulié, Peter A. Flach, and Mohammed Javeed Zaki, editors, *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,*

Paris, France, June 28 – July 1, 2009, pages 1275–1284, New York, NY, 2009. Association for Computing Machinery.

24   Louis-Philippe Morency, Rada Mihalcea, and Payal Doshi. Towards multimodal sentiment analysis: Harvesting opinions from the web. In Hervé Bourlard, Thomas S. Huang, Enrique Vidal, Daniel Gatica-Perez, Louis-Philippe Morency, and Louis-Philippe Morency, editors, *Proceedings of the 13th International Conference on Multimodal Interfaces*, pages 169–176, New York, NY, 2011. Association for Computing Machinery.

25   Sean A. Munson and Paul Resnick. The prevalence of political discourse in non-political blogs. In Lada A. Adamic, Ricardo A. Baeza-Yates, and Scott Counts, editors, *Proceedings of the Fifth International Conference on Weblogs and Social Media, Barcelona, Catalonia, Spain, July 17–21, 2011*, pages 233–240, 2011.

26   Elinor Ochs and Lisa Capps. *Living Narrative: Creating Lives in Everyday Storytelling.* Harvard University Press, Cambridge, MA, 2001.

27   Bo Pang and Lillian Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In Donia Scott, Walter Daelemans, and Marilyn A. Walker, editors, *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, pages 271–278, Stroudsburg, PA, 2004. Association for Computational Linguistics.

28   Bo Pang and Lillian Lee. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135, 2008.

29   Randolph Quirk, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik. *A Comprehensive Grammar of the English Language.* Longman, New York, 1985.

30   Hans Reichenbach. *Elements of Symbolic Logic.* Macmillan & Co., New York, 1947.

31   Ellen Riloff, Janyce Wiebe, and William Phillips. Exploiting subjectivity classification to improve information extraction. In Manuela M. Veloso and Subbarao Kambhampati, editors, *Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9–13, 2005, Pittsburgh, Pennsylvania, USA*, pages 1106–1111, Pittsburgh, PA, 2005. AAAI Press.

32   J. J. Rocchio. Relevance feedback in information retrieval. In G. Salton, editor, *The SMART Retrieval System: Experiments in Automatic Document Processing*, pages 313–323, Upper Saddle River, NJ, 1971. Prentice Hall.

33   Benjamin Snyder and Regina Barzilay. Multiple aspect ranking using the good grief algorithm in human language technologies. In Candace L. Sidner, Tanja Schultz, Matthew Stone, and ChengXiang Zhai, editors, *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, pages 300–307, 2007.

34   Sara Owsley Sood. Buzz: Mining and presenting interesting stories. *International Journal of Art and Technology*, 1(1), 2008.

35   Rebecca Steinitz. Writing diaries, reading diaries: The mechanics of memory. *Communication Review*, 2:43–58, 1997.

36   Milan Tofiloski, Julian Brooke, and Maite Taboada. A syntactic and lexical-based discourse segmenter. In *ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2–7 August 2009, Singapore, Short Papers*, pages 77–80. Association for Computational Linguistics, 2009.

37   Janyce M. Wiebe, Theresa Wilson, Rebecca Bruce, Matthew Bell, and Melanie Martin. Learning subjective language. *Computational Linguistics*, 30(3):277–308, 2004.

38   Janyce Marbury Wiebe. *Recognizing subjective sentences: a computational investigation of narrative text*. PhD thesis, State University of New York at Buffalo, 1990.

A   Narrative Annotation Example

The following example shows how a personal narrative is annotated following the scheme described in Section 4.1. Segments are enclosed in square brackets subscripted with segment indices. A label for each index is listed after the narrative.

[I consider myself a very honest person,]$_1$ [and I've always thought that truth is the best policy.]$_2$ [I am a 42-year old mother of two,]$_3$ [and my kids are the most important thing in my life.]$_4$ [A few years ago my son Tyler asked me if Santa Claus really existed.]$_5$ [He was four at the time.]$_6$ [Oh boy!]$_7$ [I just wasn't ready for that question.]$_8$ [It was a nice day outside,]$_9$ [and I told him to go out and play.]$_{10}$ [He came back only 15 minutes later]$_{11}$ [and asked me again about Santa.]$_{12}$ [I pointed to his bike,]$_{13}$ [and I asked him who gave it to him.]$_{14}$ [He said Santa did.]$_{15}$ [I nodded]$_{16}$ [and said well then,]$_{17}$ [and he gave me a huge smile.]$_{18}$ [I felt a little guilty at the time about lying to my child,]$_{19}$ [but now I know that parenting is a balancing act.]$_{20}$ [Of course the truth is important,]$_{21}$ [but nothing trumps a mother's instinct.]$_{22}$

1. Subjective Statement
2. Subjective Statement
3. Objective Statement
4. Subjective Statement
5. Story Event
6. Story Other
7. None
8. Story Private State
9. Story Other
10. Story Event
11. Story Event
12. Story Event
13. Story Event
14. Story Event
15. Story Event
16. Story Event
17. Story Event
18. Story Event
19. Story Private State
20. Subjective Statement
21. Subjective Statement
22. Subjective Statement

Using Unexpected Simplicity to Control Moral Judgments and Interest in Narratives*

Antoine Saillenfest and Jean-Louis Dessalles

Department of Computer Science and Networks (INFRES)
Telecom ParisTech
Paris, France
`{antoine.saillenfest,jean-louis.dessalles}@telecom-paristech.fr`

─── **Abstract** ───────────────────────────────────

The challenge of narrative automatic generation is to produce not only coherent, but interesting stories. This study considers the problem within the Simplicity Theory framework. According to this theory, interesting situations must be unexpectedly simple, either because they should have required complex circumstances to be produced, or because they are abnormally simple, as in coincidences. Here we consider the special case of narratives in which characters perform actions with emotional consequences. We show, using the simplicity framework, how notions such as intentions, believability, responsibility and moral judgments are linked to narrative interest.

1 Introduction

Whatever a story is used for (whether to entertain or to teach), to address the question of what makes a story interesting is of major importance. The field of computational generation of narratives has explored many ways to generate narratives. A well-formed and understandable story can make a *good* story but not necessarily an interesting one.

Many studies have addressed the question of what makes a story interesting and different factors have been listed. One of them is the reference to major life themes that elicit strong affective reactions such as death, sex, religion or politics [16, 28]. Another factor is the occurrence of unexpected or unusual events [16, 28].

Beyond these universal determining factors, narrative interest may also crucially depend on the audience's personal values, personal experiences, personal emotions or pre-existing knowledge [29]. Taking into account personal values or personal experiences is a major challenge, but it also offers opportunities for automated scenario generation, especially for the replay options they offer.

Our review of previous works on narrative interest revealed that moral considerations towards the characters play a crucial role. We investigated how interest and moral judgments are related, in the specific context of moral dilemma stories [27]. We identified a set of factors that can be manipulated to control both interest and moral response. Using a framework

---

based on the notion of Kolmogorov complexity, we proposed a unified model of interest and moral judgment that can be applied to moral dilemma stories [27]. The present paper extends these previous results. It proposes new ways of controlling narrative interest by manipulating moral judgments, intentions and responsibility in narratives.

2 Some ingredients of interestingness

Humans are narrative beings. From childhood to adulthood, they are surrounded by narratives. They use them to make sense of the world, to order events and assimilate them. This process has been called *narrative intelligence* [21]. Narratives are a cognitive tool to organize experience and understand encountered situations. They play a major role in many forms of entertainment media including theater, novels or video games. They are also widely used in educational contexts.

The process of generating narratives has been the subject of intense study in Artificial Intelligence over the last century (see [11] for a review of existing systems). Different approaches to narrative generation have been explored. The standard approach consists in pre-scripting a specific story; characters perform the same actions without any variation each time the program is run. Systems using this approach present a limited number of stories (or permutations of a single story), they show little adaptation to the user and offer few opportunities of user interaction. An alternative approach consists in generating the narrative dynamically. Systems using this latter approach are able to adapt to individual preferences, needs, abilities or values. Such systems have extended replay capabilities.

A narrative is classically defined as the recounting of a sequence of events that have a continuant subject and constitute a whole [26]. A basic form of narrative is the *fabula*, a temporally ordered sequence of events from the time the story begins to the time the story ends [1]. The fabula is only one component of a typical narrative; the other one is the *discourse*, which refers to the way the narrative is told. The process of telling a narrative consists in selecting a subset of the fabula (the *sjužet*) and in structuring the outline of the main events in order to elicit interest in the audience [1, 11].

Generating a narrative consists in producing a sequence of events which brings a world from an initial state into a final state. This narrative content is then put into words to give the discourse. This process must meet some requirements to form acceptable narratives. One of them is to be well-structured, in the sense that it respects Freytag's triangle: the narrative should have a beginning, where some conflict is introduced, a middle, containing the climax, and an ending [9]. Another requirement for the generation of an acceptable narrative is to be understandable by the audience. In particular the events should respect the causal rules of the (possibly imaginary) world and the audience must be able to infer the characters' intentionality [3].

This generation process, however, does not ensure that these well-formed and understandable narratives make interesting stories. Producing interesting stories represents a major challenge, because of the crucial role of interest in learning or entertaining contexts [14]. Some of the factors controlling narrative interest have been investigated. They include emotional responses to fundamentals such as sex, religion or death [28, 16]. Surprising or unexpected events are also essential to interest (see below).

Few models of narrative interest have been proposed. Simplicity Theory [7] is one of them. It offers a purely cognitive account of interest, based on an extension of the notion of Kolmogorov complexity to the cognitive domain [4]. The main claim of Simplicity Theory is that unexpectedness, defined as a complexity drop, plays a major role in eliciting

narrative interest. It covers various situations humans usually regard as interesting, including coincidences, fortuitous encounters or rare events [6]. Simplicity Theory has also been applied to the study of moral judgment in the context of moral dilemma stories [27].

The observation that moral dilemmas often make good stories led us to investigate which factors control both narrative interest and moral judgments. Previous studies had pointed out the crucial role played by emotions and causal reasoning in the formation of moral judgments [12]. By applying Simplicity Theory to moral dilemma stories, we could show that factors such as the length of causal chains or the interpersonal relations between the characters can be manipulated in order to control interest and moral (dis)approval [27].

The formation of judgments on characters is an integral part of the narrative experience and is involved in the perception of character believability. Believable characters must be first perceived as intentional agents by the audience, who will then approve or disapprove of their actions. Simplicity Theory offers a new way to define and control, not only the interest of a generated story, but also the believability of its characters.

We will first review the main factors that influence moral judgments in narratives and point out the importance of these judgments for the believability of its characters. Then we will briefly introduce Simplicity theory and show how it can be applied to model intentions and moral judgments. We will illustrate how some parameters in the model can be used to control narrative interest and moral judgments. Finally, we will discuss the implications for narrative generation.

3   Interesting stories with moral characters

3.1   Believability and interestingness

Believable characters in a story should be perceived as clever and their actions as intentional [5], or else the audience's suspension of disbelief will be negatively affected [3]. But being perceived as intentional is not enough. The audience also forms various moral judgments including moral approval and responsibility. Approving or disapproving of others' actions and attributing responsibility is daily routine. We do it also when experiencing narratives. Let us consider the following story fragment.

The [**Wine Story**]. *Mary is looking through the window of her manor, waiting for a taxi. John enters the room, carrying a bottle of vintage red wine. He takes two crystal glasses in the cupboard and fills them. He discretely adds a pinch of black powder in one of them and brings it to Mary. The taxi will arrive in only a few minutes. Mary starts talking about her work, waving the hand and the glass. They hear the taxi's horn. She takes a sip from the glass, grabs her handbag, drinks the glass in one gulp and goes out.*

At this point, no one knows for sure what will happen next. One assumes that John performed his act (add powder to the drink) with some specific goal in mind. Imagine that Mary dies. If Mary is perceived as a nice person by the audience, John's act will be condemned. But the context can be manipulated to change this judgment, for example if John has been threatened to be killed if he did not kill Mary. This piece of information may also change our judgment about John's responsibility.

In any case, John appears as a believable agent. But this is only one ingredient of interest. Will the story be more or less interesting if Mary does not die? If John had selfish or altruistic motives (inherit Mary's fortune vs. save many from Mary's dark plans)? John's intentions and responsibility, in the eye of the audience, are a key factor in determining narrative interest.

To generate interesting stories automatically, one must first understand which factors influence our moral perception of the characters' actions. We will first review some previous works on morality and narrative interest before proposing our own approach.

3.2 Moral Judgments, Intentions and Responsibility

During the 1950s and the 1960s, studies on moral psychology focused on reasoning. Cognitive models of information processing were the preferred framework for the study of the formation of moral judgments [17]. However, in the 1980s, the idea that moral emotions played a significant role has been put forward. Recent evidence suggests that moral judgment is a matter of emotion *and* of reasoning, though automatic emotional processes tend to dominate [12, 13].

The formation of moral judgments occurs as soon as a character in a narrative performs an action that has good or bad consequences. The character's causal reasoning before the action is an essential determining factor that affects moral judgment. Jones (1991) proposed that the certainty or probability that the effect will follow from the action highly influences the intensity of our moral judgment [15]. In a previous study [27], we could observe that the *unexpectedness* of the outcome has a major influence on the moral approval of a character's action.

Causal reasoning is of major importance not only for the approval of others' actions, but also in the process of attributing intentions and responsibility. People are generally considered to be responsible for at least some consequences of their actions [30, 19], because they are supposed to be able to anticipate these consequences. The influence of causal analysis in the way responsibility is attributed to agents has been investigated [22, 23]. For an action to appear intentional, the agent must believe that this action can be performed and will lead to a specific outcome [20].

Moral judgment depends not only on causal reasoning, but also on the affective or emotional attitude that the characters or the audience have towards the consequences of the actions performed. Jones (1991) calls it the magnitude of the consequences [15]. The emotional response to a state of affairs determines what is desired or not and to what extent. It is essential to determine the character's intention to reach some specific goal [20]. Phares and Wilson (1972) also showed that it is of major importance in the attribution of responsibility [24].

Determining if a character in a story intentionally produced a situation or is responsible for it requires that we access its mental state at the time of the action [2]. Accessing mental states means inferring the character's knowledge and emotional state. In believable stories, sufficient information should be provided for the audience to assess the characters' mental states and to form moral judgments.

3.3 From moral judgments to interest

Our review of previous works on narrative interest revealed that moral considerations towards characters play a crucial role. We decided to investigate how interest and moral judgment are related, in the specific context of moral dilemma stories [27]. We especially pointed out the different but complementary roles of unexpectedness and emotional intensity in moral judgment and narrative interest. For example, we found that readers are more likely to disapprove of an action leading to undesired outcomes when the action appears more unexpected. But at the same time, the narrative interest of the sequence increases.

Using the framework of Simplicity Theory (see next section), we proposed a purely cognitive and unified model of interest and moral judgment that can be applied to moral dilemma stories. This model, based on the notion of Kolmogorov complexity, highlights the role of unexpectedness in the elicitation of emotional responses to events. The model proposed in our previous work explains our observations and makes correct predictions in the context of moral dilemma stories, taking into account the personal values of the audience and the estimated personal values of the characters of the story.

The present article extends these previous results and proposes new way of controlling narrative interest by manipulating judgments of intentionality and responsibility.

4   The key role of unexpectedness

4.1   Unexpected simplicity

As Livia Polanyi observed thirty years ago: "the question of what [people] tell stories about has remained not only unanswered, but largely unasked." [25, p. 207]. Polanyi, as the few authors who addressed the issue, acknowledged the importance of *unexpectedness* in spontaneously reported stories "stories have as their point that something 'odd' or 'unexpected' happened." [25, p. 212]. William Labov makes a similar observation: "[If an] event becomes common enough, it is no longer a violation of an expected rule of behaviour, and it is not reportable." When reporting an event, people say "[this] was strange, uncommon, or unusual – that is, worth reporting. It was not ordinary, plain, humdrum, everyday, or run-of-the-mill." [18]. Teun van Dijk also points the importance of unexpectedness: "interestingness is usually obtained by the account of events or actions that are unexpected, deviant, extra-ordinary, or unpredictable, given the knowledge and beliefs of the audience." [31, p. 123]. The same law of unexpectedness applies to newsworthiness: "Events have to be unexpected or rare, or preferably both, to become good news." [10, p. 67].

Unexpectedness stands out as a key feature of interestingness in spontaneous conversation. This is especially true for events leading to moral evaluation. Situations of deviant behaviour, leading to positive or negative moral judgments, constitute by definition exceptions to an expected norm, and are therefore unexpected. This may contribute to explaining why conversations are replete with moral evaluation and gossip [8].

What is true of spontaneous event reports must be true of fiction, at least up to some point. People are likely to use the same cognitive dispositions when they enjoy stories about factual events and about fictitious ones. Our hypothesis is that most parameters that control interest, both in conversation and in fiction, do it by controlling unexpectedness. It is thus important to reach a formal definition of that notion.

4.2   Defining unexpectedness

Unexpectedness is intuitively associated with an impression of low probability. Rare events, remarkable happenings, oddities, exceptions, deviations from norms are all extra-ordinary. They contrast with ordinary situation which, almost by definition, conform to expectations and are regarded as probable.

Probability unfortunately does not correctly capture the notion of unexpectedness. For instance, people may assess the unexpected nature of events that they would have previously regarded as impossible, such as the fact that two nuns start running on a jogging trail. Moreover, events of equal probability, such as lottery draws, are not regarded as equally

expected. A draw like 1–2–3–4–5–6 would be regarded as highly unexpected, whereas a "normal" draw like 13–23–24–31–35–44 is just boring news for those who did not play.

Simplicity Theory [7] is an attempt to capture the notion of unexpectedness. Unexpected (and therefore interesting) situations share the property of being less complex than expected. The word 'complexity' is used here in its technical sense, meaning 'size of the most concise description'. Unexpectedness results from a contrast between expected complexity (size of the most concise explanation of the event) and observed complexity (size of the most concise unambiguous designation of the event). In other words:

> *Unexpected situations are more complex to generate than to describe.*

This definition can be applied to the lottery example: all draws are equally complex to produce (if the machine is not biased). Most of them are complex to describe as well: their best description merely enumerates the numbers. There are therefore not unexpected. The remarkable draws are those which can be concisely designated. If a draw is a natural sequence such as 1–2–3–4–5–6, numbers need not be mentioned one by one and the designation, measured in bits of information, is much more concise. Similarly, exceptions are unexpected if what makes them exceptional is simple. A running nun is exceptional if all other nuns walk at a slow pace, because 'run' is a simple characteristic. A nun whose rosary has a missing bead may be unique as well, but as far as the 'rosary-with-missing-bead' characteristic is more complex than 'running', the event is less likely to be perceived as unexpected, and therefore as interesting.

Complexity is a concept of theoretical computer science. The corresponding quantity is measured in bits. Its abstract definition is the size of the shortest computer program, expressed in binary form, that outputs the object under consideration. We must distinguish two aspects of the notion.

Generation complexity: Size of the minimal instruction set that leads to a causal generation of the event.

Description complexity: Size of the minimal set of characteristics that leads to an unambiguous designation of the event.

Unexpectedness is defined by the difference (see also the Appendix):

$$
\begin{array}{rl}
 & \text{generation complexity} \\
- & \text{description complexity} \\
\hline
= & \text{unexpectedness}
\end{array}
\tag{1}
$$

When some feature, such as the fact of running, makes a situation (running nuns) unique, describing the situation once the feature is known requires no additional information. This explains why simple unique features (running, as opposed to missing-bead-in-rosary) make the situation more interesting.

4.3 Unexpectedness and emotions

Simplicity Theory has much to say, not about emotional experience *per se*, but about emotional *intensity*. The basic idea is that emotional intensity is controlled by unexpectedness. In what follows, any reference to 'emotion' has to be understood as meaning 'emotional intensity'.

Unexpectedness plays a decisive role in the **Wine Story**. If Mary suddenly dies, her death will be totally unexpected to all witnesses but John. Of course, any one may die

any time, but Mary's spontaneous death is unlikely. The difference between actual emotion (when Mary suddenly dies) and hypothetical emotion (considering Mary's possible death) is unexpectedness (see also the Appendix).

$$
\begin{array}{rl}
 & \text{actual emotion} \\
- & \text{hypothetical emotion} \\
\hline
= & \text{unexpectedness}
\end{array}
\tag{2}
$$

From this equation, we can deduce that the following terms are interdependent:
- the observer's emotional reaction to Mary's death (actual emotion)
- the observer's concern about Mary's possible death (hypothetical emotion)
- the complexity of the simplest causal scenario leading to Mary's death that the observer may think of (generation complexity).

If the observer never anticipated Mary's death, the reference may change. If the observer compares Mary's death with her/his own, the equation will involve the following terms:
- the observer's emotional reaction to Mary's death (actual emotion)
- the observer's concern about her/his possible own death (hypothetical emotion)
- the complexity of the simplest causal scenario leading to Mary's death that the observer may think of (generation complexity)
- the minimal description of Mary for the observer (description complexity)

This explains why Mary's death will be perceived differently by different people, depending on their closeness to the victim.

4.4   Emotions and actions

If the black powder is poison, Mary's death is not unexpected from John's perspective. When Mary dies, John is not surprised and his emotion does not change. . . unless the causal link from his action (adding the powder) to the effect (Mary's death) is not perfect. In that case, the preceding equation can be rewritten as:

$$
\begin{array}{rl}
 & \text{anticipated emotion} \\
- & \text{causal unexpectedness} \\
\hline
= & \text{hypothetical emotion}
\end{array}
\tag{3}
$$

Causal unexpectedness is just an instance of the general notion of unexpectedness. It is evaluated when an agent performs an action. If the causal link to the emotional outcome involves many intermediary steps, the agent will perceive the outcome as more complex to generate. The outcome will be more unexpected and the gap between hypothetical and actual emotion will be larger. In the **Wine Story**, adding poison to the drink is not enough to provoke Mary's death. John's act must remain unnoticed, Mary should drink the glass of wine, the poison has to be effective, at a time when Mary is far from a hospital, and so on. The occurrence of Mary's death is still unexpected, and thus emotional, from John's perspective.

4.5   Intentionality, responsibility and moral judgment

The preceding results follow from Simplicity Theory's basic principles. They already provide theoretical tools to control emotional intensity when events are told in a narrative, from anyone's perspective (readers or audience, but also other characters). When it comes to

anticipating the emotional effects of characters' actions, some additional notions such as intentionality, responsibility and moral judgments must be given precise definitions as well.

Our previous works on moral judgment suggest that narrative interest and moral approval of one's action can be controlled by various factors such as the causal unexpectedness or the emotional intensity of the outcomes [27]. We will now use the framework of Simplicity Theory to define the notions of intention and responsibility.

Intentions are a matter of desires. One intends to perform an action if one desires some resulting situation to happen. The degree of intention to perform an action in order to produce some desired outcome depends on the anticipated emotional benefit. In the **Wine Story**, suppose that Mary dies because John added some poison in her glass. From John's perspective, adding poison into Mary's glass makes her death highly expected. John's degree of intention to kill Mary by adding poison depends on his expected emotional response to her death. In other words, the more John desires that Mary dies, the more his act of adding poison is judged intentional. Intention can be defined using the following equation (see also the Appendix).

$$
\begin{array}{rl}
 & \text{agent's anticipated emotional response to the outcome} \\
- & \text{causal unexpectedness (from the agent's perspective)} \\
\hline
= & \text{intention of performing the action to produce the outcome}
\end{array}
\tag{4}
$$

Some anticipated consequences of one's action may not be desired. These negative anticipated consequences tend to decrease the global intention to perform an action. In such cases, these negative outcomes are brought up knowingly. We must suppose that the action has been performed for some other purpose.

When an action has several emotional effects, as in moral dilemma, the various intentional values add up, with negative signs when the agent's emotion attached to the outcome is negative.

$$
\begin{array}{rl}
 & \text{intention of performing the action to produce the outcome 1} \\
+/- & \text{intention of performing the action to produce the outcome 2} \\
+/- & \ldots \\
\hline
= & \text{intention of performing the action}
\end{array}
\tag{5}
$$

Responsibility is defined in a similar way. There is, however, a crucial difference. The observer's emotions (or sometimes standard social emotions) must be substituted for the agent's ones. The responsibility attributed by an observer to an agent can be defined using the following equation (see also the Appendix).

$$
\begin{array}{rl}
 & \text{observer's anticipated emotional response to the outcome} \\
- & \text{causal unexpectedness (from the agent's perspective)} \\
\hline
= & \text{agent's responsibility for the outcome}
\end{array}
\tag{6}
$$

The observer's moral judgment about the agent's action is a cumulative responsibility (taking signs into account). The signs depend on the desirability or the undesirability of the outcomes from the observer's perspective.

$$
\begin{array}{rl}
 & \text{responsibility for outcome 1} \\
+/- & \text{responsibility for outcome 2} \\
+/- & \ldots \\
\hline
= & \text{moral judgment about the agent's action}
\end{array}
\tag{7}
$$

Using these notions of intention, responsibility and moral judgment based on Simplicity Theory, we can not only anticipate moral evaluations in a narrative, but also control the characters' credibility. Rational characters will perform actions that realize the best compromise between the different judgments they can anticipate. The same notions can also be used to control interest, by maximizing the observer's emotion.

5   Controlling interest and moral judgment in a narrative context

Unexpected simplicity is our best candidate to control narrative interest and moral judgments in narratives. In this section, we will see that, by modifying the context of the **Wine Story**, we can influence not only the moral judgments experienced when reading the story, but also its interestingness.

5.1   Characters' desires

In the **Wine Story**, John may attempt to poison Mary in order to inherit her fortune. If so, John's anticipated emotion concerning Mary's death depends on various factors including her wealth and John's personal feelings for her. John's desires therefore modify his intention to act (see equations 4 and 5). For John to appear as a believable character, his action must make sense, which means that he should only perform acts that are significantly intentional. Equation 4 explains why actions should be consistent with the agent's desires.

Readers' emotional reaction to Mary's death is crucial to know whether John will be blamed or praised (see equations 6 and 7). This reaction depends on various factors, such as their empathy for Mary or their personal identification with her. These effects are captured by Mary's simplicity in the reader's eyes (see equation 2), whether she is for instance a main or a peripheral character.

Anticipated emotions are crucial in the case of moral dilemmas. Imagine that John loves Mary but is forced to kill her or else he will be killed himself. John anticipates two emotions, one for Mary's death, another one for his own. If John chooses to kill Mary, then the higher his anticipated emotional response to Mary's death, the more interesting the story. In such a situation, John kills Mary with a lower degree of intentionality (equation 5). His emotional reaction to the outcome is compensated by the fact that he will not get killed himself. The dilemma also affects the observers' attribution of responsibility if they have similar emotions (equation 7). Observers (readers or other characters of the story) who are ignorant of the death threat hanging over John may form a rather different judgment. To them, John's act will be highly unexpected (see equation 2). This example highlights the importance of the way information is provided to readers and to the characters of the story.

5.2   Unexpectedness in causal chains

When reading the **Wine Story**, one can easily imagine that John is attempting to poison Mary. Now imagine that the poison turns out to have no effect on her. This introduces surprise in the story because the end of the sequence of events is unexpected. There are different ways to introduce unexpectedness in a causal sequence during the process of story generation. The action itself can be unexpected for the reader, for the acting character or for any other character of the story (see the Appendix for a commentary about inadvertence); the expected consequence may fail to happen for some unanticipated reason; or the action provokes not only its expected outcomes, but other unanticipated consequences as well.

5.2.1 Action unexpectedness

In a written story, a character's action may be unexpected for the reader. Action unexpectedness should be manipulated carefully, as it may negatively affect the reader's "suspension of disbelief". A way to make an action unexpected is to keep the agent's goals untold. However, sufficient information has to be disclosed for readers to be able to infer a causal explanation for the unexpected action. Not only readers, but also other characters in the story (as readers imagine them) may also be surprised by a character's action. Even the acting character may be surprised by her/his own action, when it has been performed inadvertently. In such case, inadvertence may considerably reduce blame or praise (equation 7 and the Appendix).

5.2.2 Causal unexpectedness

Causal unexpectedness refers to the unexpectedness of the anticipated consequence once the action has been performed. If the causal chain of events that goes from the action to its anticipated consequence consists in many intermediary events, the occurrence of the consequence will appear less expected. More causal unexpectedness decreases the attribution of intentions (equation 4) and of responsibility (equation 6), and the amplitude of the blame or praise will diminish as well (equation 7). But more causal unexpectedness increases the gap between hypothetical emotion and actual or anticipated emotion, both for the acting character and for observers (equations 2 and 3). Note that in all cases, causal unexpectedness is supposed to be assessed by the agents, with the knowledge and inference capabilities that observers grant them.

In the **Wine Story**, we suppose that John knows that the poison will be effective. The only unexpected step, from John's perspective, is to know whether Mary will drink up her glass of poisoned wine, as the unexpectedness of Mary's death once she has drunk it is nearly zero. There is room for suspense in the first phase, because of the time pressure (the taxi will arrive soon). The occurrence of "Mary drinks the wine" becomes increasingly unexpected as the time is running out.

Causal unexpectedness should also be manipulated with care, as it may negatively affect readers' suspension of disbelief. There is a risk that agents be no longer believable if their intentionality drops too much (equation 4).

5.2.3 Unexpected consequences

Of course, the course of events may differ from what characters and readers anticipate. The simplest case is when the outcome is the exact opposite of the anticipated one. In the **Wine Story**, Mary may not die nor even get sick from having drunk the wine. This turn of events is unexpected, because observers must imagine complex circumstances that may have produced the alternate ending (equation 1). Such unexpectedness must be resolved, at least for readers. This means that some explanation must be found or must be provided that provides a simpler generation process for the event.

Any action may have various consequences besides its expected consequences. These side-effects are interesting, because they are unexpected. But as for deceived anticipations, their unexpectedness must be resolved, at least for readers. Unexpected consequences are not intentional since the agent did not anticipate them (equation 4), but the agent may be judged responsible for them (equation 6).

6   Conclusion

We have shown various ways through which Simplicity Theory can be used to manipulate emotional responses and interest in narratives. The same parameters also affect the believability of characters. The main parameter underlying narrative experience is unexpectedness. It is defined as the difference between the complexity of the circumstances that brought the situation about and the amount of information required to describe it. Unexpectedness is a key factor to predict interestingness in conversational narratives. It is also involved in fictitious stories.

When actions with emotional consequences are performed in a story, unexpectedness is involved three times. Once when the consequences are unexpected, independently from the action that caused them (e.g., when a character dies). A second time if the action itself is unexpected (e.g., if the action has been performed inadvertently, as in an accident). And a third time when the causal link is regarded as unexpected by the agent (e.g., when many steps separate the action from its emotional consequences). The various roles of unexpectedness offer as many possibilities to control interest in narratives.

This paper investigated these possibilities from a theoretical point of view. We are currently planning to implement them in a narrative generation program. The main difficulty comes from the fact that the precise values of complexity are sometimes difficult to compute. Various strategies can be used to overcome the problem. Generation complexity relies on the addition of choice points (see www.simplicitytheory.org). Description complexity for imaginary objects or characters is easy to compute by ranking those objects or characters by importance. For objects or people of real life like the Eiffel Tower or Barack Obama, a rough estimate can be deduced from the number of hits on a Web search engine.

Our conviction is that computational models of narrative cannot ignore the various parameters that control unexpectedness, emotions and moral judgments. This paper was a first attempt to list these parameters, to provide tentative definitions and to consider their logical relations.

A   Appendix

Here is a formal summary of the notions presented in this article. Corresponding equations in the article and in this appendix share the same number.

Unexpected situations are more complex to generate than to describe

$$U(s) = C_w(s) - C_d(s) \tag{1}$$

$U(s)$ is the *unexpectedness* of situation s, as perceived by the observer. $C_w(s)$ measures the minimal amount of information that is necessary for the 'world' to produce s. This information evaluates the size of the minimal explication of the situation. It corresponds to the circumstances that the observer must imagine for s to happen. Description complexity $C_d(s)$ measures the quantity of information that the observer needs to unambiguously determine s.

The basic idea of Simplicity Theory is that emotional intensity is controlled by unexpectedness. It says that the expected or hypothetical emotional intensity $E_h(s)$ attached to the possible occurrence of an event s is related to the actual emotional intensity $E(s)$ attached to the effective occurrence of s:

$$E(s) - E_h(s) = U(s) \tag{2}$$

When some outcome s of an action a occurs, one can evaluate how necessary the action was, from any observer's point of view, to produce s. The previous equation can be rewritten as:

$$E(s) - U(s||a) = E_h(s) \tag{3}$$

The causal unexpectedness $U(s||a)$ is an instance of the general notion of unexpectedness. It evaluates the unexpectedness of a situation s once an action a has been performed. $E(s)$ is the emotional intensity attached to s as it has been anticipated. $E_h(s)$ is the hypothetical emotion attached to the occurrence of s.

The preceding results follow from Simplicity Theory's basic principles. They already provide theoretical tools to control emotional intensity when events are told in a narrative, from anyone's perspective (readers or audience, but also other characters).

Now we will define some additional notions such as intentionality, responsibility and moral judgment.

The degree of intention to perform an action a in order to produce some desired outcome s is the anticipated emotional benefit. For any observer O, the degree of intention $(Int^O(A, a, s))$ for an actor A is the hypothetical emotion attached to the occurrence of s from A's perspective $E_h^A(s)$ (the upper letter indicates that the calculation is calculated by O from A's perspective by inferring A's mental state).

$$Int^O(A, a, s) = E^A(s) - U^A(s||a) \tag{4}$$

Undesired but anticipated outcomes contributes to decrease the global intention to perform an action. They are produced knowingly. We must suppose that the action has been performed either accidentally $(Int^O < 0)$ or for some other purpose. Accidental actions are unexpected. This unexpectedness is a negative source of intentionality. Equation 4 should be rewritten: $Int^O(A, a, s) = E^A(s) - U^A(s||a) - U^A(a)$. For voluntary actions, $U^A(a) = 0$.

The global intention to perform an action a assigned by O to A results from the accumulation of the degrees of intention attached to all the outcomes s, with negative signs for undesired ones $(\epsilon(s) = -1)$, in the set S of all the outcomes that have been anticipated by A.

$$Int^O(A, a) = \sum_{s \in S} \epsilon(s) Int^O(A, a, s) \tag{5}$$

Responsibility is defined in a similar way. There is, however, a crucial difference. The observer's emotions $(E^O(s))$ must be substituted for the actor's ones $(E^A(s))$. An observer will find an actor more responsible if, in the same circumstances, the actor's actions elicit a more intense emotional response in the observer. Then the responsibility attributed to A by O for the occurrence of the outcome s which is a consequence of A's action a is:

$$Res^O(A, a, s) = E^O(s) - U^A(s||a) \tag{6}$$

O's moral judgment (or moral approval) MJ about A's action is a cumulative responsibility (taking signs into account). The signs depend on the desirability or the undesirability of the outcomes for the observer O.

$$MJ^O(A, a) = \sum_{s \in S} \epsilon(s) Res^O(A, a, s) \tag{7}$$

References

1   Mieke Bal. *Narratology: Introduction to the theory of narrative.* University of Toronto Press, 1997.

2   Simon Baron-Cohen. The evolution of a theory of mind. In Michael C. Corballis and Stephen E. G. Lea, editors, *The descent of mind: Psychological perspectives on hominid evolution*, pages 261–277. New York, NY: Oxford University Press, 1999.

3   Joseph Bates. The role of emotion in believable agents. *Communications of the ACM*, 37(7):122–125, 1994.

4   Nick Chater and Paul M. B. Vitányi. Simplicity: a unifying principle in cognitive science? *Trends in Cognitive Sciences*, 7(1):19–22, 2003.

5   Daniel C. Dennett. *The intentional stance.* MIT Press, Cambridge, MA, 1989.

6   Jean-Louis Dessalles. Coincidences and the encounter problem: A formal account. In B. C. Love, K. McRae, and V. M. Sloutsky, editors, *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, pages 2134–2139, Austin, TX, 2008. Cognitive Science Society.

7   Jean-Louis Dessalles. *La pertinence et ses origines cognitives : nouvelles théories.* Hermes-Science Publications, Lavoisier, 2008.

8   Robin Dunbar. *Grooming, gossip, and the evolution of language.* Harvard University Press, 1998.

9   Gustav Freytag. *Freytag's technique of the drama: an exposition of dramatic composition and art.* Scott, Foresman, 1894.

10 Johan Galtung and Mari H. Ruge. The structure of foreign news, the presentation of the Congo, Cuba and Cyprus crises in four Norwegian newspapers. *Journal of Peace Research*, 2(1):64–90, 1965.

11 Pablo Gervás. Computational approaches to storytelling and creativity. *AI Magazine*, 30(3):49–62, 2009.

12 Joshua D. Greene and Jonathan Haidt. How (and where) does moral judgment work? *Trends in Cognitive Sciences*, 6(12):517–523, 2002.

13 Joshua D. Greene, Leigh E. Nystrom, Andrew D. Engell, John M. Darley, and Jonathan D. Cohen. The neural bases of cognitive conflict and control in moral judgment. *Neuron*, 44(2):389–400, 2004.

14 Suzanne Hidi. Interest and its contribution as a mental resource for learning. *Review of Educational Research*, 60(4):549–571, 1990.

15 Thomas M. Jones. Ethical decision making by individuals in organizations: An issue-contingent model. *The Academy of Management Review*, 16(2):366–395, 1991.

16 Walter Kintsch. Learning from text, levels of comprehension, or: Why anyone would read a story anyway. *Poetics*, 9(1–3):87–98, 1980.

17 Lawrence Kohlberg. *The development of modes of moral thinking and choice in the years 10 to 16.* PhD thesis, University of Chicago, 1958.

18 William Labov. The transformation of experience in narrative syntax. In William Labov, editor, *Language in the Inner city*, pages 354–96. Philadelphia: Pennsylvania University Press, 1972.

19 David A. Lagnado and Shelley Channon. Judgments of cause and blame: The effects of intentionality and foreseeability. *Cognition*, 108(3):754–770, 2008.

20 Bertram F. Malle and Joshua Knobe. The folk concept of intentionality. *Journal of Experimental Social Psychology*, 33(2):101–121, 1997.

21 Michael Mateas and Phoebe Sengers. Narrative intelligence. In M. Mateas and P. Sengers, editors, *Narrative Intelligence: Papers from the AAAI Fall Symposium*, number FS-99-01 in AAAI Technical Reports, pages 1–10, 1999.

22 Dale T. Miller and Saku Gunasegaram. Temporal order and the perceived mutability of events: Implications for blame assignment. *Journal of Personality and Social Psychology*, 59(6):1111–1118, 1990.

23 Ahogni N'gbala and Nyla R. Branscombe. Mental simulation and causal attribution: When simulating an event does not affect fault assignment. *Journal of Experimental Social Psychology*, 31(2):139–162, 1995.

24 E. Jerry Phares and Kenneth G. Wilson. Responsibility attribution: Role of outcome severity, situational ambiguity, and internal-external control. *Journal of Personality*, 40(3):392–406, 1972.

25 Livia Polanyi. So what's the point? *Semiotica*, 25(3-4):207–242, 1979.

26 Gerald Prince. *A Dictionary of Narratology*. University of Nebraska Press, Lincoln, NE, 1987.

27 Antoine Saillenfest and Jean-Louis Dessalles. Role of Kolmogorov complexity on interest in moral dilemma stories. In N. Miyake, D. Peebles, and R. Cooper, editors, *Building Bridges Across Cognitive Sciences Around the World, Proceedings of the 34th Annual Meeting of the Cognitive Science Society, Sapporo, Japan, August 1–4, 2012*, pages 947–952, Austin, TX, 2012. Cognitive Science Society.

28 Roger C. Schank. Interestingness: Controlling inferences. *Artificial Intelligence*, 12(3):273–297, 1979.

29 Gregory Schraw and Stephen Lehman. Situational interest: A review of the literature and directions for future research. *Educational Psychology Review*, 13(1):23–52, 2001.

30 Kelly G. Shaver. *The attribution of blame: causality, responsibility, and blameworthiness*. Springer-Verlag, 1985.

31 Teun Adrianus van Dijk. Stories and racism. In Dennis K. Mumby, editor, *Narrative and Social Control: Critical Perspectives*, pages 121–142, Newbury Park, CA, 1993. Sage Publications.

Narrativity and Textuality in the Study of Stories

Moshe Simon-Shoshan

Rothberg International School
Hebrew University
Jerusalem, Israel
`mdshoshan@gmail.com`

―――― **Abstract** ――――――――――――――――――――――――――――――――――――――

This paper seeks to investigate some of the defining elements of narrative. The underlying assumption of my discussion is that the terms "narrative" and "story" do not refer to clearly defined, self-enclosed genres. Rather, they are part of a spectrum which embraces all forms of texts. Similarly, narratives and stories are not independent discourses but rather are an integral part of virtually all forms of discourse, be it day-to-day conversation or more specialized discourses. In order to analyze the relationship between narratives and other modes of discourse, we introduce the concept of narrativity. Narrativity refers to a collection of textual attributes. All texts exist along a continuum of greater or lesser narrativity, depending on the number and prominence of the narrative attributes they contain. When we refer to a text as a story, we mean that it contains a critical mass of narrativity. Most theorists of narrative have defined narrativity purely in terms of "dynamism"–that is, the extent to which a text portrays transition and change. To this I have added the quality of "specificity". Specificity refers to the extent to which a text focuses on a particular time or place, a unique event, or individual people and objects. Many if not most texts contain a certain degree of narrativity. We established, however, that in order to be considered a story the text must present a sequence of at least two interrelated events that occurred once and only once in the past. In other words, a story must have a certain degree of dynamism in that it portrays the transition from at least one event to another. It must also have specificity at least to the degree that the text narrates events that happened at a fixed time in the past. This theoretical framework allows us to chart the relationship between different types of texts within a single discourse. It also gives us a vocabulary for discussing different parts of more complex narratives which often contain elements of varying narrativity. The paper then goes on to discuss the concept of narrative structure, arguing that narrative structure is not an inherent attribute of narrative texts but a framework that the reader imposes on the text in order to make it intelligible in terms of other narratives. The structure which the reader abstracts from a given narrative will be heavily dependent on the context of the narrative with in a wider discourse.

1 Introduction

In this paper I would like to share with you some of my conclusions about the nature of narrative that have emerged from my work on narratives found in the Mishnah, the 3rd century CE Jewish legal compilation.[1] Some of the key features of Mishnaic narrative

―――――――――――――――

[1] My work on this topic if fully presented in my book [21].

make my ideas particularly relevant to efforts to model the process of story comprehension. First, stories in the Mishna are very short, many no more than a single sentence. Such brief narratives allow us to focus on the most fundamental elements of narrative, without getting distracted or confused by elements present in more complex stories. Second, Mishnaic narratives are embedded within a larger discourse, that of early rabbinic law. This allows us to consider the relationship between stories and other linguistic forms and the way in which stories, outside of autonomous literary creations, tend to integrated into some larger discourse. Finally, Mishnaic stories fall into a particular category of brief, embedded narratives, namely, anecdotes. Anecdotes are accounts that are used as rhetorical tools to reinforce a particular point. One of the fundamental problems of discussing the process of narrative comprehension is the fact that it is difficult to talk about the "meaning" of a particular story. Stories are by definition indirect means of conveying meaning. Most stories possess multiple potential meanings and implications. Which meanings any given reader will perceive in a story is highly dependent of perspective and context. This makes it very difficult to define a "correct" understanding of a story and its implications. Since anecdotes are generally deployed within a wider discourse for a specific purpose, it easier to fix a single primary meaning to a given anecdote. This in turn makes it easier to set a standard of what constitutes "understanding" a given story. I would like to deal with two main narratological issues, the definition of narrative as a linguistic product and the role of narrative structure in understanding stories.

2 Defining Narrative

Narrative is perhaps the most ubiquitous and multifarious of all literary forms. Stories, be they epic poems or modern novels, hold prominent places in the literary canons of virtually every culture. Yet storytelling is hardly the sole preserve of the belles lettres. Narratives are also an important feature of many other forms of discourse, including the study of law, medicine, history, and philosophy. William Labov went so far as to argue that "narratives are privileged forms of discourse which play a central role in almost every conversation" [15, p. 396]. We all tell stories in our day-to-day speech. Small children learn this craft as part of the normal process of language acquisition. Indeed, the potential for creating stories may well be one of the fundamental, universal characteristics of language [6, pp. 96–97]. A definition of the terms "narrative" and "story" must thus take into account the fact that these forms are often enmeshed with other linguistic structures and modes of discourse. On the other hand, such a definition must also allow for the richness and complexity of more developed, "literary" narratives. In order to understand how narrative relates to other forms of discourse, it is important to realize that the terms "narrative" and "story" do not refer to clearly defined, self-enclosed genres. Rather, they are part of a spectrum which embraces all forms of texts. This spectrum can be charted on the basis of what I call "narrativity". Narrativity refers to a collection of textual attributes. All texts exist along a continuum of greater or lesser narrativity depending on the number and prominence of the narrative attributes they contain. When we refer to a text as a "narrative" or a "story", we mean that it contains a certain critical mass of narrativity. However, the precise line between "narratives" and "non-narratives" is inherently arbitrary. I will present one possible definition of the term "story" as marking a key point of transition along the continuum of narrativity to be found in texts. Narrativity emerges from the confluence of two distinct elements in a text which I call "dynamism" and "specificity". "Dynamism" refers to the fact that narratives are fundamentally about transition, transformation, and change. "Specificity" indicates that narratives are rooted in the particular, focusing on individual characters and unique events,

and occur at demarcated points in time and space. Narrativity thus inheres in texts to the extent that they describe change and transition while at the same time focusing on the concrete, the specific, and the time-bound. In their definitions of narrative, narratologists overwhelmingly tend to focus on dynamism at the expense of specificity. For example, E. M. Forster, in his classic work Aspects of the Novel, suggests the following sentence as an example of a minimal story:[2]

The king died and then the queen died of grief. [10, p. 116]

In his analysis of this text, Forster focuses exclusively on its dynamic aspects. For him, the salient qualities that make this text a story are that it is a "narrative of events arranged in their time sequence . . . the emphasis falling on causality". [10, p. 116] This statement succinctly encapsulates all of the formal traits that collectively define the minimum requirements for dynamism in a narrative or story. First, Forster presumes that the constituent elements of narratives are "events".[3] For a text to be dynamic, something has to happen. Narratives are first and foremost the representation of happenings. Events are thus the fundamental building blocks of narratives. To be sure, it is not always a simple matter to break down a story into discrete events.[4] However, in relatively simple narratives such as those found in the Mishnah, a definition of "event" is possible. In such narratives, each time a narrative presents a verb that describes some action or change of state we have a representation of a narrative event. The complete event is represented by the entire clause in which the verb appears. See [17, p. 17]. Narrativity can thus be associated with the presence of dynamic verbs in a text.[5] In the case of Forster's story, we have two events each centered on the verb "died". The first event is "the king died". The second is "the queen died of grief". Returning to Forster's statement that a narrative consists of "events arranged in their time sequence", the next operative word in this phrase is "time". The dynamic nature of narratives demands that they portray the passage of time. This is done through the representation of at least two events in sequence, creating the illusion of a seamless continuum of time moving inexorably forward. In the case of the death of the king and queen, the reader does not merely experience two discrete events. Rather, through the phrase "and then", the reader follows the passage of time from the death of the king through the death of the queen. Finally, we come to the element of narrative which Forster calls "causality". Forster argues that the sentence "The king died and then the queen died" is not, despite its narrative elements, a narrative. In order to transform this text into a narrative, we must add the final clause, "of grief". True dynamism requires not only that the events follow each other sequentially but also that they be inherently interrelated. It is not sufficient that the death of the queen chronologically follows that of the king. The queen's death must follow from that of the king. This is what

---

[2] Forster makes use of a different set of terms than I do. Forster calls this text a "plot", in contrast to a "story", which for him refers to a set of events that lack causality. For the sake of clarity, I have replaced Forster's terms with my own.

[3] Aristotle already declared that "the plot is the mimesis of the action–for I use 'plot' to denote the construction of events". [12, p. 39].

[4] Suzanne Fleischman presents a useful survey of the critical discussion surrounding the question of "events" in narrative [8, pp. 97–100].

[5] However, as I argue in greater detail in my book, stative verbs still have a place in narratives. This is because, first, some stative verbs can be used to describe changes in state or status such as in the sentence "The king died and the queen was very sad.". Here the verb "was sad" clearly indicates that the queen only became sad after the death of the king and hence describes a change or transformation. Furthermore, sentences with little to no narrativity can certainly play a role in more complex stories. See [21, pp. 17–18, 20–21].

Forster and others call the need for "causality" in a narrative.[6] The term "causality" suggests a degree of inevitability and determinism that does not necessarily reflect the contingent manner in which narrative events often unfold. A better formulation of this requirement is presented by Binder and Weisberg, who state that a narrative must present "one event as standing in some relation of significance to a later event such that one is made meaningful by the other". [4, p. 221]. The authors make this statement in explicating the position of Arthur Danto. Taking my cue from these scholars, I will refer to the interrelationship of events in a narrative rather than using the more problematic term "causality". Taking Forster's definition we might thus posit that in order to be considered a "story" a text must possess the following "dynamic" characteristics: (1) It must be a representation of events (2) It must present two or more events in sequence and (3) these events must be inherently interrelated in such a way as to portray some change in the world represented by the text. Something is missing however. This quality of dynamism does not sufficiently describe the characteristics of texts we call "stories" or "narratives". To illustrate the problem, let us alter slightly Forster's paradigmatic story as follows:

> Kings die, and then their queens die of grief.

This text also fits Forster's definition of a story and it contains all of the dynamic elements that we have established as necessary for our definition of a story as well. Yet, most readers would agree that this text is in some way less of a story than Forster's original formulation. The missing element in this text is the quality that I call "specificity". Rather than referring to a specific, if totally anonymous king and queen like Forster's story, this new text purports to describe a general phenomenon which applies to any number of monarchs and their consorts in different times and places. Though we are not told either the names of the king and queen described by Forster, or when and where they ruled, the story clearing indicates that it tells of one and only one king and one and only one queen . The events described clearly happened once and only once. Stories thus need to do more than portray change. The must portray change taking place in a specific context in time and space, rather than in a generalized situation. It is not necessary that the time and locale of the story be disclosed. The important thing is that the story portrays one-time events regarding a definable group of individuals or objects. On the basis of this criterion we can fine tune our previous definition of a story as follows:

> A story is any representation of a sequence of at least two interrelated events that occurred once and only once in the past.[7]

---

[6] Aristotle anticipates the need for causality at the end of Book 9 of his Poetics [12].

[7] Fleischman presents two similar definitions of a story which compliment my own. First she defines narrative in strictly linguistic terms. She defines the "constituent properties"of narrative as being "past time reference, perfective aspect, and a distanced, objective perspective on events that are realis, semelfactive (unique occurrence), and sequentially ordered". [12, p. 55]. Later, Fleischman quotes Susan Herring's previously unpublished description of a "prototypical narrative":

> The prototypical past tense narrative is concerned with events rather than static description, and the events are not narrated in random order but rather in a sequence which is iconic with the temporal order in which they actually occurred. ... Further, the completion of one event is implied by the inception of that which follows, a fact, which may give rise to an interpretation of aspectual perfectivity for the (simple) past tense where no other value is specifically indicated. ... The prototypical narrative is factual and time-bound, in that it chronicles a unique set of events, which took place at a specific point (or over a specific bounded interval) in time. There is also a sense in which the ideal narrator is objective, maintaining distance between him or herself and the events narrated in order to relate them as they actually occurred, in linear order with a minimum of personal evaluation or digression. It is this complex of features which, in the absence of indications to contrary, the "narrative past" typically evokes [8, p. 101].

Most stories contain more than this bare bones level of specificity. Stories in which the characters and their motives, as well as the setting and background of the story, are described in great detail would thus be considered to have a greater level of narrativity then less descriptive narratives. Similarly, stories that focus on specific individuals contain more specificity than those which portray large groups. One of the genre that presents the greatest level of specificity is certainly the novel. It is therefore hardly surprising that critics tend to identify specificity as a distinguishing feature of modern realistic fiction, rather than as a fundamental aspect of narrativity (see, e.g., [14, pp. 19–27] and [2, pp. 11–17]). However, the fact that we intuitively differentiate between Forster's example and my more general text demonstrates that specificity is something that we associate even with some of the simplest of narrative texts. One of the most important indicators of specificity is grammatical tense. Numerous critics have pointed out that the past tense, or more specifically the preterit, is the primary tense for storytelling.[8] Events in stories must happen once and only once at a definable point in time. Only the past tense can fully provide the sort of concreteness and specificity necessary for stories. Another way of further defining the need for specificity is that a story must recount events using verbs in the realis mood (see [15, p. 400]). That is, the events described by them are represented as having been realized in the material world. The reality portrayed by realis verbs may be fictional, nonfictional, or some combination of the two. The important thing is that the reader is called upon to imagine an actual situation, event, or story (see [20, pp. 98–99]). Realis accounts are to be distinguished from irrealis accounts, which "are verbalizations of experience that is unrealized either because it is predicated on taking place in the future or because it is in some sense hypothetical". [8, p. 104]. Unlike the need for dynamism, which is emphasized in one form or another by virtually all narrative theorists, the need for specificity has generally been minimized or even ignored by most students of narrative. Gerald Prince and Wendy Steiner are among the few narrative theorists who emphasize the need for specificity in narrative ([18, pp. 61–76] [22]).[9] I am aware of only one writer who presents a definition of narrative that focuses on specificity almost to the exclusion of dynamism; G. A. Gaballa writes: "A story is a specific event carried out by particular characters in a particular place." [11, p. 5]. It is hardly a coincidence that both Steiner and Gaballa are interested in narrative expressed through painting. Steiner emphasizes that, as an essentially atemporal medium, painting can possess only limited dynamism. The flip side of this is the clichéd observation that a picture is worth a thousand words. It is precisely a painting's ability to simultaneously present a vast number of details that makes it especially suited to the expression of specificity, even beyond that of written texts. The notion that a still life or landscape painting might possess narrativity might seem to stretch the normal uses of the word "narrative" and "story" beyond recognition. It is precisely my intent to provoke the reader to rethink conventional understandings of these terms. Ultimately narrative is about more than action and change. It is also about representing and engaging the particular and unique aspects of individuals, objects, and

---

Both Fleischman and Herring add an additional narrative attribute to the two I propose, namely, that the events be narrated in an objective manner. This opens up the possibility of a third general category of narrative attributes, namely, that narratives are narrated. See also [13, pp. 209–32].

[8] Perhaps the earliest systematic attempt to demonstrate the integral relationship between the past tense and narrative was undertaken by Emile Benveniste with regard to the French language in [3]. Fleischman offers a more thorough, cross-linguistic study of the role of tense in narrative and the past tense in particular. See [8]. See also Roland Barthes' comments on the French passé simple in [1, p. 34]. Barthes emphasizes the role of the past tense in what I have called the dynamic nature of narrative.

[9] Fludernik's concepts of "narrativity" and "narrativization" are also ultimately rooted in the specific nature of the narrative experience, though she works from very different premises, [9, pp. 20–43].

situations. On the basis of these definitions it is possible to chart the relative narrativity of a set of texts along the axes of dynamism and specificity using primarily linguistic markers. Verbs that suggest action and change, especially when they coordinate the relationship between two nouns, or when they come in series, tend to mark dynamism. Specificity is marked by realis forms which limit time and place and the use of terms that reflect individuals or narrow categories. Closely related to stories are texts which contain dynamism but lack sufficient specificity to be stories, such as iterative narratives that tell of events multiple times and non-realis descriptions of interconnected events, such as hypothetical scenarios on the one hand and texts which are specific but lack dynamism such as detailed descriptions of static situations on the other. At the low end of the narrativity spectrum, we have general laws and principles which define general statuses and are meant to transcend time and place. Such an approach could have important implications for the study of how people identify and understand stories. By identifying the characteristics of stories in linguistic terms, we should be able to isolate those processes which are used to understand stories and the narrativity which inheres in a wide variety of texts that do not qualify as stories. The comprehension of stories thus becomes a sub-set of a wider problem. Similarly, it should be possible to teach a computer to recognize stories and to determine the relative narrativity of a text on the basis of these criteria, and perhaps, to interpret the text accordingly as well.

3   Narrativity in the Mishnah

In my book, *Stories of the Law: Narrative Discourse and the Construction of Authority in the Mishnah*, I use the above definitions of narrative, story and narrativity to chart the range of forms used by the Mishnah, the early third century CE rabbinic legal text, to present legal rulings and principles. This Mishnah is distinguished by the way it intermixes forms of varying levels of narrativity within a single passage. I argue that this has significant implications for the way in which the Mishnah conceptualizes law and jurisprudence. To give the reader a sense of how my concept of narrativity can be applied to an individual text, I will present and give examples of the basic categories of mishnaic formulations in order from least to greatest narrativity. Many of these examples are quite technical and space does not allow for a full explications of the legal concepts that stand stand behind them. I believe however, that these text's basic lingiustic forms and their significance will still be accessible to the general reader. Mishnaic formulations break down into two basic categories, irrealis texts and realis texts. Irrealis texts are those that present hypothetical situations or actions. Since stories must be realis texts, which refer to an actual event in the past, irrealis texts are inherently limited in their narrativity. At most they can be narratives, representing a hypothetical sequence of actions. This category in turn divides into two subcategories: apodictic and casuistic formulations. Apodictic formulations state the law in an absolute manner, such as: "It is prohibited to do X" or "Y must be done." They generally contain only a single verb, and hence are generally not narratives. Their exact level of narrativity depends on a variety of factors, primarily the specificity and dynamism reflected in the verb forms used in the individual statement. A few examples of apodictic formulations from the tractate Shabbat, dealing with the laws of the Sabbath include:

1. The standard of one who bleaches, hackles, dyes, or spins [wool] is a full double sit (*Shabbat* 13:4).
2. Any knot that is not permanent entails no culpability (*Shabbat* 15:2).
3. [We] may tie a bucket [over a well] with a strap (*Shabbat* 15:2).
4. Aristocrats may anoint their wounds with rose oil (*Shabbat* 14:4).

Note that the first two examples state general principles and do not describe any sort of action, while the second two declare the permissibility of a specific action. Since they describe a hypothetical event, the latter two examples possess a higher level of narrativity.

The other primary form of irrealis formulation is the casuistic statement. These are "if ... then ..." statements that establish the law in a given situation. By definition they consist of two parts, the description of the case and the ruling. These two parts almost always constitute two interconnected events and are therefore narratives. Once again, the exact level of narrativity will depend on the verb forms used and other factors. Examples, once again from the laws of the Sabbath include:

1. [If] a fire broke out on a Sabbath night / food for three meals may be saved (*Shabbat* 16:2).
2. [If a gentile] made a stairway [on the Sabbath to descend by it [from a ship] / an Israelite may descend after him (*Shabbat* 16:8).

These texts can easily be transformed into stories by making a few changes in tense and mood.

Realis texts possess an inherently high degree of specificity since they refer to a specific event in the past. Not all realis texts, however, refer to a onetime event or events. The Mishnah sometimes presents individual events that occurred repeatedly such as,

R. Eleazar ben Azariah's cow used to go out with a strap between its horns [on the Sabbath] (*Shabbat* 5:4).

More frequently, the Mishnah will present a series of events that were repeatedly enacted in sequence. I call these "ritual narratives" because they generally portray cultic procedures in the Jerusalem Temple. Such texts are often indistinguishable from actual stories on the grammatical level. It is only context that allows the reader to determine that the Mishnah is presenting events that took place repeatedly and not a series of one-time events.

Sometimes, the Mishnah presents a single one time event such as,

It once happened that R. Gamliel said to his servant Tevi, "Go out and roast us the Passover offering on the grill" (*Pesahim* 7:2).

These texts have a high level of specificity but lack significant dynamism since they do not portray a chain of interrelated events.

Finally, we have full-fledged stories. Mishnaic stories portray the rabbis as issuing rulings and legal enactments or establishing precedents through their own public behavior. Examples include:

1. It once happened that R. Gamliel and the elders were traveling on a ship, when a gentile made a stairway for going down, and R. Gamliel and the elders descended by it (*Shabbat* 16:8).
2. It happened that the people of Tiberias placed a cold water pipe into a channel of hot water. The sages said to them: "On Shabbat, water heated thus is like any other water heated on Shabbat—it is forbidden to use it for washing or drinking. On festivals, it is like any other water heated on festivals—it is forbidden to use it for washing but permitted for drinking" (*Shabbat* 3:4).
3. Originally, they received testimony of the new moon from anyone. When the sectarians became corrupted, it was ordained that testimony should be received only from persons known [to the court] (*Rosh Hashannah* 2:1).

Note that the first story presents a precedent set by the actions of a great rabbi, the second presents a ruling regarding a specific case, while the third tells of a decree issued by the rabbis in response to a particular historical circumstance. In all of these cases, stories embed the law in specific human and historical contexts and ground it in the individual judgment of specific rabbis. This contrasts sharply with laws presented using low narratvity forms, especially apodictic formulations. These abstract forms present the law as being rooted in timeless principles rather than contingent circumstances and individual judgment.

Throughout the Mishnah these different types of formulations are juxtaposed and interwoven, creating an environment of constantly shifting narrativity and internal dialog between forms. Stories and narratives in the Mishnah are part of a larger discourse which contains a full range of narrativity. They can only be fully understood within this wider context.

4 Narrative Structure

One thing the above framework does not account for is the role of structure in narrative. The question of narrative structure has been central to the endeavor of narrative theory going back to the early twentieth century (see especially [19]). Traditional narratologists have generally viewed structure as a fixed attribute of the narrative text. In fact, virtually any text can be described using a variety of structural models.[10] Narrative structures are artificial devices that we use, consciously or unconsciously, in order to interpret stories. They are a way of retelling the story, focusing on those aspects of the story on which the interpreter sees as most important. The purpose of abstracting a structure from a narrative is to establish a basis of comparison with other stories. Comparing and contrasting stories with each other is perhaps the primary way through which we interpret them. Establishing a narrative structure is a process of removing the specificity from a text to a certain degree, leaving behind a series of interrelated events or sets of events described in a more general manner. By reducing the specificity of a narrative, we potentially reveal its wider significance or meaning. Take for example the following story:

> I was driving on Highway 1 from Jerusalem to Tel Aviv at 160 kilometers per hour. I was stopped by the police and given an 800 shekel ticket and had my license suspended for three months.

This story can be reduced to the following structure: (1) driving from Jerusalem to Tel Aviv, (2) speeding, (3) being stopped by police, and (4) receiving a large fine and a suspended license. This structure establishes a causal relationship between speeding on the Jerusalem–Tel Aviv highway and being stopped by the police and punished. The moral of this story can thus be expressed as: "Don't speed on Highway 1 in Israel; you risk being pulled over, getting a hefty ticket, and having your license suspended." This story could also be translated into a less specific, more abstract structure such as: (1) driving well above the speed limit, (2) being stopped by police, (3) receiving a ticket. Reading the story through this structure, we receive a much broader lesson that applies far beyond a single highway in a particular country: "If you speed, you may be pulled over and fined." Finally, we can remove all specificity from the story with the following structure: (1) violating the law, (2) getting caught, (3) being punished severely. Now the moral of the story is even broader: "Crime doesn't pay." As it stands, this anecdote carries all of the meanings listed above and more. However, when individuals tell anecdotes such as this in the course of a conversation or some

---

[10] On the priority of text over structure in the study of narrative see, [7, pp. 27–37]. See also [5, pp. 3–36].

other discourse, they usually have one meaning in mind as the primary one which they seek to communicate. Identifying which meaning is relevant is crucial to properly understanding a story in its context. The context tells the listener or reader which narrative structure to impose on the text. Thus if the context is "driving in Israel" the first structure would be appropriate. The second structure might be activated if the story were found in the midst of a discussion of speeding or traffic tickets in general and so on. It could also be that this story is not meant to make a more general point but rather serves to explain a state of affairs, such as the fact that the teller is in a bad mood. In such a case the interpreter would need to understand that this narrative is incomplete. The fact that the teller is in a bad mood, which was revealed before the story was told, is in fact is the final event of the story. By configuring the story in this manner the interpreter comes to understand that there is a causal link between the events described in the story and the teller's mood. In the case of the anecdotes in the Mishnah which I studied, things were a little more complicated. Read in their immediate context these stories serve to teach a particular legal ruling or principle. Yet, when we consider these stories in the context of each other, they express a different message based on a different narrative structure. These stories all portray rabbis as teaching and handing down authoritative rulings. Viewed in this light, the stories collectively give a message that the rabbis are the authoritative transmitters and interpreters of the law. This illustrates the way in which applying multiple narrative structures to a single story can be neccesary to gain a full appreciation of the place of a story in its wider context.[11] On the whole, stories told in the context of a larger text, speech or conversation are deployed to make a specific point. Understanding that point by abstracting the proper structure from the story is crucial to competence in understanding narrative. It seems to me that it might be valuable for researchers seeking to understand or model the process of narrative comprehension to focus on anecdotes as they appear in the context of everyday speech. This would clarify what is meant by "understanding" a story. In the case of anecdotes, first and foremost, narrative comprehension means understanding the role of the story in the wider discourse in which it appears.

5   Conclusion

In sum, I would like to suggest the following postulates about stories:

1. The categories of "story" and "narrative" are essentially artificial. All texts exist along a continuum of narrativity.
2. Narratives and narrativity cannot be understood by focusing solely on the way in which they portray change. The fact that stories focus on the specific, on individuals and events bound by time and place is equally important to the concepts of narrative and narativity.
3. Stories very frequently appear as part of larger discourses. Understanding how stories function in these contexts is an import part of understanding how stories work.
4. Stories are inherently polysemous. They can be reduced to multiple different structures of meaning. Context provides the cues for knowing which structure was primarily intended by the storyteller.

---

[11] The issue of context has played an important role in post-classical narratology and social scientific study of narrative. For a survey of some of the literature, see [16].

------- **References** -------

1 Roland Barthes. *Writing Degree Zero.* Hill and Wang, New York, 1968. transl. A. Lavers and C. Smith.

2 Roland Barthes. The reality effect. In Tzvetan Todorov, editor, *French Literary Theory Today: A Reader*, pages 173–204, Cambridge, UK, 1982. Cambridge University Press.

3 Emile Benveniste. *The Correlation of Tense in the French Verb.* Number 8 in Miami Linguistics Series. University of Miami, Coral Gables, FL, 1971. transl. M. E. Meek.

4 Guyora Binder and Robert Weisberg. *Narrative Criticism of Law.* Princeton University Press, Princeton, NJ, 2000.

5 Peter Brooks. *Reading for the Plot: Design and Intension in Narrative.* Knopf, New York, 1984.

6 Jerome Bruner. *Making Stories: Law, Literature, Life.* Harvard University Press, Cambridge, MA, 2002.

7 Jonathan Culler. Fabula and sjuzhet in the analysis of narrative: Some American discussions. *Poetics Today*, 1(3):175–191, 1980.

8 Suzanne Fleischman. *Tense and Narrativity: From Medieval Performance to Modern Fiction.* University of Texas Press, Austin, 1990.

9 Monika Fludernik. *Towards a 'Natural' Narratology.* Routledge, London, 1996.

10 Edward Morgan Forster. *Aspects of the Novel.* Edward Arnold and Company, London, 1927.

11 Gaballah Ali Gaballah. *Narrative in Egyptian Art.* Verlag Philipp von Zabern, Mainz, 1976.

12 Stephen Halliwell, editor. *Aristotle. Poetics.* Loeb Classical Library. Harvard University Press, Cambridge, MA, 1995.

13 Barbara Herrnstein Smith. Narrative versions, narrative theories. In William J. Thomas Mitchell, editor, *On Narrative*, pages 173–204, Chicago, 1980. University of Chicago Press.

14 Roman Jakobson. On realism in art. In Krystyna Pomorska and Stephen Rudy, editors, *Language in Literature*, pages 173–204, Cambridge, MA, 1987. Belknap Press.

15 William Labov. Some further steps in narrative analysis. *Journal of Narrative and Life History*, 7:175–191, 1997.

16 Ansgar Nünning. Surveying contextualist and cultural narratologies: Towards an outline of approaches, concepts and potentials. In Sandra Heinen and Roy Sommer, editors, *Narratology in the Age of Cross-Disciplinary Narrative Research*, pages 173–204, Berlin, 2009. Walter de Gruyter.

17 Gerald Prince. *A Grammar of Stories.* Mouton, The Hague, 1973.

18 Gerald Prince. Narrativity. In Karl Menges and Daniel Rancour-Laferriere, editors, *Axia: Davis Symposium on Literary Evaluation*, pages 173–204, Stuttgart, 1981. Akademischer Verlag Hans-Dieter Heinz.

19 Vladimir Yakovlevich Propp. *Morphology of the Folktale.* University of Texas Press, Austin, TX, 2nd edition, 1968. transl. L. Scott.

20 Paul Ricœur. *Time and Narrative Volume 2.* University of Chicago Press, Chicago, 1985. transl. K. McLaughlin and D. Pellauer.

21 Moshe Simon-Shoshan. *Stories of the Law: Narrative Discourse and the Construction of Authority in the Mishnah.* Oxford University Press, New York, 2012.

22 Wendy Steiner. *Pictures of Romance: Form Against Content in Painting and Literature.* University of Chicago Press, Chicago, 1988.

Social Narrative Adaptation using Crowdsourcing*

Sigal Sina[1], Avi Rosenfeld[2], and Sarit Kraus[1]

1 Department of Computer Science
   Bar-Ilan University
   Ramat-Gan, Israel
   `sinasi@macs.biu.ac.il, sarit@cs.biu.ac.il`
2 Department of Industrial Engineering
   Jerusalem College of Technology
   Jerusalem, Israel
   `rosenfa@jct.ac.il`

──── **Abstract** ────

In this paper we present SNACS, a novel method for creating **S**ocial **N**arratives that can be **A**dapted using information from **C**rowd**s**ourcing. Previous methods for automatic narrative generation require that the primary author explicitly detail nearly all parts of the story, including details about the narrative. This is also the case for narratives within computer games, educational tools and Embodied Conversational Agents (ECA). While such narratives are well written, they clearly require significant time and cost overheads. SNACS is a hybrid narrative generation method that merges partially formed preexisting narratives with new input from crowdsourcing techniques. We compared the automatically generated narratives with those that were created solely by people, and with those that were generated semi-automatically by a state-of-the-art narrative planner. We empirically found that SNACS was effective as people found narratives generated by SNACS to be as realistic and consistent as those manually created by the people or the narrative planner. Yet, the automatically generated narratives were created with much lower time overheads and were significantly more diversified, making them more suitable for many applications.

1 Introduction

Narratives correlate and relate events in our world and represent an important part of human experience. Family members and friends share their experiences via stories. Teachers and leaders tell stories to convey ideas while entertainers tell stories for fun. There is a growing need for conversational agents to understand and validate narratives in settings as diverse as military training, interactive computer games, elderly companion agents and educational, pedagogical and tutoring agents [14, 16, 19, 20, 23, 28, 30].

Due to the significance of this problem, many studies have analyzed storytelling, automatic story generation and interactive narratives. Early works considered a whole story

creation without getting any real-time input from the user [16, 28], while later works focus on interactive narratives where the user is part of the story and can affect the plotline [5, 20, 21, 23, 30]. Some previous works use hybrid methods whereby a predefined plot is created and an autonomic agent can later add to or adapt the plot in real-time [19, 22]. However, these studies focus on the general plot of the story but not on the story's details, which were almost exclusively manually created by content experts.

Our goal is to generate realistic narratives as easily and quickly as possible, but with varied content. The narratives we create must be of sufficient quality that they be believable, and thus be useful. Specifically, we present SNACS, an algorithm for generating everyday, social narratives that is customized for a specific storyteller and narrative type. The goal is for SNACS to generate everyday, social narratives that are (a) reasonable (that match the storyteller profile and the input constraints and limitations), (b) consistent (does not contain any contradictions) and (c) realistic (does not raise doubts about the credibility of the storyteller and is compatible with common knowledge implications).

As is the case with several recent works [11, 13, 19, 22], SNACS' algorithm implements a hybrid method which constitutes a "fill and adjust" semi-automatic narrative generation method. However, unique to our work, we add content as follows: We start by generating a dataset of daily, social narratives from Amazon Mechanical Turk workers using a semi-structured form. The dataset contains a collection of these acquired narratives written in natural language along with key attributes from within the narrative and demographic data of the worker. SNACS is a novel algorithm that decides how to use this information dataset in order to output a personalized narrative by adjusting its content according to the requested needs of a specific situation. This paper focuses on describing how the SNACS algorithm operates.

We tested SNACS on four types of social narratives, where the first two types were related to entertainment activities – going out to see a movie or eating at a restaurant, and the other two types were related to errand activities – buying groceries or going to the dry cleaners. We present results that show that SNACS produces narratives which are as reliable, consistent and realistic as the original narratives and a previously defined planning-based approach [17, 18]. Yet our narratives were easier to generate and were judged to be significantly more diversified than the planning-based approach.

2   SNACS Overview, Definitions and Algorithm

We propose and build a system which generates new everyday social narratives using previously stored narratives collected using crowdsourcing techniques. To help clarify and illustrate SNACS' definitions and the algorithm flow, please refer to motivating example 1, a narrative which was generated by SNACS and example 2, the narrative upon which SNACS was based. In example 1, we wish to provide a believable original narrative for a 31-year-old female, who is single and has no children. SNACS generated this narrative based on the narrative in example 2 which was written by a 26-year-old female, who is married and has one child.

▶ **Example 1.** "My sister, my nephew and I went to eat at a restaurant on Friday evening. We went to the nearby city and ate at "Maggiano's Little Italy". This restaurant is one of our favorites so we go there often. Walking in, everything smelled so good and we were greeted at the door and promptly seated since we had called ahead of time. Our waiter was nice and we ordered a nice glass of wine and juice for my nephew. We ordered bruschetta and ordered our main course pasta dishes. My nephew had some mac and cheese. It was very relaxing and the food was amazing. I love going to "Maggiano's Little Italy"."

▶ **Example 2.** "My husband and I went to Olive Garden with our son. This was this past weekend. We went for dinner on Saturday night. We didn't want to cook and wanted to get out of the house. Walking in, everything smelled so good and we were greeted at the door and promptly seated since we had called ahead of time. Our waiter was nice and we ordered a nice glass of wine and juice for my son. We ordered bruschetta and ordered our main course pasta dishes. My son had some mac and cheese. It was very relaxing and the food was amazing. I love going to Olive Garden."

In this section, we describe how this and other narratives can be generated. Note that while both narratives are similar in some ways (e.g., both are narratives about going to restaurants, the lengths of the stories are similar, and both children had mac and cheese), several key differences exist between stories (e.g., one story refers to a person's son, the names of the restaurants are different, and were visited at different times). To explain the process by which some details are retained and others are tailored to the new situation, we begin by providing definitions for SNACS' algorithm. We then describe how information provided through crowdsourcing is used to create a semi-natural language dataset. Last we describe how this dataset can be used to generate a believable narrative to a new storyteller, given her demographic properties.

2.1 Definitions

Throughout this paper, we use the following terms to describe the algorithms that we developed. We notate **original** to refer to elements from the source story upon which we base the new narrative. We notate **generated** for elements related to the narrative that we wish to create. We use examples 1 and 2 to illustrate the definitions and accordingly example 1 is the **generated** narrative while example 2 is the **original** narrative.

Narrative Dataset(DS) – is a collection of narrative records R for a specific narrative type, such as see-a-movie or eat-at-a-restaurant.
Example 2 is an instance of such a record within DS.

Narrative Record (R) – is a triple $\langle P, A, S \rangle$ where: P is the storyteller profile, A is the narrative attributes vector, S is the narrative natural language presentation.
A detailed description of the three fields in R immediately follows. The key to the SNACS' algorithm is how to best select the most appropriate original narrative record from within the entire set of DS. For that every record R is also associated with a tagging vector named ILV (described below), which is the base of the SNACS' selection mechanism. The original narrative record is then modified to create the generated narrative record.

Storyteller Profile (P) – describes the storyteller's properties and consists of gender, age, personal status and number of children.
In examples 1 and 2, the original storyteller profile was: Female, age 26, Married, one child, and the generated storyteller profile was: Female, age 31, Single, no children.

Narrative Attributes Vector (A) – contains a vector of attributes which accompany the narrative. It contains general attributes such as participants, a day, part of day, duration and location. This vector also contains information specific to the narrative domain. It can contain optional values, and thus can be full or partial. Examples of optional values in the restaurant narrative include: location, part of day and participants (e.g., a person can eat at a restaurant alone, but can also go with a spouse, friend or children).

For example, within examples 1 and 2 above, the restaurant narrative attributes are day, part of day, restaurant name, restaurant type, location and participants. Thus, we represent the original narrative attributes vector as: ⟨day (Saturday), part of day (night), restaurant name (Olive Garden), restaurant type (Italian American cuisine), location (downtown) and participants (spouse and son)⟩. The generated narrative attributes vector is: ⟨day (Friday), part of day (evening), restaurant name (Maggiano's Little Italy), restaurant type (Italian American cuisine), location (nearby city) and participants (sister and nephew)⟩.

Natural Language Narrative Presentation (S) – contains a detailed description of an everyday, social activity written in natural language. The natural language (NL) presentation is composed of three parts:

1. The narrative `introduction` – which describes the main facts of the activity, such as who went, when, what are the main object names (which movie/restaurant), where and why.
2. The narrative `body` – which describes the experience in detail, what was the course of events and what happened during the experience.
3. The experience `perception` – which describes how good or bad the experience was from the storyteller's perspective.

We intentionally split the social activity's detailed description into these three parts. This semi-structured natural language writing is very applicable when describing social, everyday situations and and it also centralizes most of the event specific details in the introduction part. This facilitates us to adjust the narrative to a new storyteller profile and attributes vector during the narratives generation.

The original narrative presentation in example 2 is composed of the following three parts:

1. The narrative introduction – "My husband and I went to Olive Garden with our son. This was this past weekend. We went for dinner on Saturday night. We didn't want to cook and wanted to get out of the house."
2. The narrative body – "Walking in, everything smelled so good and we were greeted at the door and promptly seated since we had called ahead of time. Our waiter was nice and we ordered a nice glass of wine and juice for my son. We ordered bruschetta and ordered our main course pasta dishes. My son had some mac and cheese."
3. The experience perception – "It was very relaxing and the food was amazing. I love going to Olive Garden."

Logical Constraints and Common-Sense implications (CS) – contains a set of handwritten, logical rules which the generated narrative should fulfill in order to be reasonable, consistent and realistic. For example:

(a) If the activity occurs late at night, then it shouldn't include small children.
(b) For the see-a-movie activity, the participants should include children if the selected movie type is a children's movie.
(c) The storyteller's profile should match the activity's participants.

Note that we assume that only the generated story need worry about CS, as we assume that the records within DS that were generated through crowdsourcing already fulfill these requirements.

Similarity Attributes Vector (SA) – is a vector of N selected attributes from the storyteller's profile and narrative attributes vector and is defined as SA =

$\langle SA_1, \ldots, SA_N \rangle$. SNACS uses this vector to assign values for attributes within a generated narrative as we will explain in this paper.

To generate Social Narratives, SNACS uses a vector of 7 ($N = 7$) known attributes values. These 7 attributes are: gender, age, number of children, personal status, participants, type and part of day. Note that these 7 attributes include all of the profile information attributes (P), but only a subset of the narrative attributes (A).

Similarity Level (SL) – we defined three similarity levels: **same**, **similar** and **other**. We coded these similarity levels as 1, 2 and 3 respectively.

Importance Level Vector (ILV) – a vector ILV contains N values, corresponds to the vector SA, and is defined as $ILV = \langle ILV_1, \ldots, ILV_N \rangle$. Each value in ILV is a value SL and is used to represent the importance of the compatibility of a given attribute SA_i within the narrative body and the experience perception parts of the narrative.

Within SNACS, every record R within the dataset DS has a vector ILV. These values control how much importance should be given to having similarity between the original and generated narratives. Accordingly, if a given attribute within a narrative R can be changed without violating any common sense implications (CS), then the value for SA_i is **other**. At the other extreme, if that attribute is critical and even small variations can make the story implausible, then the value for SA_i is **same**. As we will see in the next section, SNACS considers two ways in which the vector ILV can be built for every record – either a fixed value across all records within DS for a narrative type (which we will call `SNACS-Bst`) or utilizing a content expert to manually tag every record within DS (which we will call `SNACS-Tag`).

The fixed (automatic) ILV contains the **same** value for the gender attribute and a **similar** value for all of the other attributes, i.e., $\langle 1, 2, 2, 2, 2, 2, 2 \rangle$. For the narrative in example 2, the manual ILV is $\langle 3, 3, 3, 3, 2, 1, 2 \rangle$, i.e., $\langle$gender (other), age (other), number of children (other), personal status (other), participants (similar), type (same), part of day (similar)$\rangle$. Note that the type attribute got the **same** value (which means a specific detail) as the narrative contains the sentence "We ordered bruschetta and ordered our main course pasta dishes."

Matching Level Vector (MLV) – a vector MLV contains N values, corresponds to the vector SA, and is defined as $MLV = \langle MLV_1, \ldots, MLV_N \rangle$. Each value in MLV is a value SL and is used to represent the matching level between the original and generated values of a given attribute SA_i. Within SNACS, we built a comparison method for each attribute, which gets as input two values and returns one of the three similarity levels, SL. In case one of the values is missing, it returns the **similar** value, as we assume SNACS will fill this attribute with a similar value. As we will see in the next section, SNACS uses the vector MLV to select the best candidate narrative record R from the DS, in both the `SNACS-Bst` and `SNACS-Tag` variations of the algorithm.

For example, we consider the number-of-children attribute to be the **same** if the difference between the original profile and the generated profile is 1 or less, **similar** if the difference is less than or equal 3 and **other** if one person has children and the other does not or when the difference is greater than 3.

The MLV between examples 1 and 2 will be $\langle 1, 3, 3, 3, 2, 1, 2 \rangle$, i.e., $\langle$gender (same), age (similar), number of children (other), personal status (other), participants (similar), type (same), part of day (similar)$\rangle$.

Compatibility Measure (CM) – The novelty of SNACS is its ability to generate new narratives based on original, acquired narratives. SNACS uses this measure to select the best candidate for the given storyteller profile and the attributes vector.

The compatibility measure *CM* of a narrative record *R*, given the vector *ILV*, a storyteller profile *P* and a (partial) attributes vector *A*, is calculated as a weighted sum of scores that depends on the values of the given *ILV* and the calculated *MLV*. We describe this calculation later in this section.

The Problem Statement.   Given these definitions, we describe how to generate a narrative, *S*, written in natural language. Narrative *S* is built from a storyteller profile *P*, a (partial) vector of attributes *A* (e.g., time, location and participants) and a set of logical implications *CS* (data constraints and common knowledge implications). Narrative *S* must not only be written in natural language but must also be:

(a) reasonable (it matches the storyteller profile and the input attributes),

(b) consistent (does not contain any contradictions) and

(c) realistic (does not raise doubts about the credibility of the storyteller and is compatible with common knowledge implications).

2.2   The SNACS Algorithm

The key to SNACS' success is having a varied set of narrative records within DS from which it can generate new narratives. We used Amazon Mechanical Turk (`http://www.mturk.com/`), a crowdsourcing web service that coordinates the supply and demand of tasks which require human intelligence to create DS. Amazon Mechanical Turk (AMT) has become an important tool for running experiments with human subjects and was established as a viable method for data collection [1, 15]. Our previous experience in running experiments on AMT has demonstrated that this is an effective medium for collecting data about various tasks [2, 3].

We crowdsourced the creation of DS as follows: We built a dedicated, semi-structured questionnaire on AMT to collect the narrative records – the profile, P (gender, age, personal status and number of children), the narrative attributes vector, A (story attributes) and the narrative presentation in natural language S. In this questionnaire the AMT workers were first asked to provide their profiles. Then, the workers were asked to describe daily, social narratives in natural language in as much detail as possible, according to the the three narrative parts – introduction, body and perception. Lastly, the workers were presented with a list of specific questions used to collect the narrative attributes vector, such as "What was the name of the movie/restaurant?", "With whom did you go?" and, "on what day?". The completed record was then stored at the narratives dataset. As we receive new narratives as input from crowdsourcing, we store every record for a given type of narrative (e.g., going to a movie or eating at a restaurant) within the dataset DS_k for that story type. To demonstrate the generality of SNACS, we created four such datasets ($k = 4$) – see-a-movie, eat-at-a-restaurant, buying-groceries and dry-cleaning. We crowsourced 10 narrative records for each narrative type.

Once we wish to generate a new narrative for a listener (user) for a new storyteller profile, we must select the most appropriate narrative record (the original narrative) from DS given that storyteller's profile. We then generate the missing narrative attributes in the narrative attributes vector according to the new storyteller and the chosen record. Finally, we generate the narrative's natural language presentation for the given storyteller profile and the generated attributes vector based on the original chosen narrative's natural language presentation. The process for these steps is formally presented as a SNACS' algorithm and further described later in this section.

Algorithm 1 An Algorithm to create **S**ocial **N**arratives through **A**daptation of **C**rowd-**S**ourcing narratives **(SNACS)**

Require: Storyteller profile P and a (partial) narrative attributes vector A
Require: Logical common knowledge constraints CS
Require: Narrative dataset DS
Ensure: Reasonable, consistent and realistic narrative record R
  1: Select original candidate narrative record OR from dataset DS based on P and A
  2: Create a new narrative attributes vector NA and complete it according to P, A, CS and OR
  3: Create a new narrative record NR from OR, P and NA
  4: Replace the original content of the introduction (first part) in NR with a new generated introduction based on P and NA
  5: Adjust the second and third parts (the body and perception) in NR
  6: **return** The updated NR

The input and output. SNACS gets as input from the listener (user) a storyteller profile P and a vector of optional attributes A (such as participants, or the date). It returns as output a new narrative record NR which contains a reasonable, consistent and realistic narrative presentation S written in natural language.

The algorithm's logical process. The algorithm first selects a candidate narrative record from the dataset (line 1). We present three variations of this selection process below. Then, SNACS completes the missing narrative attributes (line 2). SNACS generates attributes which are similar to the selected, original narrative attributes and matches them to the new storyteller profile. It starts with the participant's generation, who went and how many people participated in the event. It generates the objects' names (movie, restaurant, location) and time frame attributes. For example, if in the original narrative someone went to see a children's movie with his daughter and the new storyteller has no children, the algorithm can choose to include his niece/nephew among the participants. Next, the algorithm generates the narrative's natural language presentation. First, it replaces the original narrative introduction (line 4), i.e., its first part (who went, when, where, why), with a newly generated introduction according to the new profile and the new vector of attributes. This is done by using several predefined Natural Language Generation (NLG) templates with parameters as we describe later in this section.

Finally (line 5), SNACS applies some adjustments to the body and perception parts of the narrative's natural language presentation (the second and third parts). This is done by replacing the references of the original attributes' vector to the new corresponding narrative attributes' vector. In our example, the original participants were a husband and son where the generated participants are sister and nephew; and the original restaurant's name was "Olive Garden" where the generated restaurant's name is "Maggiano's Little Italy". The SNACS algorithm will replace the reference of **My son** with **My nephew** and the reference of **Olive Garden** with **Maggiano's Little Italy** in the following narrative body and perception: "...We ordered bruschetta and ordered our main course pasta dishes. **My son** had some mac and cheese. It was very relaxing and the food was amazing. I love going to **Olive Garden**."

Original narrative selection.   We implemented three variations of the SNACS algorithm which differ in how the original candidate narrative record is chosen (at line 1): `SNACS-Any`, `SNACS-Bst`, `SNACS-Tag`.

The `SNACS-Any` variation is a baseline measure that randomly chooses one narrative record from DS. No further logic is performed to check how appropriate that choice is. In contrast, both the `SNACS-Bst` and `SNACS-Tag` variations use a compatibility measure (CM) previously defined. In both of these variations, this measure is used to select which candidate from among all records in DS will serve as the base for the generated narrative. As such, the record with the highest CM value is chosen as the original record.

However, these two variations differ as to how the importance level vector (ILV) is generated. Within the `SNACS-Bst` variation, the algorithm uses the a fixed (automatic) ILV $\langle 1, 2, 2, 2, 2, 2, 2 \rangle$ for the CM calculation. In contrast, in the `SNACS-Tag` variation, a content expert manually tags every record within DS to create the vector ILV of the importance of every attributes. Referring back to example 2, the expert tagging generates an ILV of $\langle 1, 3, 3, 3, 2, 1, 2 \rangle$. Note that the `SNACS-Tag` variation is much more time consuming than `SNACS-Bst`. As our results will show in Section 4, `SNACS-Bst` and `SNACS-tag` often choose the same record to serve as the base of the narrative. This indicates that we can use the much simpler set values in `SNACS-Bst` and avoid the time of manually tagging the narratives' attributes. Within the `SNACS-Bst` and `SNACS-Tag` algorithms a compatibility measure (CM) is then used to select which record in DS will serve as the original narrative. We use 3×3 matrix scores for all the possible combinations of ILV_i and MLV_i values. For each narrative record R, the algorithm first calculates the MLV for the given storyteller's profile P and input attributes vector A. It then calculates the record's compatibility measure CM as a weighted sum of the corresponding score from the 3×3 matrix based on the ILV (fixed or manual) and the calculated MLV values. Finally, it returns the candidate with the highest compatibility measure for the requested storyteller's profile and the generated narrative attributes vector.

The narrative introduction generation.   We generate the narrative introduction (at line 4) using SimpleNLP [9], a Natural Language Generation (NLG), template-based surface realization. Realization is a subtask of NLG, which involves creating an actual text in a human language from a syntactic representation. SimpleNLG is an NLG system that allows the user to specify a sentence by providing its content words and its grammatical roles (such as the subject or verb). SimpleNLG is implemented as a Java library and it allows the user to define flexible templates by using programming variables in the sentence specification. The variable parts of the templates could be filled in with different values. We had a few NLG templates for each narrative type, which were randomly chosen during the introduction generation. For example, two of the NLG templates we use to build a narrative introduction for the see-a-movie narrative are:

(a) Last $\langle time \rangle$ I went to a movie with my $\langle with \rangle$. We went to see the movie $\langle movie \rangle$ at $\langle theater \rangle$.

(b) My $\langle with \rangle$ and I went to see the movie $\langle movie \rangle$ on $\langle time \rangle$. My $\langle with \rangle$ had seen the trailer and wanted to see this movie ever since.

Each template can generate a few variations according to the chosen attributes. For example, the first part of above template (a): "Last $\langle time \rangle$ I went to a movie with my $\langle with \rangle$" where the participants are a wife and son and the time is Sunday afternoon, can generate a few variations, such as:

(a) Last Sunday I went to a movie with my family.

(b) Last weekend I went to a movie with my family.

(c) Last Sunday afternoon I went to a movie with my wife and my son.

3   Comparison with a State-of-the-Art Generator

In order to validate the significance of SNACS, we implemented a planning-based algorithm narrative generator for the see-a-movie and eat-at-a-restaurant narratives. One of the most common approaches for narrative generation is planning-based story generation systems [16, 23, 21, 30, 5]. The planning-based approach uses a causality-driven search to link a series of primitive actions in order to achieve a goal or to perform a task. For the domain of narratives of every day activities, hierarchical scripts can capture common ways to perform the activity. Therefore, we implemented the planner-based generator using a Hierarchical Task Network (HTN). HTN is one of the best-known approaches for modeling expressive planning knowledge for complex environments. It is a natural representation for domains where high-level tasks are decomposed into simpler tasks until a sequence of primitive actions solving the high-level tasks is generated. We used the state-of-the-art SHOP2 planner [18], a well known HTN planner, which has been evaluated and integrated in many real world planning applications and domains including: evacuation planning, fighting forest fires, evaluation of enemy threats and manufacturing processes [17].

Several key similarities and differences exist between generating narratives with SNACS and using any planner, such as SHOP2. First, in both implementations, we assume that the storyteller's profile (P) and a partial vector for A are given. Additionally, we assume that along with P and A, a set of logical constraints (CS) exists that constrains how any narrative can be built, and that Natural Language Generation (NLG) templates will assist with creating the narratives' introduction. However, key differences exist between the Planner and SNACS regrading how the planner will choose the missing narrative attributes and descriptions within the body and the perception of the narrative as we now detail.

Specifying the Planner Input. The HTN planner domain is built according to a **plot graph** of basic actions. A plot graph [31] is a script-like structure, a partial ordering of basic actions/events that defines a space of possible action/event sequences that can unfold during a given situation. For example, the basic actions for see-a-movie include: travel to theater, choose a movie, buy tickets, buy snacks, find seats, see the movie and talk about movie. In our implementation of SHOP2, we manually built the graph plots for two narrative types – see-a-movie and eat-at-a-restaurant. As is the case in SNACS, the planner algorithm starts by filling in the attributes vector (A) according to the logical constraints (described below) and then it searches for a valid plan for the given storyteller profile (P) and the narrative attributes vector (A). As we wanted to get a richer narrative which includes a detailed description of an everyday, social activity written in natural language, we gave the planner an option to tailor natural language descriptions in the basic actions portion of the narrative, as we describe below. Note that in the SNACS algorithm, this step was not necessary as we automatically got the narrative's detailed descriptions from the original narrative. For example, for the "Buy Snacks" action within the see-a-movie narrative type, we gave the planner an option to choose one of the following descriptions:

- We smelled the popcorn and went to buy some.
- We went to get popcorn but it was too expensive.
- We got soda and popcorn from the concession stand.
- We bought some popcorn and drinks and were annoyed at how expensive it was.
- No matter what the price, I just can't see a movie without a big tub of popcorn.
- Snacks at a movie are a complete waste of money! I just bring a snack from home and save the money.

- We decided to buy a large soda to stay cool because the air conditioning wasn't working in the movie theater.
- Though I am supposed to be on a diet, I just couldn't resist buying a few chocolates to eat during the movie.

We defined for each basic action, we defined a set number, 10–15, of different descriptions that were tailored to the specific narrative. These descriptions were manually handwritten by two experts. These experts needed approximately one hour (half an hour per expert) to write the set of descriptions for each basic action option. The experts' data insured that the implemented SHOP2 planner had a variety of descriptions with which to build narratives. Part of these descriptions were also manually tagged by the system's designer with specific tags, such as movie or restaurant types. Note that this tagging is part of the manually supported process needed by this approach in addition to the time spent by the content experts. This tagging gave the generator an option to choose between a generic description which can be associated with any movie/restaurant type or a specific description which can be associated with the current selected movie/restaurant type. For example, for the "Talk about Movie" action we used both generic descriptions, such as:

- It was a nice movie though the end was a bit disappointing.
- I enjoyed the movie though there was a lot of noise from the back seats during the show.
- It was a refreshing and interesting movie even though it was too long for my taste.
- The movie wasn't bad. While I have seen better movies, it was overall a nice way to spend an evening.
- The movie was very nice. I arrived not expecting much and really enjoyed it.

and specific descriptions per movie type, such as

- It was the most stunning film, filled with action and spectacular effects. (action movie)
- I enjoyed the movie. A nice story and very moving. The film was very dynamic, constantly evolving and interesting. (dramatic movie)
- The movie was so funny I couldn't stop laughing. (comedy)
- I highly recommend this film to all ages; whether you are a kid or an adult you will find enjoyment in this film. (children's movie)

The HTN generator flow. The SHOP2-based narrative generator starts filling in the missing attributes vector of the narrative for the given storyteller profile. As the SHOP2-based narrative generator isn't based on an original narrative as in the SNACS algorithm, it uses random selection when there are no logical constraints. The planner implementation first selects the timeframe (when). Next, it chooses the participants, who went and how many people participated in the event, according to the given storyteller profile, the selected time frame and the logical constraints. For example, as in SNACS, if the selected time is night, the narrative planner will not choose children as participants. Lastly, it chooses the objects' names (movie, restaurant, location) and the activity reason (why). In this step, the planner also considers the logical constraints, such as don't choose a children's movie if the participants don't include a child. Once it has all of the activity attributes' values, it looks for a valid the plan according to the domain plot graph. When it has several options from which to choose, it uses a random tie-breaker to select one option, whether it is between different branches of the plot graph or between the basic action descriptions.

The HTN generator output. It is important to note that the output of this planning algorithm is not natural language. The planner algorithm returns as output a raw, semi-structured narrative plan which can been seen (partially) below:

```
(detail action p2 went-to-watch-a-movie)
(detail profile p2 (female 28 married 3))
(detail activity when (4 18 week 3))
(detail activity with kids)
(detail activity name puss-in-boots)
(detail activity loc regal-cinemas)
(detail activity why (my children saw the trailer ...))
(detail action p2 (bought-ticket at-cinema))
(detail action p2 buy-snacks)
(detail action-desc snacks (we bought some popcorn ...))
(detail action p2 found-seat)
(detail action-desc seat (we decided to sit in the first row ...))
(detail action p2 watched-the-movie)
(detail action-desc how (the kids enjoyed the ...))
```

In order to convert the semi-structured plan into a natural language narrative presentation, we implemented a dedicated realizator (SimpleNLP based) which receives this raw, semi-structured narrative plan as an input and returns the narrative presentation written in natural language. This realizator first generates the narrative introduction based on the chosen narrative attributes vector and the storyteller profile using the same NLG templates the SNACS generator used. Next, it generates the event activity descriptions for the body and perceptions parts. For this step, it uses dedicated templates, which we didn't need for the SNACS algorithm as we based these parts on the original narrative that we had in our dataset. These templates were manually predefined for each one of the activity actions. Each template generates the basic description of the action in natural language and combines this output with its' associate chosen description, if one exist in the plan. For instance, the following two lines of the plan: (1) (detail action p2 found-seat) and (2) (detail action-desc seat (we decided to sit in the first row . . .)) will generate the following natural language sentences: "We went in to find seats and watched the movie. We decided to sit in the first row because the theater was crowded and I didn't want anyone to block my view."

Overall, the HTN-based narrative generator has an inherently higher cost associated comparing to the SNACS algorithm for the following reasons: Both SNACS and the planner do have the steps of building of the narrative introduction templates and the implementation of the logical constraints. However, the planning-based algorithm implementation also required some additional manual steps. These steps include: the manual building of the plot graph; the writing, associating and tagging of several detailed descriptions for each basic action; and writing a specific realizator for each basic action. Each one of these steps requires both time and resources from a content expert or a system's designer. In fact, because of this cost overhead, we only used the SHOP2 planner in order to define two entertainment narrative types – going out to see a movie or eating at a restaurant. We intentionally did not implement the HTN-based narrative generator for the errand narrative types – buying groceries and dropping off or picking up dry cleaning. Nonetheless, the narratives produced by SNACS were as good as those developed by this costly process, as our results detail in the next section.

4 Experimental Setup and Results

The evaluation of the effectiveness of our narrative generation algorithm SNACS was based on several instances of feedback from Amazon Mechanical Turk participants. The people were presented with narratives generated by our algorithm, the time intensive HTN planning-based generator, the manually handwritten narratives, and the random baselines (describe below).

4.1 Experimental Setup

Specifically, our evaluation was as follows: For each narrative type, we generate 4 storyteller profiles. For each profile, we then generate narratives using both SNACS and the HTN algorithms. As was previously mentioned, we implemented the HTN-based narrative generator only for the entertainment activities – going out to see a movie or eating at a restaurant, and did not implement it for the errand activities – buying groceries and dropping off or picking up dry cleaning, due to the inherent high cost associated with this approach. We also randomly selected four narratives out of a pool of ten original narratives.

The AMT participants were presented with a questionnaire, in which they were asked to grade each of the narratives together with their associated storyteller profiles. The questionnaire contains grades with values from 1 (Least) to 6 (Most). We intentionally chose a scale with an even number of values as we didn't want to allow the users to choose a middle value. The questionnaire asked the participants to grade six aspects of the narratives:

- Authentic – Did you find the story authentic?
- Reasonable – Did you find the story reasonable?
- Profile – Does the story match the storyteller profile (gender, age, personal status, number of children)?
- Coherency – Did you find the story coherent? Does the story make sense overall?
- Fluency – What is the fluency level of the story?
- Grammar – What is the level of the grammar in each sentence?

Each questionnaire contains between 8–10 narratives and was fill out by 8–10 users, to ensure we had a least 30 independent grades per narrative type and narrative generation algorithm. In order to better understand the participants' responses, we also asked them to explain their choices in free text. This ensured that subjects answered truthfully as we were able to manually check these open questions before accepting a participant's grades. Additionally, as we estimate that completing a questionnaire takes 8–12 minutes, we filtered out questionnaires which were filled out within less than 4 minutes as we assumed that the responses were not valid.

Random baselines. We also implemented two random algorithms as baselines to ensure the validity of our experiments. The first random algorithm, to which we will refer as `Rnd-SNACS`, uses the SNACS generator infrastructure. The `Rnd-SNACS` algorithm randomly chooses one of the narratives in the collection and then it randomly chooses the participants, the time frame, location and objects' names. This version ignores the current storyteller profile, the original use profile, the original narrative attributes' vector and the logic constraints and implications. The second random algorithm, to which we will refer as `Rnd-Planner`, uses the planning-based generator. Here, we removed the logical constraints on the participants, time frame and movie or restaurant types from the domain and the problem instance of the original version and use instead random selections to fill in the attributes vector and to choose

■ **Table 1** Average Grades.

Algorithm	Movie			Restaurant		
	Mean	**Std**	**N**	**Mean**	**Std**	**N**
Original	4.75758	1.14040	33	4.68981	1.02365	36
Planner	4.52083	1.09639	32	4.25000	1.20780	32
Rnd-Planner	3.71905	1.52097	35	3.35784	1.45998	34
Rnd-SNACS	2.75238	0.90782	35	3.06481	1.17937	36
SNACS-Any	4.47475	1.06484	33	4.30303	0.98817	33
SNACS-Bst	4.42857	1.14567	35	4.52688	1.08277	31
SNACS-Tag	4.43750	0.9482	32	4.46774	1.11589	31

■ **Table 2** Errands Average Grades.

Algorithm	Buy Groceries			Dry Cleaning		
	Mean	**Std**	**N**	**Mean**	**Std**	**N**
Original	4.09375	1.25291	32	4.51010	1.23106	33
Rnd-SNACS	3.59896	1.35391	32	3.74479	1.28263	32
SNACS-Any	3.77451	1.39315	34	4.63333	1.21725	35
SNACS-Bst	4.37879	1.45492	33	4.75269	0.94953	31
SNACS-Tag	4.83333	1.00179	32	4.70707	1.15096	33

the detailed descriptions. As the logical constraints and Common-Sense implications (CS) part was removed from these algorithms, we posit that these random baselines will generate poor stories. For instance, within the eat-at-a-restaurant narrative type, the algorithm `Rnd-SNACS` generated a contradiction regarding the time: "On Thursday **night**. . ." and then later in the story "We went in when the doors opened at **11:00AM** . . . ". Similarly, in a different generated narrative the `Rnd-Planner` algorithm generated a contradiction regarding how crowded the restaurant was: "We couldn't find a seat and **had to wait** for more than fifteen minutes" yet later in the same story this algorithm generated, "I liked that fact that it was kind of **empty**".

4.2 User Feedback Results

Table 1 shows that our novel SNACS algorithm generated revised story events which were rated by the AMT participants as being as consistent, believable and realistic as the original narratives and those produced by time-intensive planning technique. This is done without the costly overhead involved with manually creating narratives or using a planner-based approach.

Table 1 presents the average grade that AMT participants gave for the entertainment narratives. It is clear that the both of the random variations `Rnd-SNACS` and `Rnd-Planner` got lower grades, between 2.75 and 3.72, which are significantly lower than the generation algorithms and the original narratives' grades (specifically, the ANOVA test for the movie narratives of `Rnd-Planner` compared to `original`, `Planner` and `SNACS-Bst` had a much smaller than 0.05 threshold level with $p = 1.95 \cdot 10^{-4}$, $4.99 \cdot 10^{-3}$ and $9.30 \cdot 10^{-3}$ respectively). These results also show that all non-random generation methods got very similar grades

(4.43–4.47) for the movie narratives and that the `SNACS-Bst` and `SNACS-Tag` grades (4.46–4.52) are slightly better for the restaurant narratives than `SNACS-Any` method, which got an average grade of 4.30. Comparing to the SNACS based generator, the `Planner` generator got slightly better grades for the movie narratives (4.52) and slightly worse grades for the restaurant narratives (4.25). Although the `original` grades, 4.76 and 4.39, for the movie and restaurant narratives are slightly higher than all of the other methods, overall there is no significant difference between all of the SNACS-based algorithms or the planning-based generator or the original narratives.

Table 2 shows the average grades of the errands narratives, Again, the results show that the random variation `Rnd-SNACS` got lower grades for both the buy groceries and dry cleaning activities. For the buy groceries activity, `SNACS-Tag` got a best grade of 4.83 while `SNACS-Bst` got a best grade of 4.75 for the dry cleaning activity. In both cases, there is no significant difference between the grades of `SNACS-Tag`, `SNACS-Bst` and the `original` narratives.

We also tested each grade separately and the results were very similar and can be found in Table 3 and Table 4. Overall, for all four narrative types, there is no significant difference between `SNACS-Tag`, `SNACS-Bst` and the `original` narratives.

Table 3 Stories Aspects Grades.

Narrative Type	Algorithm	AvgGrade	Authentic	Reasonable	Profile	Coherent	Fluency	Grammar
Movie	Original	4.75758	4.97	4.94	5.27	4.88	4.48	4.00
	Planner	4.52083	4.56	4.56	4.59	4.56	4.69	4.16
	Rnd-Planner	3.71905	3.00	3.09	4.09	3.51	4.06	4.57
	Rnd-SNACS	2.75238	2.20	2.40	1.66	3.14	3.37	3.74
	SNACS-Any	4.47475	4.39	4.64	4.36	4.64	4.55	4.27
	SNACS-Bst	4.42857	4.51	4.57	4.14	4.63	4.57	4.14
	SNACS-Tag	4.43750	4.88	5.09	3.84	4.56	4.06	4.19
Restaurant	Original	4.68981	4.72	4.89	4.94	4.53	4.75	4.31
	Planner	4.25000	4.09	4.09	4.06	4.19	4.62	4.44
	Rnd-Planner	3.35784	3.03	3.12	3.62	3.00	3.59	3.79
	Rnd-SNACS	3.06481	2.44	2.42	3.22	3.47	3.44	3.39
	SNACS-Any	4.30303	4.39	4.45	4.42	4.18	4.18	4.18
	SNACS-Bst	4.52688	4.71	4.77	4.48	4.74	4.35	4.10
	SNACS-Tag	4.46744	4.61	4.94	4.65	4.45	4.10	4.06

Table 4 Errands Aspects Grades.

Narrative Type	Algorithm	AvgGrade	Authentic	Reasonable	Profile	Coherent	Fluency	Grammar
Buy Groceries	Original	4.09375	4.06	4.25	4.47	4.22	3.75	3.81
	Rnd-SNACS	3.59896	3.84	3.56	3.47	3.66	3.53	3.53
	SNACS-Any	3.77451	4.03	3.97	3.47	4.03	3.79	3.35
	SNACS-Bst	4.37879	4.64	4.48	4.58	4.30	4.21	4.06
	SNACS-Tag	4.83333	4.87	4.91	4.91	4.84	4.69	4.78
Dry Cleaning	Original	4.51010	4.61	4.64	4.52	4.58	4.39	4.33
	Rnd-SNACS	3.74479	3.63	3.50	3.78	3.84	4.00	3.72
	SNACS-Any	4.63333	4.51	4.49	4.63	4.80	4.57	4.80
	SNACS-Bst	4.75269	5.03	4.94	4.84	4.84	4.58	4.29
	SNACS-Tag	4.70707	4.58	4.61	4.88	4.91	4.58	4.70

4.3 Lexical Variability Results

We assume that different people have different communication types. As a result, stories for different groups will need to use fundamentally different types of language. For example, if we consider a narrative for elderly women, we want the content to focus on activities with grandchildren or a daily exercise routine, whereas if we consider a narrative for a small boy who is bedridden for medical reasons we might instead focus on activities with parents, school or his friends. Besides the essential purpose of generating a good and realistic narrative, we wanted a method which can generate a variety of narratives for different profiles with high lexical variability.

It is possible to objectively measure the level of variability between narratives. The basic lexical variability measure is the Lexical Overlap (LO) method [25, 7] which measures the lexical overlap between sentences, i.e., the cardinality of the set of words occurring in both sentences. Other studies suggest a variation of this method for text documents and produce lexical matching similarity scores which were based on the number of lexical units that occur in both input segments. More recent studies [4, 6, 10] show that adding semantic analysis, such as WordNet [8] and other statistical corpus data, can improve the accuracy of classification of similar documents over the basic methods of lexical matching/overlap.

In our case, we know that the documents are semantically similar as they all describe similar social situations. As such, we only need to measure the narratives' lexical variability and therefore we decided to use lexical matching methods. We chose the Ratio and Cosine models for measuring lexical overlap because recent work found these to be the most effective [12]. The Ratio model [29] is a binary model which only considers the occurrences of the words in each document. The Cosine model [24] is a count model which also considers the number of occurrences of each word within the document. All narrative evaluation scores have been measured without the Stop Words, which is the most commonly preferred method. We use the set of English Stopwords taken from `http://www.textfixer.com/resources/common-english-words.txt`. The similarity scores are defined as follows:

- The Ratio Model – a binary similarity model (Tversky's (1977)[29]):

$$S_{ij}^{\text{Ratio}} = \frac{a_{ij}}{a_{ij} + b_{ij} + c_{ij}}$$

- The Cosine model – a count similarity model (Rorvig's (1999)[24]):

$$S_{ij}^{\text{Cosine}} = \frac{\sum_k x_{ik} x_{jk}}{(\sum_k x_{ik}^2 \sum_k x_{jk}^2)^{\frac{1}{2}}}$$

where:

- x_{ik} counts the number of times that the k-th word occurs in the document i.
- t_{ik} denotes whether the k-th word occurs in the document i, i.e:
  $t_{ik} = 1$ if $x_{ik} > 0$ and $t_{ik} = 0$ if $x_{ik} = 0$.
- For the i-th and j-th documents, the count $a_{ij} = \sum_k t_{ik} t_{jk}$ is the number of words that are common to both documents.
- For the i-th and j-th documents, the counts $b_{ij} = \sum_k t_{ik}(1 - t_{jk})$ and $c_{ij} = \sum_k (1 - t_{ik}) t_{jk}$ are the distinctive words that one document has but the other does not.

SNACS' major advantage over hand-crafting stories is the time saved. However, as we note here, SNACS also produced significantly more varied and diversified stories than those based on the planning-based generation algorithm. Table 5 demonstrates this claim by

Table 5 Similarity Scores.

Algorithm	Movie		Restaurant	
	Ratio	Cosine	Ratio	Cosine
Original	0.08915	0.35634	0.07192	0.15028
Planner	0.20149	0.55960	0.26066	0.52533
Rnd-Planner	0.19184	0.59161	0.26240	0.51384
RND-SNACS	0.09156	0.40798	0.11768	0.33108
SNACS-Any	0.08417	0.35171	0.12598	0.30366
SNACS-Bst	0.09910	0.51831	0.11548	0.32730
SNACS-Tag	0.10248	0.50804	0.10795	0.32415

presenting the similarity measures, ratio and cosine, for all of the algorithms. As expected, the `original` narratives, which are completely hand-written, got the best scores (lower is better) in both measures. The `original` narratives got a ratio score of 0.089 and 0.072 and a cosine score of 0.36 and 0.15 for the movie and restaurant narratives respectively.

The average ratio score for the SNACS algorithms (`SNACS-Any, SNACS-Bst, SNACS-Tag` and `Rnd-SNACS`) is 0.094 for the movie narratives and 0.117 for the restaurant narratives. The average ratio score for the planning-based algorithms (`Planner` and `Rnd-Planner`) is 0.197 for the movie narratives and 0.262 for the restaurant narratives, which are significant worse than the SNACS algorithms. This difference in the results was seen also in the cosine scores. The average cosine scores are 0.447 and 0.322 for the SNACS algorithms and 0.576 and 0.520 for the planner based algorithms for the movie and restaurant narratives respectively. When looking at the generated narratives, the average ratio scores for all of the SNACS algorithms (`SNACS-Any, SNACS-Bst, SNACS-Tag` and `Rnd-SNACS`) were significantly lower (which is better) than the planning-based algorithms (`Planner` and `Rnd-Planner`), for the movie narratives and for the restaurant narratives (the ANOVA tests for mean difference of the entertainment narratives' ratio and cosine scores of `Planner` compared to `SNACS-Tag` at the 0.05 significance level being $p = 7.65 \cdot 10^{-4}$ and $2.31 \cdot 10^{-3}$ respectively). Overall, the SNACS algorithms have generated more variable and versatile narratives.

5 Related Work

Studies of automated storytelling have focused on creating systems to generate novel narratives based on prior authored knowledge and logical representations of narrative structure. The two most common approaches to story generation are case-based reasoning and planning. The case-based story generators, such as [28, 27, 11], construct novel narratives by reusing prior narratives, or cases, while planning-based story generation systems, such as [16, 23, 21], use a causality-driven search to link a series of primitive actions in order to achieve a goal. These methods for automatic story generation require that a person primarily author explicitly detail most parts of the narrative, including details about the story and the domain within which it takes place. While such narratives are well written, manually writing makes the development process costly in both time and resources. We compared our automatically generated narratives with those that were planning-based generated and found that overall people found the automatically generated narratives as realistic and consistent as those manually created by content experts. The main advantages of SNACS are the simplicity and the rapidity with which it achieves its results. Furthermore, the narratives

generated with our method are much more varied.

Early studies such as Tale-Spin [16] and Minstrel [28] focused on the creation of a standard representation that can be used to describe characters and their goals, the story's setting, etc. but didn't get any real-time input from the user. More recent studies focus on interactive narrative, which enables the player to make decisions which directly affect the direction and/or outcome of the narrative experience being delivered by the computer system [5, 20, 21, 23, 30]. As the human author is not present at run-time, authoring interactive narratives is often a process of anticipating user actions in different contexts and using computational mechanisms. These studies are similar to ours in that they need to adjust the story to the new constraints presented by the user, but still differ from the current study. They all focus on the creative side of the story and on the high level of the story's plot and characters' goals, while the domain and the story's detailed description were mostly manually written, and we use the wisdom of the crowd to acquire our domain.

Recently, a few studies have started explore ways to automatically build the knowledge base. The SayAnything [26] application constructs new stories from fragments of stories mined from online blogs. The user and computer take turns writing sentences of a fictional narrative, where the sentences contributed by the computer are extracted from Internet weblogs. Our study differs from this work; SNACS generates new content based on the narrative collection while this work only selects and reuses appropriate narrative fragments. In [11] the authors apply case-based reasoning techniques to build an intelligent authoring tool that learns cases from human storytellers who enter dialogue-based narratives via a custom interface. The cases can only be expressed in terms of a known set of possible pre-defined actions. They also use an utterances library which was manually pre-authored and pre-tagged with sets of utterances for each possible dialogue act and therefore limited to a given domain. The authors in [13] introduce an approach to automatically learn script-like sociocultural knowledge from crowdsourced narratives. They describe a semi-automated process by which they query human workers to write natural language narrative examples of a specific, given situation and learn the set of events that can occur and its typical ordering. Our approach also acquires its knowledge from crowdsourced narratives, but it still differs from this work. We use the crowd as the source for the detailed description of situation narratives and thus we focus more on the overall situation and less on the specific, atomic events that compose it.

6   Conclusions and Future Work

In this paper we presented SNACS, a novel approach for generating everyday activities narratives using crowdsourcing. Instead of manually handwriting, which makes the development process costly in both time and resources, SNACS uses crowdsourcing to create a varied base of stories. We compared SNACS to manually handwritten narratives, those generated by SHOP2, a state-of-the-art HTN planner [18] and random baselines. Our evaluation showed that our SNACS narratives were rated as being as believable and consistent as those which are manually handwritten or created from the HTN planner. Yet, the SNACS-based narratives had the main advantage of ease and the rapidity with which these narratives were generated. A second significant advantage to SNACS is that these narratives are rated to be significantly more diversified than those from the HTN planner – all while still maintaining their believability. SNACS is also more adaptable to creating new narratives and it can be easily extended to produce narratives for new domains. In future work we would like to integrate the SNACS algorithm into a dialogue-based application with a virtual human. We

would like to apply our algorithm to different types of applications, such a virtual agent that uses narratives to help interact better with the elderly and for dialogue-based training systems.

References

1 Amazon Mechanical Turk services. `http://www.mturk.com/`.

2 A. Azaria, Y. Aumann, and S. Kraus. Automated strategies for determining rewards for human work. In Jörg Hoffmann and Bart Selman, editors, *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22–26, 2012, Toronto, Ontario, Canada*, pages 1514–1521, 2012.

3 A. Azaria, Z. Rabinovich, S. Kraus, C. V. Goldman, and O. Tsimhoni. Giving advice to people in path selection problems. In Wiebe van der Hoek, Lin Padgham, Vincent Conitzer, and Michael Winikoff, editors, *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012, Volume 1*, pages 459–466. International Foundation for Autonomous Agents and Multiagent Systems, 2012.

4 D. Bär, T. Zesch, and I. Gurevych. A reflective view on text similarity. In Galia Angelova, Kalina Bontcheva, Ruslan Mitkov, and Nicolas Nicolov, editors, *Recent Advances in Natural Language Processing, RANLP 2011, 12–14 September, 2011, Hissar, Bulgaria*, pages 515–520, 2011.

5 M. Cavazza, F. Charles, and S. J. Mead. Character-based interactive storytelling. *IEEE Intelligent Systems*, 17(4):17–24, 2002.

6 C. Corley and R. Mihalcea. Measuring the semantic similarity of texts. In *EMSEE '05, Proceedings of the ACL Workshop on Empirical Modeling of Semantic Equivalence and Entailment*, pages 13–18. Association for Computational Linguistics, 2005.

7 M. Damashek. Gauging similarity with n-grams: Language-independent categorization of text. *Science*, 267(5199):843–848, 1995.

8 C. Fellbaum. *WordNet: An electronic lexical database*. MIT Press, 1998.

9 A. Gatt and E. Reiter. SimpleNLG: a realisation engine for practical applications. In Emiel Krahmer and Mariët Theune, editors, *ENLG'09, Proc. of the 12th Europ. Workshop on Natural Language Generation*, pages 90–93. Assoc. for Computational Linguistics, 2009.

10 O. Glickman, E. Shnarch, and I. Dagan. Lexical reference: a semantic matching subtask. In Dan Jurafsky and Éric Gaussier, editors, *EMNLP 2007, Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, 22-23 July 2006, Sydney, Australia*, pages 172–179. Association for Computational Linguistics, 2006.

11 S. Hajarnis, C. Leber, H. Ai, M. Riedl, and A. Ram. A case base planning approach for dialogue generation in digital movie design. In A. Ram and N. Wiratunga, editors, *Case-Based Reasoning Research and Development*, pages 452–466, 2011.

12 M. D. Lee, B. M. Pincombe, M. B. Welsh, and B. Bara. An empirical evaluation of models of text document similarity. In B. G. Bara, L. Barsalou, and M. Bucciarelli, editors, *Proceedings of the 27th annual meeting of the Cognitive Science Society*, pages 1254–1259. Erlbaum, 2005.

13 B. Li, S. Lee-Urban, D. S. Appling, and M. O. Riedl. Automatically learning to tell stories about social situations from the crowd. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 142–151, İstanbul, 2012.

14 B. Magerko, R. E. Wray, L. S. Holt, and B. Stensrud. Customizing interactive training through individualized content and increased engagement. In *The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC), Volume 2005: One Team, One Fight, One Training Future*, St. Petersburg, Florida, 2005. Simulation Systems and Applications Inc.

15 W. Mason and S. Suri. Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods*, 44:1–23, 2012. 10.3758/s13428-011-0124-6.

16 J. R. Meehan. TALE-SPIN, an interactive program that writes stories. In D. R. Reddy, editor, *Proceedings of the 5th International Joint Conference on Artificial Intelligence. Cambridge, MA, August 1977*, pages 91–98. William Kaufmann, 1977.

17 D. S. Nau, T. C. Au, O. Ilghami, U. Kuter, H. Muñoz-Avila, J. W. Murdock, D. Wu, and F. Yaman. Applications of SHOP and SHOP2. *IEEE Intelligent Systems*, 20(2):34–41, 2005.

18 D. S. Nau, T. C. Au, O. Ilghami, U. Kuter, J. W. Murdock, D. Wu, and F. Yaman. SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research*, 20:379–404, 2003.

19 J. Niehaus, B. Li, and M. O. Riedl. Automated scenario adaptation in support of intelligent tutoring systems. In R. C. Murray and P. M. McCarthy, editors, *Proc. of the 24th Int'l Florida Artificial Intelligence Research Society Conference*, pages 531–536, 2010.

20 M. O. Riedl. Incorporating authorial intent into generative narrative systems. In Sandy Louchart, Manish Mehta, and David L. Roberts, editors, *Intelligent Narrative Technologies II. Papers from the AAAI Spring Symposium*, number SS-09-06 in AAAI Technical Reports, pages 91–94, Menlo Park, 2009. AAAI Press.

21 M. O. Riedl and A. Stern. Believable agents and intelligent story adaptation for interactive storytelling. In S. Göbel, R. Malkewitz, and I. A. Iurgel, editors, *Technologies for Interactive Digital Storytelling and Entertainment, Third International Conference, TIDSE 2006, Darmstadt, Germany, December 4–6, 2006*, number 4326 in Lecture Notes in Computer Science, pages 1–12. Springer, 2006.

22 M. O. Riedl, A. Stern, and D. M. Dini. Mixing story and simulation in interactive narrative. In *Proc. of the 2nd Artificial Intelligence and Interactive Digital Entertainment Conf., June 20–23, 2006, Marina del Rey, California*, pages 149–150. The AAAI Press, 2006.

23 M. O. Riedl and R. M. Young. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:217–268, 2010.

24 M. E. Rorvig. Images of similarity: A visual exploration of optimal similarity metrics and scaling properties of TREC topic-document sets. *Journal of the American Society for Information Science*, 50(8):639–651, 1999.

25 G. Salton. *Automatic Text Processing: The Transformation, Analysis, and Retrieval of.* Addison-Wesley, 1989.

26 R. Swanson and A. Gordon. Say Anything: A massively collaborative open domain story writing companion. In U. Spierling and N. Szilas, editors, *Interactive Storytelling, First Joint International Conference on Interactive Digital Storytelling, ICIDS 2008, Erfurt, Germany, November 26-29, 2008, Proceedings*, number 5334 in Lecture Notes in Computer Science, pages 32–40, 2008.

27 B. Tearse, N. Wardrip-Fruin, and M. Mateas. Minstrel remixed: Procedurally generating stories. In G. M. Youngblood and V. Bulitko, editors, *Proceedings of the Sixth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE 2010, October 11-13, 2010, Stanford, California, USA*, pages 192–197, 2010.

28 S. Turner. MINSTREL: a computer model of creativity and storytelling. Technical Report UCLA-AI-92-04, University of California at Los Angeles, 1993.

29 A. Tversky. Features of similarity. *Psychological Review*, 84(4):327–352, 1977.

30 S. G. Ware and R. M. Young. CPOCL: A Narrative Planner Supporting Conflict. In V. Bulitko and M. O. Riedl, editors, *Proceedings of the Seventh AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE 2011, October 10-14, 2011, Stanford, California, USA*, pages 97–102, 2011.

31 Peter Weyhrauch. *Guiding Interactive Drama*. PhD thesis, Carnegie Mellon University, 1997.

Towards a Computational Model of Dramatic Tension*

Nicolas Szilas and Urs Richle

TECFA, Faculté de Psychologie et des sciences de l'éducation
University of Geneva
Genève, Switzerland
`Nicolas.Szilas@unige.ch`

─── **Abstract** ───

One of the approaches to generate narrative consists in modeling narrative in terms of a deep structure, as introduced by narrative theories in the middle of the 20th century.

This papers revisits this computational approach, and raises the central issue of dramatic tension: Would it be possible to build a computational model of dramatic tension, where tension could be managed according to the well known ascending/descending dramatic curve?

The paper describes a new computational model of narrative, based on a set of structural narrative elements (goals, tasks, obstacles, side-effects), a hierarchical and modular approach, a paradox-based model of dramatic tension and a solution for managing endings.

The papers illustrates this theoretical model with a full example.

1 Structures and Paradox

There exists a multitude of computational models dedicated to the narrative generation. A promising approach is based on structuralist theory. The principle is rather simple: starting from Lévi-Strauss study of myths [23], the same way a unique deep structure of the myth can produce many different specific myths across communities and cultures, a unique core structure of a narrative could produce, when simulated, a large variety of different stories, across different executions. The key feature of such structures – and this is important to insist since the term "structure" shares different and sometimes contradictory meanings [38] – is their *atemporal* nature [23]. Such deep structures contain various interconnected narrative elements that, at least in part, do not specify a *specific* temporal relation between these elements. For example, stating that two characters are in conflict or that the narrative revolves around two thematic elements are atemporal specifications. These specifications later produce temporally ordered actions or events.

There has been some attempts to build systems based on these principles. In Black Sheep [20], the six actant model of Greimas [17] was implemented, but the model remained very basic, so the benefit in terms of variations could not be observed. In *IDtension* [36, 34], a specific structural model was developed for Interactive Drama, integrating the dramatic

notion of ethical conflict. Although a full scale and playable story could be implemented [18], it appeared quite difficult to write deep structures, and these structures were often designed according to branching or conditional approaches.

A deep narrative structure is a compact set of narrative elements that describes the core issue in a narrative. This structure is then processed to produce narrative events. Such a processing is possible if two components are implemented:

- An "Action Generator", able to generate possible actions, based on a deep structure.
- A "Narrative Sequencer", able to select actions among the many ones that are generated by the Action Generator, based on how the narrative is perceived by the audience. This is typically implemented with a model of the user [2, 35, 42].

This Narrative Sequencer is very challenging to build, because it is hard to find an agreement within narrative theories regarding the set of criteria that determine that a sequence of events is perceived as a narrative or not. Which criteria should be considered first, to guide the narrative generation? Closure, conflict, emotions, etc.? We suggest to consider *dramatic tension* as a good candidate for the model of the user, because it is widely recognized and intuitively observed that in a drama, tension increases and decreases during the narrative. However, it is not straightforward to implement the concept of dramatic tension, because this concept is not well defined in narrative theories. Modeling tension consists in evaluating the tension of generated actions, which implies the construction of a computational model of the dramatic tension. This is different from tagging events with a tension level, as in the interactive drama *Façade* [25]. The tagging approach only works when narrative events are not generated but pre-written.

In dramaturgy and screenwriting, the concept of tension is associated with the dramatic curve, that describes dramatic tension as a function of time: dramatic tension increases up to the climax, occurring at approximatively two thirds into the narrative and then decreases until the end of the narrative. This curve was introduced by G. Freytag in the mid nineteen century [24]. But what is it exactly that increases and then decreases during the typical drama? What makes the tension increase or decrease? This is usually not clearly defined [5]. There exist however other related concepts that are better defined in various narrative theories: suspense, conflict and paradox.

Suspense can be rigorously defined [8]. It occurs when a certain story state (target state) is strongly hoped by the audience and when among the various outcomes that can be anticipated, the target state is perceived to have very little chance to occur. (Reciprocally, suspense can be defined via a fear of a negative state that has a high chance to occur.) Suspense can be produced either at the story or at the discourse level.

Conflict denotes the "struggle in which the actors are engaged" [29]. For many theorists and writers of drama, this is a core concept [13, 14, 22, 26], and it is common to read "drama is conflict". This is, as tension, a very common term. The main conceptual difficulty is that there are different types of conflicts, which, when gathered into one single word, create a fuzzy concept. Generally speaking, conflict arises when the two goals cannot be reached at the same time. Depending on the kinds of goals, different kinds of conflict can be distinguished, such as conflict of interest, ethical conflict, external vs internal conflict.

Finally, dramatic tension can be the result of the paradoxical nature of narrative [28]. According to B. Nichols, any narrative is based on a paradox, that is a logical impossibility, as illustrated in classical paradoxes: "Epimenides was a Cretan who said, 'Cretans always lie'." (paradox of Epimenides). The narrative paradox is the juxtaposition of two contradictory terms, that successive narrative actions attempt to resolve. Trying to "resolve the unresolvable" along the major part of a story inevitably creates tension. In some case, the story can ends

by operating a major shift in the paradox, where the paradox's term are re-interpreted with a new point of view. In this case, the paradox is not solved but rather dissolved: it does not exist anymore. In other cases, the paradox remains until the end. This approach finds it roots in the analysis of the myth by C. Lévi-Strauss [23] and can be related to, outside the narratology realm, paradoxical injunction [41] and humor [21].

As an example, in a love story, the male hero must accomplish a "bad" activity to conquer his love, but if he performs this activity, she might learn it and refuse him. Nichols provides many examples in feature films. For example, for *The Birds*, from A. Hitchcock: "If I am to win Mitch, I must become part of his family, but if I become part of his family, I can not win Mitch".

There exist computational models of both suspense [11] and conflict [4, 6, 31, 34, 40]. In this paper, we want to explore the modeling of paradox. What is particularly interesting in this theory is that a certain configuration of narrative elements can generate an infinite temporal behavior, when characters try to solve the paradox but, by definition, cannot solve it. If the paradox could be expressed as a structural property of the above-mentioned deep narrative structures, then we would have an elegant and powerful way to generate a flow of actions that are narratively relevant, since they express the fundamental struggle of the narrative.

2 A computational model of narrative based on dramatic tension

2.1 General architecture

Based on the narratological discussion that precedes, we will sketch a computational model able to generate a flow of events that are narratively relevant, this relevance being assessed according to a modeling of the dramatic tension. This computational model will meet the three following requirements:

First, the model must be able to easily identify what is at stake in the narrative, that is the fundamental paradox of the narrative. Other elements should be implemented as well – without them, the narrative would not work – and they should serve to express the fundamental paradox.

Second, the model must be able to expand elements in the structures into a variety and quantity of narrative actions, which enables to make the deep structure work. This will be the role of the "Action Generator" mentioned above.

Third, the concept of dramatic tension must be formalized, so that 1) at each moment, it is possible to evaluate the current tension; 2) for each possible action, it is possible to associate the impact on the tension and 3) a strategy for managing this tension is designed.

While the focus of this paper is on the deep narrative structure, other components also play an essential role and will be described as well. In particular, it is important to keep in mind that the elements of the structure are not, in many cases, the actions that are performed. The latter are generated from the former by the Action Generator (see Section 2.5).

2.2 The core components of the Structural Model

The basic elements of the structural model will be largely inspired from previous research [36, 37], although alternative representations could be explored [9].

Four simple elements, along with their interrelations will constitute the basic "bricks" to build a structure and are called *nodes*:

Goals: A goal represents a state of affairs that a character may wish to reach. This concept is omnipresent in dramaturgy [22, 26, 39] and also intervenes is some recent definitions of narrative [30], since it corresponds to the fundamental notions of characters' actions and intentions. A large number of computational models of narrative implement goals [1, 4, 10, 27, 36, 43], in particular agent-based models, since goals are a fundamental concept in Artificial Intelligence.

The character who has a goal is called the *actor* of the goal. In the discrete case (default), the goal is either reached or not reached. In the continuous case, the goal is associated to an *achievement value*, that varies between 0 and 1, 1 meaning that the goal is fully reached. When a goal is activated (added in the active structure, see below), the achievement value is set to 0. Each time a task connected to the goal succeeds, the achievement value is incremented by a fixed parameter called the *increment*.

Tasks: A task represents a concrete action that a character can perform to reach a goal. It corresponds to the notion of action or operator in several AI or robotic frameworks. However, in the current model, actions are based on tasks but cannot be reduced to tasks, as explained in Section 2.5.

When performed, a task leads either to a success, in which case the goal is reached (see above), or to a failure, in which case the goal in not reached. By default, a performance of a task is successful. But obstacle attached to the task may lead to a failure.

Obstacles: An obstacle represents an event attached to a task that may hinder the reaching of a goal, when the task is attempted by a character. This is where the model departs from classical AI representations, due to the nature of narrative: while the focus in robotics is to reach a goal as efficently as possible, the focus of narrative is on the difficulties met by the actors when they try to reach their goals. Obstacle is a classical way of representing these difficulties, and is borrowed from dramaturgy and screenwriting [22, 26]. Obstacles implement the important notion of failure that is central in the narrative theory of C. Bremond [7].

An obstacle is attached to a task. If an obstacle *triggers* after an attempt to perform the task, the goal is not reached and it may impact the storyworld (see relations below). An obstacle is associated with a probability of triggering, called the *chance*. It is comprised between 0 and 1 (1 by default). This value is author-defined, according to the frequency of occurrence the author wants to give to the obstacle. For some obstacles, the chance to trigger is undecided: the decision to trigger them is at the discretion of the narrative engine. These obstacles are called *free obstacles*.

Side-effects: A side-effect represents an event attached to a task that may trigger when the task is performed by a character. But contrary to an obstacle it does not prevent the reaching of a goal. Side-effects thus implements a variant of a task. Side-effects are not a prominent concept in narrative theory or dramaturgy. However, it is included in the model because we lacked such a concept, when a creative author was writing stories with IDtension [12], which implements obstacles but not side-effects. More precisely with obstacle, plenty of variants of a task can be written, that occur unpredictably when a character starts to perform a task, but they all lead to a failure. Symmetrically, it would be relevant to write variants of tasks which occur unpredictably even when the task succeeds. When a side-effect is not known in advance by the character, it corresponds to an involuntary action, as described by C. Bremond: An agent undertakes a task but performs at the same time an "involuntary action" [7, p. 237].

As obstacles, side-effects are associated to a chance (1 by default).

While AI-oriented models of narrative use world states to express the consequences and preconditions of action execution, the proposed model avoids this kind of representation that is cumbersome for creative authors. Rather, the dynamic of goals, tasks, obstacles and side-effects is entirely described with *relations* (or *links*) between these nodes. The following relations have been identified:

Reaching: A *reaching relation* connects a task to the goal it enables to reach. In the continuous case, it contains the value added to the achievement value, called the *achievement increment*.

Attachment: An *attachment relation* connects an obstacle or a side effect to the task that may trigger it.

Sub-goaling: A *sub-goaling relation* connects an obstacle to a goal: it indicates that the obstacle (called the triggering obstacle), when it triggers, activates the goal (called the sub-goal). The task remains blocked by the obstacle until the sub-goal is reached. In the continuous case, the task remains blocked by the obstacle until the achievement value of the sub-goal reaches a certain threshold, called the *achievement threshold*. The achievement threshold is contained in the relation. The relation may also contain the actor of the sub-goal, if different from the actor of the obstacle.

Inhibiting: An *inhibiting relation* connects a goal to an obstacle or side-effect. The chance of the target obstacle or side-effect is set to 0 if and only if the source goal (called the inhibiting goal) is reached. In the continuous case (chance between 0 and 1), the chance is modulated, according to a parameter called the *inhibiting factor*, comprised between 0 (no modulation) and 1 (maximum modulation). Furthermore, the resulting chance also depends on the achievement value of the inhibiting goal. Symmetrically, an *exciting relation* may also be used, that will not be detailed here.

Needing: A *needing relation* connects a side-effect to a goal, different from the goal reached by the task. When the side-effect is triggered, it activates the target goal (instantiates it in the active structure), if it is not already activated. If it is already activated, it sets its achievement value to 0. In the continuous case, the target goal achievement value is diminished by a certain value, called the *decrement*. The relation may also contain the actor of the triggered goal, if different from the actor of the side-effect.

An *abstract structure* is a graph containing nodes and relations. It corresponds to what the author writes, in terms of abstract narrative content (data needed to display this content is described separately). It can be observed that such a structure does not contain any characters, and therefore does not represent any concrete narrative. For the narrative to become concrete, some goals, and their associated elements in the structure, must be *activated*. Activated elements constitute the *active structure*, which contains some elements described in the abstract structure, completed by the information about actors. The active structure is one instantiation of the abstract structure. Algorithms that build and manage the active structure from the abstract structure will not be detailed here. They are derived from [37]. In particular, these algorithms specify what changes in the target element when a relation fires.

Abstract structures do not contain variables, except the actors. In this model, it is therefore not possible to represent the fact that a character can steal an object in general. This choice is motivated by the search of a simpler model, for a first version. Variables being a powerful way of producing variations in a computer-generated narrative, they will be considered for a later version of the model.

■ **Figure 1** Graphical representation of the elements of an abstract structure. The type of relation is indicated by three letters: *rea* for reaching, *sub* for sub-goaling, *inh* for inhibiting and *nee* for needing. When the relation is associated to a value, this value is represented after the three letter code, between parenthesis. The attachment relation is not depicted, but simply represented by positioning the obstacle or side-effect on the task it is attached to. The number next to the obstacle is the associated chance. If the letter *F* is written instead, it means that this is a free obstacle (see text).

The nodes and relations can be represented visually, which is essential for the authoring process. Figure 1 depicts an example of an abstract structure containing all types of nodes and relations introduced above. Further examples are provided in Figures 2 and 5.

2.3   Adding hierarchy

When modeling a narrative at the higer level, it is important to keep the structure simple, avoiding the resulting graph to become undecipherable. At the same time, the narrative structure needs a significant amount of content to be able to express variety and richness. To solve this dilemma, one modeling technique should be added: *modularity*. It consists in grouping subsets of elements into modules, and in reasoning at the level of theses modules rather than at the level of the elements themselves. This creates a hierarchy between the lower level (simple elements, inside a module), and the upper level (modules). This approach is classical in Artificial Intelligence and has proven useful, both for computational and ergonomical reasons; see for example Hierarchical Finite State Machines [16], hierarchical petri nets [3], hierarchical neural networks [19], etc. Only two levels of hierarchy are considered in this paper.

The concept of *task module* is introduced. A task module is a subset of tasks, goals, obstacles and side-effects. This subset as a whole can connect to a goal the same way a single task can connect to a goal, via a reaching link. If an internal task with no goal attached is finished, it fires the reaching link (other tasks do not need to succeed). In the continuous case, the increment of the reaching link is specified for each internal task.

Other relations and firing mechanisms are also introduced, to and from a task module:

- Because a task module contains side-effects and obstacles, it can also connect to a goal via a needing link.
- An internal side-effect can also fire the reaching link of the enclosing task module.
- An internal side-effect or goal can be modulated by an external goal, via an inhibiting or exciting relation.

Figure 2 Graphical representation of the hierarchical and modular model for structures. Small grey arrows pointing up denote the decrement associated to the needing relation (*nee*) from the task module to the Goal *A*. Small grey arrows pointing to the right denote the achievement increment associated to the external reaching link (*rea*) from the task module to the goal *B*.

An illustration of a structure with a task module is proposed in Figure 2. In this example, the core structure is still clearly visible, while the story has gained in complexity.

2.4 Structures in tension

From the above description, an abstract structure, with its set of nodes and relations, could perfectly be created in such a way it would not present any dramatic interest. For example, the author can choose 10 tasks, connected to 10 goals, and activate the 10 goals at the beginning. As a result, the story would be a random succession of unconnected actions, which does not fit with what is usually considered as an interesting dramatic situation: in this case, the unity of action is typically lacking. So far, the model is underspecified. Therefore, the key question is: what is a good structure?

We postulated previously that a good structure is a structure that contains tension in it, and that such a structure contains some *circularity* [38]. This circularity expresses characterization of a dramatic situation as a "system of forces in internal tension"("un système de forces en tension intérieure" [32, p. 42], our translation). To further formalize and specify this circularity, we propose to think the structure in terms of paradox, as introduced in Section 1.

The structure depicted in Figure 2 contains circularity. It says that if a character wishes to reach the goal *A*, he or she meets an obstacle than leads him or her to wish to reach the goal *B*. But to reach the goal *B*, one hinders the goal *A* (needing relation from the task module to the goal *A*). In short, if this character starts to reach goal *A*, he or she does not reach goal *A*. As formulated by Nichols [28], this impossible situation creates the dynamics of narrative, when characters try to solve this paradox.
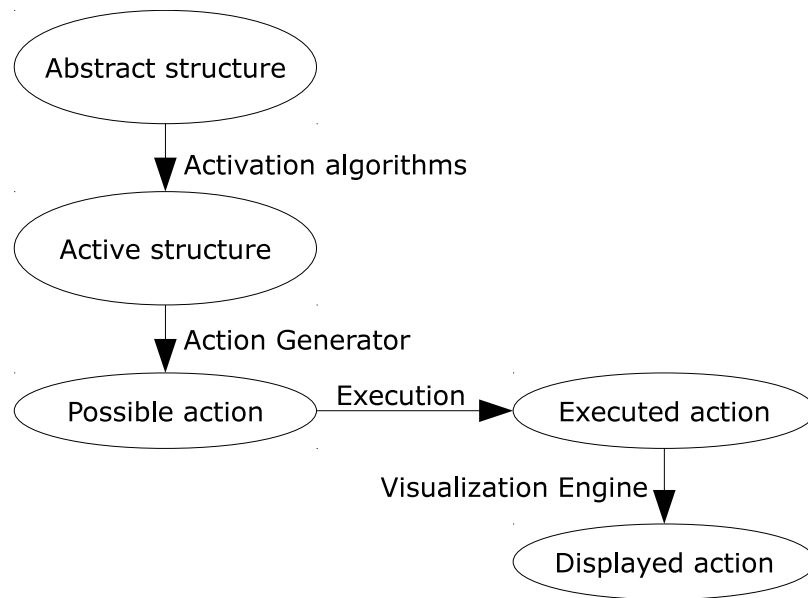
◼ **Figure 3** Causal chain of an action, from the abstract structure to its display.

Let's fill the structure depicted in Figure 2 with some narrative content. John desires that princess Mary accepts him as a future spouse (goal A). For that purpose, he proposes to her (task 1), but she cannot accept, because he is not a prince (obstacle). So he decides to become a fake prince (goal B, triggered as a sub-goal), which he can perform via different tasks, grouped into a task module. Performing these tasks of appearing like a prince makes him a liar, which in turn makes Mary not want to marry him, via the side-effect (Adapted from the animated movie *Aladdin*, Disney, 1992).

2.5 Playing around the structure

Tasks, obstacles and side-effects do not equate to what is executed and finally displayed to the user. The engine needs to calculate an action, derived from the active structure by applying a predicate on one or more of its constituting elements. The first predicate is PERFORM. It takes as parameters the task to perform and, optionally, the list of obstacles and side-effects that will trigger during the execution. An example of performance action is: PERFORM(ProposeTo(John,Mary)).

When an action is executed, it is added to the *story*. An executed action is possibly *displayed*, in which case the player can perceive it (visually, textually, auditory, etc.). To sum up, when the player perceives an action, it comes from elements in the abstract structures, which have been activated, then transformed into an action, which has been executed and finally displayed (figure 3).

Many other types of action can be found in narrative. In previous research [36], we have identified influences (encourage/dissuade), sanctions (felicitate/condemn), delegations (ask for help, propose help, accept help), etc. But reducing the types of actions to a predetermined set of possibilities proved limiting, in some contexts [33]. Therefore, we leave it open to the author which action types should be implemented, in addition to the PERFORM one. This supposes that a sort of language is designed that enables an author to program action types. For example, an author might specify the action type ASK_ADVICE, which would

generate the action: ASK_ADVICE(John,Bill,ProposeTo(John,Mary)), meaning that John asks advice to Bill concerning the idea of proposing to Mary.

These predicate-based actions are generated by author-defined rules that specify possible actions, that is actions that are logically compatible with the current active structure and the story so far. For example, a rule would specify that if a character has a goal, then he or she can ask advice to any other character, regarding any task that reaches that goal. These rules are called *possibility rules,* they are managed by the Action Generator.

Other rules must specify priorities among possible actions, so that the engine can select which actions, among all possible actions, are finally executed. For example, a rule would favor the action of asking for advice regarding a task, just after the character has been attached to the corresponding goal, in the active structure. These rules are called *preference rules*, they are managed by the Narrative Sequencer. Preference rules could be refined with a mechanism that looks ahead at the consequences of possible actions, using planning technology [43].

Finally, the visualization engine transforms the predicate-based action into a perceivable output that may include text, image, sound, three-dimensional simulation, etc. In particular, some text generation algorithms are necessary at this level.

2.6 Tension measurement and management

In the previous sections, we have provided a construction kit for structures (Sections 2.2 and 2.3), a sketch of a second construction kit for actions (Section 2.5) and a guideline for creating interesting core structures, based on the notion of paradox in narrative (Section 2.4). The next step is to enable the Narrative Sequencer to select an action according to an evaluation of the dramatic tension.

Following the view that a narrative paradox, by its logical impossibility, arouses the fundamental dynamics of narrative, we postulate that tension is created by this movement. More precisely, we propose that one important strategy to increase the dramatic tension is to force the focus of the audience to switch from one proposition in the paradox to another one. Effectively, it can be supposed that these switches activate the paradox by a phenomenon of recency: the audience lives the current position while recalling the previous one.

For example, in a two goal structure such as the one depicted in Figure 3, tension will increase if the action of mentioning the goal *A* to the user is followed by the mentioning of the possible failure of a task reaching this goal. Dramatic tension would be, at least in part, a matter of contrast between successive actions.

This can be represented by a "paradox space", in which each action can be positioned, depending on its relation to the success or failure of the goals (see Figure 5). By construction of the paradox, it is impossible to fully reach the goal *A*. But an action can still be at the right side of the paradox space – success of goal *A* –, because this position is temporary, or partial (achievement value is lower than one), or anticipated, or because this success is simply evoked, in a dialogue. Therefore, axes in the paradox space do not measure the achievement value of the respective goals, but a certain relation between the action and the goal, as illustrated by the following examples:

- If a character encourages another one to perform a task by mentioning the reaching of the goal, the action is placed in the positive half-plane of the space, regarding this goal.
- If a character dissuades another one to perform a task by mentioning the failure of the other goal due to a side-effect, the action is placed in the negative half-plane of the space, regarding this second goal.
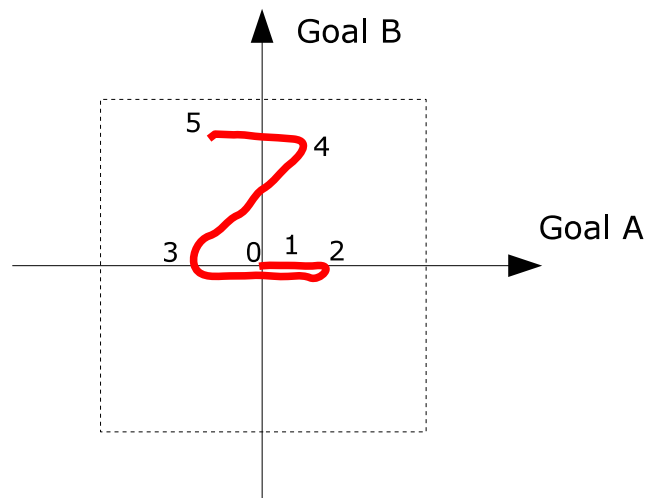
■ **Figure 4** The paradox space, and a possible trajectory in this space (see text).

━ If a character delegates a goal to another one, the action is placed in the positive half-plane of the space, regarding this goal (because it generates hope that the goal will be reached).
━ Etc.

Precise estimation of the position is not described in this paper because it needs further investigation and formalization.

Let us illustrate a trajectory in the paradox space with the help of Figure 4. The trajectory is the succession of the coordinates of visualized actions in the space (only actions related to goal *A* or goal *B* are reported). Some complex actions contain several coordinates in the space. The trajectory is represented continuously for the purpose of illustration. The situation starts at the centre (point 0): the story has not started and the user's perception is neutral. The user is then considering goal *A* (point 1), the user attempts (point 2) to reach this goal but fails (point 3). After that, the user is faced with the possibility to solve the problem raised by the obstacle by undertaking an action that reaches the goal *B* (point 4), but while succeeding the task, a side-effect makes him or her fail regarding goal *A* (point 5), etc.

We will not give formulas of tension calculation at this stage, but only describe the main features of this calculation:

━ It is initialized at 0.
━ When, in the paradox space, the trajectory crosses an axis between two displayed actions, the tension is increased, proportionally to the distance between the two successive actions in the paradox space.
━ The dramatic tension decays with time, meaning that if nothing happens or if executed actions do not impact the position in the paradox space, then tension decreases progressively.

In order to increase the tension, the Narrative Sequencer will adopt a strategy that consists in regularly choosing an action that crosses an axis in the paradox space.

At the same time, it is necessary to visualize the action differently, according to the level of dramatic tension the Narrative Sequencer wants to produce. This is difficult to do in a generic manner, and it highly depends on the medium. For example, if the medium allows

dynamic facial expression, then the emotions can be of higher valence when the tension is higher.

2.7 Endings

How does the story end? From a narrative perspective, for the story to end, the paradox needs to vanish. In a happy ending, something changes radically in the situation of the characters, which makes the paradox suddenly inexistent, and the main goal can be reached. For example, the law that states that a princess should necessarily marry a prince is abrogated. In a sad ending, the tension rises but the paradox remains. The tension becomes unbearable for the main character, until the main goal of the paradox disappears: the hero dies, the hero renounces, the main goal does not exist anymore (without being solved) – e.g., the princess marries another prince. In both happy and sad endings, the resolution of the paradox lies, by definition, outside of the paradox (what changed in the paradox logic? How the goal disappeared?).

Therefore, it is certainly not possible to automatically fully generate an ending, based on the terms of the paradox and general knowledge. At the general level, *types of ending* must be written by the author. Nevertheless, this does not mean that each ending is fully pre-written, because each ending can generates variations (see Section 3.4).

According to the computational model presented so far, it is possible to classify several types of endings. We will reason with the two goal case, as depicted in Figure 4. A first line of separation is naturally the sad vs. happy ending, which could be matched with the success or failure of the main goal. Another line of separation could be the success or failure of the second goal. In Figure 4, this makes four types of endings, corresponding to the four quadrants.

The role of the author is to write a certain number of endings, that she classifies in the four quadrants depicted in Figure 4. Finally, we need to specify how the narrative engine will select one of the four quadrants. The solution that we propose is in two steps:

First, each ending needs to be adaptive, so that, for example, the reason why one character changes its view should be explained by a specific action that the user has performed. More precisely, each ending is a script with some steps being activated according to a set of preconditions on the user's actions. This means that each ending has two or more variants. Furthermore, at least one ending has a variant by default, which can be chosen even if no condition on the user's actions is met.

Second, when the dramatic tension has reached a certain level (and possibly after a certain duration), an ending is triggered, among all possible endings for which one variant's preconditions are met. If several endings meet this criterion, one is chosen randomly. If no ending meets this criterion, one ending that has a default variant is chosen randomly.

3 A full example

3.1 Setting

This example is inspired by a research project that aims at helping teenagers cope when one of their parents suffers from a Traumatic Brain Injury, by offering them an interactive narrative related to their situation [18].

The user plays Frank, age 15 year, whose father, Paul, suffers from a Traumatic Brain Injury. Julia, a classmate of Frank, visits Frank at home to get advice on some math problems. When she arrives, and later in the scenario, the user has the possibility to either explain her

the family situation, or to directly help her with the math problems, which is a good thing if he wants to be appreciated by her (his goal in the story). Not explaining the situation might create problems later in the story when Paul interrupts the math session, while explaining too much might make Frank appears less cool, in front of his classmate.

From this context, the fundamental paradox is: "If I want Julia to like me, I need to help her with her math and to avoid any problem with my father I need to explain her the situation; but if I explain her the situation, then she will not like me". The structure is atemporal in that it does not say that "get math book" should occur before or after "explain accident". Of course, temporal relations can be immediately deduced from this structure, such as the execution of "explain accident" occurs before the reaching of the corresponding goal, "Julia prepared", but this corresponds to only one possible trajectory. For example, delegating the goal "Julia prepared" could occur before any attempt to perform "explain accident".

3.2   Structure

The corresponding abstract structure is depicted in Figure 6. Thanks to the two task modules, the core structure is clearly identified. The first task module contains two goals: one for helping Julia in math, and the other to deal with the father, who is otherwise constantly interrupting them. Obstacles and side-effects in the structure often fire external links, which gives them a central role in the narrative, beyond the role of providing a variant.

When instantiating the structures, characters are added: Frank (the teenager), Julia (the classmate), Paul (the father), Lili (the sister) and Olivia (the grandmother). The goal "appreciated by Julia" is instantiated, with Frank as actor. This automatically instantiates the goal "Julia ok with math" and the task "compliment Julia", to constitute the active structure. This active structure will be completed by other goals and task during the execution. Note that in this simple scenario, other characters have no goal at the beginning – they behave in reaction to Frank's actions.

3.3   Actions types and rules

In order to set in motion the active structure, action types must be defined, in addition to the default action type PERFORM. Each action type takes as parameters elements of the active structure to build a specific action. For this example, the following action types have been designed:

- Asking advice concerning a task (ASK_ADVICE): A character asks advice concerning the performance of an action.
- Asking advice concerning two tasks (ASK_ADVICE2): A character presents two alternatives to another character, in order to get advice.
- Influences regarding a task (ENCOURAGE/DISSUADE): A character encourages or dissuades another character to perform a task. Several variants are possible: influence referring to the consequence of the task (suffix: _CSQ), influence referring to a previous performance of the task (suffix: _PREV), influence referring to more than one consequences (suffix: CSQN).
- Comparing two tasks (COMPARE): In a dialog, a character compares two alternative tasks that reach the same goal.
- Help seeking (DELEGATE): A character delegates a goal to another character who will undertake the goal. It is supposed that whenever the second character has managed to

Figure 5 A full example of a narrative structure. The small arrow between "Paul angry" and "Julia vexed" means that the triggering of the side-effect is only considered if the obstacle has triggered. In the story, when Paul is angry, he becomes particularly aggressive and says bad words against Julia, so she might be vexed. But her reaction depends on wether she is prepared or not for this kind of situation (second main goal).

reach the goal, the first characters is immediately aware of it. Note that not all goals are delegable (defined by the author).

- Accept or refuse to help (ACCEPT_DLG, REFUSE_DLG): A character who has been asked for help accepts or refuses to help.
- Inform about a regret (REGRET): A character tells another character that s/he should not have performed a given task.
- Confirm regret (CONFIRM_REGRET): A character agrees with another character that the latter should not have performed a task.

For example, the action type "COMPARE" could generate the action COMPARE(Olivia, Frank, divert(Frank), reason_Paul(Frank)), which may be displayed with the following dialog line: "You could divert Paul but Julia might not appreciate your behavior. You could also try to reason Paul but he will certainly not listen to you".

Possibility rules and preference rules must also be defined, in order to transform elements in the active structure into executed actions. They have not yet been specified for this scenario.

3.4   Endings

Four families of endings can be distinguished, corresponding to the four quadrants described above, with the two following goals: "appreciated by Julia" and "Julia prepared" (respectively goals *A* and *B* in Figure 4):

Quadrant 1: Frank has pleased Julia by helping in math. For that, he has successfully explained the family situation to her. Finally, Julia finds Frank very brave and she even admires him, instead of being disgusted, vexed or annoyed.

Quadrant 2: Frank has managed to explain the family situation to Julia, but, as a consequence, Julia find him uninteresting, regardless his effort with math.

Quadrant 3: Despite Frank's efforts, Julia does not feel closer to Frank after the visit, neither has frank managed to explain her the family situation.

Quadrant 4: Frank has not managed (or even tried) to explain the family situation to Julia, because he did not want to offend her. As a result, because Paul always interrupted them, Frank has been unable to do anything to explain the math. But Julia was charmed by Frank's behavior in such a situation.

This example clearly distinguishes happy endings (quadrants 1 and 4) from sad endings (quadrants 2 and 3). It also illustrates that the happy endings need creative thinking to re-interpret the paradox and annihilate it (Julia finds Frank brave or charming). Figure 6 illustrates one of the ending, described as a script.

3.5   Output story

From the abstract structure (Section 3.2), the action types and the associated rules (section 3.3) and the story endings (Section 3.4), a story can be generated. The active structure is updated after each executed action. The following depicts a handmade outcome of this process, covering the beginning of the story, after the introduction. The text visualization has been written according to what a template-based text generation system could provide.

We first present the **formal action**, then the text-based visualisation.

Julia to Frank: "I gotta go now. Thanks for the explanation but I think I'll manage it differently next time." [event:ending2] "Bye!" Julia leaves. Olivia to Frank: "You know, it is difficult for a young girl to feel comfortable with your father". Frank leaves the room, slamming the door.

where the part in brackets refers to the description attached to an event that pushed in the Quadrant 2 during the story:

- For "explain accident": "you couldn't stop explaining your father's accident... you are becoming a drag"

- For "provide medical details": "It is so disgusting these brain operations you describe!"

- For "explain with book": "I came to study math, not medicine!" (showing the medicine book Frank showed her)

Figure 6 Script for ending 2.

1. **PERFORM(get math book(Frank)):**
   - obstacle "Paul interrupts" triggered

   Frank to Julia: "I would like to get you a great math book that will help you". He walks towards his bedroom to get the math book, but Paul comes to him and declares: "Frank, who is this stranger in our house?". Frank answers: "It's Julia Dad, my classmate, she is here to study math with me."

   Frank decides to do something to satisfy Paul's curiosity.

   Comment: *This adds the goal "Paul satisfied" into the active structure (visualized opposite as a narrator sentence)*

2. **PERFORM(divert(Frank)):**
   - obstacle "Fails" triggered
   - side-effect "Julia surprised" triggered

   Frank: "Dad, well, would you please look at the DVD player, it does not seem to work properly."

   Paul: "Oh no, I prefer staying here, with you and your friend."

   Julia: "Frank, you really have a strange way to talk to your father..."

   Frank decides to get Julia aware of the situation, so that she understands his behavior with his father.

3. **ASK_ADVICE(Frank, Olivia, explain accident(Frank))**

   Frank to Olivia: "Grandma, do you think I should shortly explain Dad's accident to Julia?"

4. **ENCOURAGE_CSQ(Olivia, Frank, provide medical details(Frank))** Olivia: "I think you should rather give her all medical details concerning your father's problem. Explain her what a traumatic brain injury is, and she has better chance to understand our family.

   Comment: *The consequence related to the encouragement is the goal "Julia prepared".*

5. **DISSUADE_CSQ(Lili,Frank, provide medical details(Frank))**

   Lili: "No no! If you give Julia all medical details, she will certainly be disgusted!" The consequence related to the dissuasion is the side-effect "disgusted".

6. **PERFORM(provide medical details(Frank)):**

   – obstacle "disgusted" triggered

   Frank to Paul: "My father had an accident. His head was hit, and after several operations in the brain, he could not fully recover. His temporal lobe is damaged. Now, my father suffers from several troubles, including loss of memory, change of mood, etc.". Julia makes a disgust expression and answers "Ah... yeah".

7. **PERFORM(correct homework(Frank)):**

   – obstacle "Paul interrupts" triggered

   Frank to Julia: "Come on Julia, let's see your homework, I would like to help you on this". Julia seems ok with this but Paul comes to her and declares: "We are just about to have dinner. I don't think we have time for you. Please leave our house!".

8. **COMPARE(Lili,Frank, divert(Frank), rebuff Paul(Frank))**
   Lili to Frank: "If you want to get rid of Dad's annoying remarks, you could distract him with something, but, if you have not explained the situation to Julia, she might not appreciate. You could also rebuff Dad, but he could become angry and say bad things to Julia.

9. **PERFORM(divert(Frank)):** obstacle "fails" triggered
   Frank to Paul: "Dad, well, would you go to the grocery store and buy something for dinner?".
   Paul: "Oh no, I prefer staying here, with you and your friend."

10. **PERFORM(rebuff Paul(Frank)):** success
    Frank to Paul: "Dad, that's enough! Could you please leave us alone?!"
    Paul: "Ok ok, sorry."

11. **PERFORM(correct homework(Frank)):** success
    Frank to Julia: "Come on Julia, let's see your homework, I would like to help you on this". Julia: "Great!". Five minutes later, Julia appears much friendly to Frank: "Thank you for your help, I appreciate it." Compared to action 7, no obstacle is triggered: Paul is "satisfied".

12. **PERFORM(get math book(Frank)):** obstacle "Paul interrupts" triggered
    Frank to Julia: "I would like to get you a great math book that will help you". He walks towards his bedroom to get the math book, but Paul comes to him and declare: "Frank, you can't leave a pretty friend alone in the living room like that!".

13. **PERFORM(rebuff Paul(Frank)):**

    – obstacle "Paul angry" triggered

    Frank to Paul: "Dad, that's enough! Could you please leave us alone?!" Paul: "What? You want to stay with this stupid friend of yours, while your father needs you?!"
    Comment: *Julia gets vexed by Paul's remark.*

14. **DELEGATE(Frank, Olivia, Julia prepared)**
    Frank to Olivia: "Grandma, I need your help. Would you please help me to explain Julia about Dad and our family?"

The above example shows a mixture of various action types. They enable to extend the expressiveness of a core narrative structure, by both producing narrative content and emphasizing the underlying paradox in the structure. The corresponding and tentative trajectory in the paradox space is represented in Figure 7.

Figure 7 Tentative trajectory in the paradox space, for the story described in Table 3.5.

4 Conclusion and future work

In this paper, a novel computational model of narrative has been proposed, which comprises several components: a hierarchical model of narrative structure, a set of rules to select the next action, an algorithm to estimate the dramatic tension based on the concept of paradox, and a set of rules to select the ending. The two main contributions of this model are 1) the hierarchical nature of the model, which enables to model complex structures without loosing a global view of the main "mechanics" of the narrative (its paradox) and 2) a novel approach for modeling the dramatic tension, which renews the more usual approaches that are based on conflict or suspense.

This computational model has been only partially described above, as many components need to be more formally described: the specification of the trajectory into the paradox space, the assessment of the dramatic tension, the action selection mechanism and the endings' specification and management. In order to successfully specify these components and refine the model, the next step will consist in paper prototyping the system, according to a methodology that has proved efficient in game design [15]. In the context of interactive narrative, the model along with the proposed story will be played by the user, while the narrative management will be carried out by a "game master", who will simulate the future narrative engine. Sessions will be recorded and later analyzed.

The story described in Section 3, while going beyond the level of a "toy" example, remains relatively simple. Once the model has been validated with this story, several improvements of the model will be considered: more than one paradox in a story (e.g., one paradox per main character), more than two levels in the hierarchal structure, parametrization of narrative elements, more variations in endings.

References

1 Ruth S. Aylett, Sandy Louchart, João Ferreira Dias, Ana Paiva, and Marco Vala. FearNot! – an experiment in emergent narrative. In Themis Panayiotopoulos, Jonathan Gratch, Ruth Aylett, Daniel Ballin, Patrick Olivier, and Thomas Rist, editors, *Intelligent Virtual Agents, 5th International Working Conference, IVA 2005, Kos, Greece, September 12–14, 2005, Proceedings*, number 3661 in Lecture Notes in Computer Science, pages 305–316. Springer, 2005.

2 Paul Bailey. Searching for storiness: Story-generation from a reader's perspective. In M. Mateas and P. Sengers, editors, *Narrative Intelligence: Papers from the AAAI Fall Symposium*, number FS-99-01 in AAAI Technical Reports, pages 157–163. AAAI Press, 1999.

3 Daniel Balas, Cyril Brom, Adam Abonyi, and Jakub Gemrot. Hierarchical Petri nets for story plots featuring virtual humans. In Christian Darken and Michael Mateas, editors, *Proceedings of the Fourth Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*, pages 2–9. The AAAI Press, 2008.

4 Heather Barber and Daniel Kudenko. Dynamic generation of dilemma-based interactive narratives. In Jonathan Schaeffer and Michael Mateas, editors, *Proceedings of the Third Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, pages 2–7, Menlo Park, CA, 2007. AAAI Press.

5 Raphaël Baroni. Incomplétudes stratégiques du discours littéraire et tension dramatique. *Littérature*, 127:105–127, 2002.

6 Cristina Battaglino and Rossana Damiano. Emotional appraisal of moral dilemma in characters. In David Oyarzun, Federico Peinado, R. Michael Young, Ane Elizalde, and Gonzalo Méndez, editors, *Interactive Storytelling, 5th International Conference, ICIDS 2012, San Sebastián, Spain, November 12–15, 2012. Proceedings*, number 7648 in Lecture Notes in Computer Science, pages 150–161. Springer, 2012.

7 Claude Bremond. *Logique du récit*. Seuil, Paris, 1973.

8 Noel Carroll. *Beyond Aesthetics: Philosophical Essays*. Cambridge University Press, Cambridge, 2001.

9 Mario Cataldi, Rossana Damiano, Vincenzo Lombardo, and Antonio Pizzo. An agent-based annotation model for narrative media. In Wiebe van der Hoek, Lin Padgham, Vincent Conitzer, and Michael Winikoff, editors, *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages 2010–2012. ACM Press, 2012.

10 Marc Cavazza, Fred Charles, and Steven J. Mead. Characters in search of an author: AI-based virtual storytelling. In Olivier Balet, Gérard Subsol, and Patrice Torguet, editors, *International Conference on Virtual Storytelling (ICVS 2001)*, number 2197 in Lecture Notes in Computer Science, pages 145–154, Heidelberg, 2001. Springer.

11 Yun-Gyung Cheong and R. Michael Young. Narrative generation for suspense: Modeling and evaluation. In Ulrike Spierling and Nicolas Szilas, editors, *Interactive Storytelling, First Joint International Conference on Interactive Digital Storytelling, ICIDS 2008, Erfurt, Germany, November 26–29, 2008, Proceedings*, number 5334 in Lecture Notes in Computer Science, pages 144–155. Springer, 2008.

12 Jean E. Dumas, Nicolas Szilas, Urs Richle, and Thomas Boggini. Interactive simulations to help teenagers cope when a parent has a traumatic brain injury. *Computers in Entertainment*, 8(2), 2010.

13 Lajos Egri. *The art of dramatic writing*. Simon & Shuster, New York, 1946.

14 Syd Field. *Screenplay? The Foundations of Screenwriting*. Dell Publishing, New York, 1984.

15 Tracy Fullerton. *Game Design Workshop: A Playcentric Approach to Creating Innovative Games*. Morgan Kaufmann, 2008.

16 Patrick Gebhard, Gregor Mehlmann, and Michael Kipp. Visual SceneMaker – a tool for authoring interactive virtual characters. *Journal on Multimodal User Interfaces*, 6(1-2):3–11, 2011.

17 Algirdas Julien Greimas. *Sémantique structurale*. Presses Universitaires de France, Paris, 1966.

18 Nicolas Habonneau, Nicolas Szilas, Urs Richle, and Jean Dumas. 3D simulated interactive drama for teenagers coping with a traumatic brain injury in a parent. In David Oyarzun, Federico Peinado, R. Michael Young, Ane Elizalde, and Gonzalo Méndez, editors, *Interactive Storytelling, 5th International Conference, ICIDS 2012, San Sebastián, Spain, November 12–15, 2012. Proceedings*, number 7648 in Lecture Notes in Computer Science, pages 174–182. Springer, 2012.

19 Michael I. Jordan and Robert A. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6(2):181–214, 1994.

20 Martin Klesen, Janek Szatkowski, and Niels Lehmann. The black sheep: Interactive improvisation in a 3D virtual world. In Giorgio De Michelis and Ulrich Hoppe, editors, *Building Tomorrow Today: Community, Design and Technology, I3 Annual Conference, Jönköping, 13-15 September 2000, Papers, Interactives Performances, Exhibiitons*, pages 77–80, 2000.

21 Arthur Koestler. *The Act of Creation*. Pan Books, London, 1964.

22 Yves Lavandier. *La dramaturgie*. Le clown et l'enfant, Cergy, France, 1997.

23 Claude Lévi-Strauss. *Anthropologie structurale*. Plon, Paris, 1958.

24 Elias J. MacEwan. *Freytag's Technique of the Drama: An Exposition of Dramatic Composition and Art*. Scott, Foresman and Company, Chicago, third edition, 1900.

25 Michael Mateas and Andrew Stern. Integrating plot, character and natural language processing in the interactive drama façade. In Stefan Göbel, Norbert Braun, Ulrike Spierling, Johanna Dechau, and Holger Diener, editors, *Technologies for Interactive Digital Storytelling and Entertainment, TIDSE 03 Proceedings*, number 9 in Computer Graphik Edition, pages 139–151, Darmstadt, 2003. Fraunhofer IRB.

26 Robert McKee. *Story: Substance, Structure, Style, and the Principles of Screenwriting*. Harper Collins, New York, 1997.

27 James Meehan. TALE-SPIN. In Roger C. Schank and Christopher K. Riesbeck, editors, *Inside computer understanding: Five programs plus miniatures*, pages 197–226, Hillsdale, NJ, 1981. Erlbaum.

28 Bill Nichols. *Ideology and the image*. Indiana University Press, Bloomington, IN, 1981.

29 Gerald Prince. *A Dictionary of Narratology*. University of Nebraska Press, Lincoln, NE, 1987.

30 Marie-Laure Ryan. Introduction. In Marie-Laure Ryan, editor, *Narrative Across Media*, pages 1–40, Lincoln and London, 2004. University of Nebraska Press.

31 Nikita Sgouros. Dynamic generation, management and resolution of interactive plots. *Artificial Intelligence*, 107(1):29–62, 1999.

32 E. Souriau. *Les deux cent mille situations dramatiques*. Flammarion, Paris, 1950.

33 Nicolas Szilas. Interactive drama on computer: Beyond linear narrative. In M. Mateas and P. Sengers, editors, *Narrative Intelligence: Papers from the AAAI Fall Symposium*, number FS-99-01 in AAAI Technical Reports, pages 150–156. AAAI Press, 1999.

34 Nicolas Szilas. A new approach to interactive drama: From intelligent characters to an intelligent virtual narrator. In John Laird and Michael van Lent, editors, *Proceedings of the AAAI Spring Symposium on AI and Interactive Entertainment*, number SS-01-02 in AAAI Technical Reports, 2001.

35 Nicolas Szilas. A computational model of an intelligent narrator for interactive narratives. *Applied Artificial Intelligence*, 21(8):753–801, 2007.

36 Nicolas Szilas. Requirements for computational models of interactive narrative. In Mark Alan Finlayson, editor, *Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium*, number FS-10-04 in AAAI Technical Reports, pages 62–68, Menlo Park, California, 2010. AAAI Press.

37 Nicolas Szilas, Urs Richle, and Jean E. Dumas. Structural writing, a design principle for interactive drama. In David Oyarzun, Federico Peinado, R. Michael Young, Ane Elizalde, and Gonzalo Méndez, editors, *Interactive Storytelling, 5th International Conference, ICIDS 2012, San Sebastián, Spain, November 12–15, 2012. Proceedings*, number 7648 in Lecture Notes in Computer Science, pages 72–83. Springer, 2012.

38 Nicolas Szilas, Urs Richle, and Paolo Petta. Structures for interactive narrative: An authored-centred approach. Technical Report TECFA12-1, TECFA-FPSE, University of Geneva, 2012.

39 Eugene Vale. *The technique of screenplay writing*. Grosset & Dunlap, New York, 1973.

40 Stephen G. Ware, R. Michael Young, Brent E. Harrison, and David L. Roberts. Four quantitative metrics describing narrative conflict. In David Oyarzun, Federico Peinado, R. Michael Young, Ane Elizalde, and Gonzalo Méndez, editors, *Interactive Storytelling, 5th International Conference, ICIDS 2012, San Sebastián, Spain, November 12–15, 2012. Proceedings*, number 7648 in Lecture Notes in Computer Science, pages 18–29. Springer, 2012.

41 Paul Watzlawick, Janet B. Bevelas, and Don D. Jackson. *Pragmatics of Human Communication*. W. W. Norton & Company, New York, 1967.

42 Peter Weyhrauch. *Guiding Interactive Drama*. PhD thesis, Carnegie Mellon University, 1997.

43 R. Michael Young, Mark O. Riedl, Mark Branly, Arnav Jhala, R. J. Martin, and C. J. Saretto. An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development*, 1(1):51–70, 2004.

Writing Consistent Stories based on Structured Multi-Authored Narrative Spaces

Alan Tapscott, Joaquim Colàs, Ayman Moghnieh, and Josep Blat

Interactive Technologies Group
Universitat Pompeu Fabra
Barcelona, Spain
`{joaquim.colas,alan.tapscott,ayman.moghnieh,josep.blat}@upf.edu`

───── **Abstract** ───

Multi-authoring is currently a common practice in the field of contemporary storytelling but producing consistent stories that share a common narrative space when multiple authors are involved is not a trivial task. Inconsistencies, which are not always well-received by readers are sometimes expensive to fix. In this work we attempt to improve the consistency of stories and narrative spaces by introducing a set of rules based on a formal model. Such a model takes into account the reader's concept of consistency in storytelling, and acts as a framework for building tools to construct stories grounded in a common narrative space with a reinforced sense of consistency. We define a model (the Setting) and deploy it through a tool (CrossTale); both based on previous research, and discuss some user evaluation, with an in-depth analysis of the results and their implications.

1 Introduction

The evolution of digital interactive media and information technologies has been instrumental in the development of systems that bring together authors and readers to compose and consume multi-authored stories though multiple media. In this context, the audience is not only interested in rich narratives, but also wants to participate in their development by adding and sharing their very own creations, compositions, and ideas. Nowadays people actively publish and share thousands of creative works (blogs, stories, songs. . .) on the web, often related to other original creations through relations that range from mere inspiration to direct referencing. Some of the works may be further developed by more authors, who expand their content, structure, and knowledge value through original creation and composition processes. On the other hand, there is an emerging interest to support collaborative creation, composition, and consumption of multi-authored narratives that may grow in a shared information space for prosumers and professionals alike. We use a basic definition of information space: "The set of concepts and relations amongst them held by an information system" [14]. We believe narrative spaces are information spaces that ground all media based on the same characters, situations, plots or other casually interlinked entities, hence introducing a certain degree of consistency to the set as a whole. Narrative spaces are especially worth analyzing when dealing with collaborative storytelling since they establish many of the rules for the interaction among authors. The authors' awareness and interpretation of the narrative space will heavily

condition their interaction with it. Fans often expand the narrative space of their favorite entertainment franchises by introducing their own stories deeply rooted in a well-established narrative universe and its mythology, creating rich networks of fan-fiction (Fanfiction.net, referenced later, gathers hundreds of thousands of users around hundreds of franchises) that coexist with the official material. Also, Web and information technologies provide momentum to complex entertainment franchises created by dozens of authors to span across multiple media. In this context, there are certain tools that support readers and writers who are contributing to well-established narrative spaces. Articy Draft [1] and Celtx [5] are both collaborative tools meant for creative story development and represent good examples of this emerging trend. Such tools may be created by the same company delivering the content, but this content is often the result of a collaborative effort undertaken by an author community. It is worth noting that tools do not merely intend to support the construction of a story in the sequential, traditional way. They provide mechanisms that allow for free, divergent exploration of all the related information, supporting the non-linear growth of narrative spaces.

There are some examples of narrative spaces worth mentioning. Fanfiction web site [8] is devoted to fan-developed stories within the narrative space defined by specific franchises, and provides a good example of amateur and professional authors creating stories in the same narrative space. Most of these stories, however, do not take into account the contributions of their fellow fan authors, only the original, canonical one. Another example is the website that holds most the information related to the A Song of Ice and Fire book series in its many articles [19]. It is fed with content from multiple contributors, properly structured and published in a readable way. The site also publishes articles that cover most of the books related to the canonical narrative space and a text-based roleplaying game that allows players to introduce their own creations (e.g., characters, locations, and other elements around the original canonical narrative). Players can interact with each other while expanding the original setting. This site has the approval of the author of A Song of Ice and Fire who is known to be vocal against common fan fiction developed without consent. On the other hand, he created Wild Cards [12], a book series written by multiple authors under his editorial control. Chris Crawford's Storytron [7] is an interesting approach to developing a commercial tool that would allow users to design interactive stories. Although it is currently on-hold due to problems regarding the learning curve (i.e., the complexity of building a whole interactive story with the tool), this approach is interesting in terms of decomposing the narrative space into a set of unitary elements, and defining the logic that relates them. Storyjacker [10] is another interesting example closely related to the tradition of the Exquisite Corpse writing technique. This game proposes that its players first read a flash fiction (roughly between two or three hundred words) created by another writer with an explicit editorial challenge attached to it. Players rewrite the text answering the challenge and pass the result to the next player, introducing a new challenge of their own. While this approach is a game, the writing dynamics of its multi-author design are interesting and not very far away from what we propose in this paper.

In this paper we try to understand if there are people interested in writing stories collaboratively in a consistent way and provide them with an appropriate tool for that purpose. First we discuss our focus on enhancing consistency, especially how it is perceived by authors and readers, followed by a brief state of the art of previous research on multi-authored narratives for similar scenarios. Next we describe some users' experiments we conducted. These experiments were designed to test mechanisms developed to increase narrative consistency. We then analyze and discuss the resulting experimental data. Finally we discuss these findings in relation to the approach proposed and introduce future research.

2   Supporting Narrative consistency

Complexity can easily scale with the developing size of narrative spaces, possibly increasing the difficulty of reading and authoring stories based on them. Each element in a narrative space, such as a character, location or event, is linked to other elements in the same narrative space through causal relations, providing a sense of continuity and consistency. Modifications introduced to the narrative space may cause contradictions in the logic of the network of elements and causal links. This often leads to plot holes that may compromise the story's global consistency, sometimes invalidating the primary causal links that represent the foundation of fundamental plot threads, and potentially hindering the experience of authors and consumers. Stories containing plot holes also tend to have a bad reception amongst sophisticated readers [17].

If consistency is a key factor when dealing with multi-authored storytelling, some sort of mechanism designed to monitor and enhance its presence could result in a better experience for its readers. This work pursues a suitable method to assist multiple authors in developing narrative spaces with enhanced consistency. This might lead to stories which are more satisfactory to develop collaboratively in these narrative spaces and are also more enjoyable to read. When analyzing narrative spaces and their unfolding stories, we distinguish between two kinds of consistency measurements:

- Firstly structural as the level of agreement among the elements of the narrative space with respect to each other. This can be measured if the narrative space is mapped to a computational structure of some sort by validating the narrative space information against a formal model.
- Secondly reader-perceived as the level of consistency associated by readers to a specific story. This is most often obtained by asking readers to rate it after having read it.

We think this distinction is necessary because of the subjective nature of some stories along with the existence of some literary techniques, such as the use of biased narrators that describe reality through perception and language. Having two different measures of consistency is invaluable when trying to relate both kinds of consistency. By analyzing the content of a narrative space and mapping it to a computable and evaluable structure, we can provide some recommendations or guidelines to increase the structural consistency of a narrative space. To some extent, starting with Propp and his structural approach to narrative [16], the field of semiotics is grounded on similar principles and has been an active discipline for decades. Its theoretical foundation, specifically the syntactic branch that deals with formal structures, has been a source of inspiration for our work. Deconstructing a narrative space into a computational structure based on a suitable model can be a challenging discretization process. We do not propose a model that attempts to do this. Instead, the model we propose is based on observations regarding the author and reader perception and interpretation of consistency. Every author has a personal way to tell stories. This means that the perception of a story's consistency depends on the technique and structure of its discourse – not to mention the influence of genre. Readers may find a story consistent or inconsistent regardless of the raw material from the narrative space used by the author. Also every reader's perception is heavily influenced by factors such as his/her cultural, academic and social background, which can be difficult to control and keep track of. The most obvious way to measure the user-perceived consistency of a specific story is to ask different readers to rate it. There are other more indirect methods, such as asking specific questions to check if the reader understands the story or to observe the reading procedure, trying to encode it into meaningful data. We have found these measurements difficult to operationalize and correlate to the reader-perceived consistency level. Our goal in this research is to determine

whether monitoring and enhancing the structural consistency of a narrative space implies that stories based on it are perceived as more consistent by readers.

3   Related Works

We now analyze related literature and discuss its implications for our goal. Meehan's TaleSpin is a system that generates stories via carefully crafted processes that operate at a fine level on story data [13]. It was one of the first attempts to model narratives as computational systems. Since it automatically generates stories, it holds a certain notion of computational causality and consistency. We also pursue a formal model with such notions, but Meehan's approach seems too constraining to support an open definition of a story. Brenda Laurel's doctoral dissertation described a complex framework for drama management [11] and is considered by some as the beginning for the many successful approaches that deal with structured narrative spaces. While it is meant for abstract depictions of large narrative spaces, it also provides a systematic representation for them. A key factor is its ability to introduce highly dynamic narrative structures. These structures support complex stories that hide the formal complexity from readers, something we wish to introduce in our approach. Thue [18] proposes an interesting approach that formally structures the story, favoring consistency monitoring and analysis. Player Modeling is a simple concept that attempts to personalize the story through several profiling techniques, enabling some of the user's personality traits to have certain impact on the resulting experience. Understanding the reader's perception of consistency is a concern we share. Some other approaches use a strictly formal definition to model stories. For instance, Cavazza proposed a character-based approach [4] that was adapted and improved by Pizzi to model a part of Madame Bovary [15]. This line of work is grounded on planning and the field of artificial intelligence. Interestingly enough, it deals with complex aspects of human nature such as emotions and feelings. The AI planning used in [15] is concerned with optimality, seeking to reach a target with economic operations and may not be adequate for our approach. We believe storytelling should encourage causal links, but not necessarily in an optimal way. They represent, however, some of the most intricate and complete attempts to discretize the narrative structure into a formal model, a goal we also pursue. Next we discuss some existing formats and recent tools that allow modeling narrative entities independently from their story, that keep track of the flow of complex events, that impose constraints or rules to preserve consistency, that keep track of plot meta-data (such as character motivations, feelings or the literary theme and mood), and that are suited for collaborative development of a story. This discussion inspired the conception of our tool.

- Traditional scripts are often created by a single or a couple of authors. Large media franchises and episodic shows sometimes need to become heavily interrelated. Game of Thrones [2] is a good example of a TV show that has heavily interrelated scripts written by multiple authors. To some extent they represent one of the most popular instances of a multi-authored narrative with a strong need for consistency.
- There is a certain tradition of background books in rich fiction series, providing concept art, character profiles or even maps depicting fictional lands. These books, far from narrating a story in the traditional sense, describe a specific part of a fictional universe. We found these works interesting because they represent a set of characters, themes and plots in their original, protean form, not necessarily attached to the linear context of a traditional tale. They are often written by authors who were not creators of the original concepts, and represent an example of collaborative authoring.

Table 1 Multi-Authoring Narrative Supports Comparison.

	Traditional Script	Background Book	P&P RPG Source Book	Wikipedia/Wikia
Atomic Narrative Elements	Not formally	Yes, clearly differentiated	Yes, clearly differentiated	Yes, clearly differentiated
History Log	Sequence is implicit on its description	Yes, mostly inside individual element descriptions	Yes, mostly inside individual element descriptions	Yes, mostly inside individual element descriptions
Consistency Constraints	No	No	Yes, enforced by the game's rules	No
Plot Meta-Data	No	Yes, mostly inside individual element descriptions	Yes, within plot meta-data	Yes, only attached to individual element descriptions
Suited for Collaborative Developement	No	No	No	Yes

- RPG books, such as ADD Monster Manual [9] or Vampire: The Requiem Coteries [3] are interesting examples of narrative entities modeled independently. They provide a growing organic framework for authors to build their own adventures and share them with friends, adopting the role of a live storyteller in tabletop gaming sessions. The source material in these books can be used to enrich the session experience by introducing new characters, object or plot threads. While fairly similar to background books, RPG books provide guidelines that allow content to be used in the arbitrary context of a game with rules, which introduces a high degree of formality to the information.

- Certain tools such as Wikipedia or Wikia are effective means of storing and organizing data from a specific narrative space. Although they are commonly used to structure already-existing background information, they represent some of the most popular tools that support collaborative writing. Their capability to deal with individual entities such as characters or locations is the trait we find more interesting. On the other hand, entities commonly depicted as linear, such as stories or plot threads, are not very intuitive to understand and follow using these tools. As shown in table 1, most of them possess some of the traits we introduced earlier, but no tool has got all of them as far as we know. Incorporating existing mechanisms that seem appropriate is part of our efforts to design a tool with all these traits.

4 Experiments

We carried out three experiments to understand better narrative spaces and the stories based on them in terms of user perception. For each of them we introduce its purpose, any tool specifically designed for it, the experimental design in depth and the most significant results.

4.1 Experiment I: Understanding the Sharing of Narrative Spaces

Our first experiment aimed at understanding how users perceive a narrative space and its associated stories while contributing and navigating through it. We intended to understand their mental model and to measure it. Some arbitrary conventions were introduced, such as an initial set of scenes already connected or a limited set of characters and objects. This was done to encourage participation, providing a certain sense of narrative immersion and to reduce the creativity required from subjects in order to participate. A fairy tale was chosen,

■ **Figure 1** A Story Wall.

including its most canonical elements (e.g., a king, a Princess, a castle, and a dragon among others) along with others far away (e.g., a robot, aliens, and a starship among others).

4.1.1   A Collaborative Story Wall

A large glass wall was used as a space to develop and visualize a collaborative narrative (figure 1) composed of scenes and transitions. Our purpose was to provide a canvas for authors to freely interact with the story. The scenes were sheets of paper with a collage of images (obtained by mixing characters and props picked from a set) and text written to describe the scene more explicitly. Scenes could be added anywhere on the wall and connected by transition arrows drawn on the glass, as a directional indicator, providing a sequential order by connecting scenes. The experimenters provided an initial story as a starting point for users who, in succession, could modify what was on the wall: change or delete scenes, alter the structure (erasing and drawing transition arrows, and moving scenes to new positions), and place their own scenes in any point of the unfolding story. We introduced 7 initial scenes narrating the beginning of the kidnapping and rescue of the Princess.

4.1.2   Experiment I in Detail

16 subjects were invited to participate in the experiment one after another sequentially. There was no special consideration in the demographics involved. A non-imposed average elapsed time of 12 minutes was measured.

Subjects were asked to read the existing narrative which was the result of the accumulative modifications made by previous subjects on the initial set provided. They were also interviewed after they finished reading the existing narrative on how they've had chosen to read the story (order, objects and concepts they had followed, etc.), along with their opinion on some specific matters such as the literary value and consistency perceived.

Next they were offered the possibility to contribute to the narrative, and allowed to modify or delete previous scenes, to alter the structure of the story structuring (erasing and drawing transition arrows, and moving scenes to new positions), and to place their own scenes at any desired point. Finally all subjects answered a series of questions designed to learn more on how they interacted with the story, such as the nature of their contributions (according to them) along with their driving motivation or purpose. We also asked some open questions on some subjects such as if it was a fun experience or if they would enjoy doing the same with their friends through a social network.
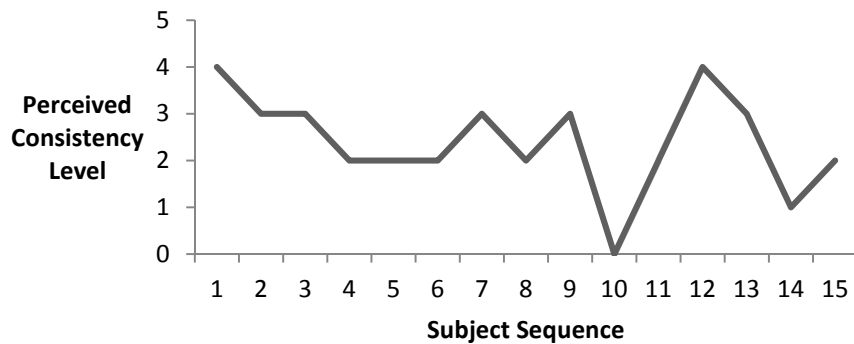
> **Figure 2** Story on a Wall reader-perceived consistency level.

4.1.3 Results

The story resulting from the experiment contained 29 scenes connected through two main branches that converged towards their end. Each participant added either one or two scenes to the growing narrative. No subject eliminated scenes from previous participants, but modifications on existing scenes were common: half of the participants inserted their scenes between existing ones and/or altered the direction of arrows; over one third created convergence between two or more isolated branches (for example two characters gathering at one point, or one event affecting the story of another author). A few subjects claimed to focus exclusively on solving inconsistencies during the authoring phase of the experiment. Maintaining consistency in the evolving narrative was stated as the principal reason for 8 out of the 15 contributors. The notion of conflicting scenes was stated 4 times as something disliked in the interviews. According to subjects all of the changes made to previously existing elements were for the sake of consistency. Other contributions were centered mainly on extending existing plot arcs instead of creating new ones. Consistency seemed to be key in user motivation and overall experience. The subject-perceived level of narrative consistency (figure 2) tends to be on the middle-high portion of the scale but decays slowly. As the initial story is different for each user, the results cannot be easily compared but subsequent experiments allow for comparison. In this experiment we were mainly interested in observing the interaction between authors and the story.

According to the interviews, the literary value of the narrative concerned little the subjects. Interestingly, individual scenes and small narrative branches had greater entertainment value than the overall narrative. Since the sequence of events can only be guessed through the spatial layout of the scene and the arrows network, some conflicting notions appeared on what was happening before, after, or simultaneously to a given scene when dealing with parallel stories. This suggested that scenes could be arranged in some sort of linear organizational structure to provide an improved sense of sequence and causality. Our close observation of how scenes related with each other and how participants authored existing characters, revealed that each character was considered the same entity throughout the whole narrative, almost always labeled with the same name. The experiment also showed that the authors faced a complexity which scaled very strongly if they tried to maintain the structural consistency of the story. The more scenes it contained, the harder it was to introduce new material without contradicting or violating existing established facts. On the other hand, the decreasing reader-perceived consistency of stories containing a large amount of scenes indicated that the reading process became more difficult as well. Some people were motivated by the unfolding implicit collaboration, and nobody stated openly to be bothered by it. In fact,

contributing to the narrative was not mandatory but all of the subjects added scenes, and they actively searched for an interesting entry point and modified the whole context, changing and rearranging scenes connected to their contributions, instead of just attaching them to the end of a story thread. More than half of the subjects expressed their interest in repeating the process later and many of them returned after their contribution to see how the narrative was evolving. A good number of people who just happened to pass by stopped to read the whole story, many of whom asked to participate in subsequent iterations of the experiment.

4.2 Experiment II: Measuring the Impact of Consistency Constraints

The purpose of the second experiment was to measure the impact of an underlying formal model to user contributions and their overall interaction with a multi-authored non-linear narrative. This formal model was designed to provide structural consistency to the narrative space, hopefully reinforcing key factors that enhance the production of stories that are perceived as more consistent. We introduced some constraints into the interaction to prevent subjects from creating scenes that somehow violated the rules proposed by the underlying model. We used a platform we developed, [6], to be used on a connected laptop, which meant changing to a much more private environment.

4.2.1 A Setting that provides an Underlying Consistency Model

The Setting tries to provide an underlying formal model that resembles the author's mental construction of a narrative space. We used data from the previous experiment to map their understanding of the story into an assessable and measurable model, through a process that can be found in our previous publication [6]. The Setting serves to monitor and enhance the consistency of the stories unfolding within it. It provides a common ground for authors to interact by building stories in the same narrative space. Its informal definition is the following:

- The Setting contains timeframes and locations on a grid.
- Timeframes have an implicit order.
- Every location is at a certain distance of other locations. The distance from A to B is the minimum number of locations needed to go from A to B.
- Every scene takes place in a location and contains one or many characters and zero or many objects.
- Scenes can belong to plot storylines or character storylines.
- Storylines contain one or many scenes.
- Character storylines contain all the scenes that contain a specific character.
- Plot storylines contain all the scenes tagged to design a specific plot.
- Characters may only appear once per timeframe in a scene.
- Characters may only appear once in the same scene.
- Characters may not move between non-adjacent locations (distance > 1) in a single timeframe.

These rules were designed to provide a certain sense of consistency, which can be measured, monitored and enhanced, on the basis of the results of the Story on a Wall experiment, attempting to predict and enforce the factors actively pursued by users through their contribution. Our goal was not to evaluate this definition as a generalist model capable of describing any narrative; instead we wished to measure the impact of using a formal model in a multi-authoring scenario in terms of the consistency of the resulting stories.

Figure 3 CrossTale Interface.

4.2.2 Introducing the platform

CrossTale (see figure 3) is a software prototype whose main window follows a distribution similar to that of Story Wall, adding the rules imposed by the Setting. In fact, its main context is a dashboard with two axes, one for time and one for place. Users can scroll at will to navigate the dashboard. By selecting existing scenes they can view their images and read the descriptive texts. Specific characters and storylines can be selected, enabling users to read all the scenes involving that character or storyline in a sequential order. The grid also highlights scenes belonging to the selected entity and connects them with an arrow line to reflect their sequential order. There is also a secondary context that enables users to create scenes, providing a set of components (characters, objects and plot storyline tags) along with a visual representation of the location where the scene takes place and a text box to introduce the description. These scenes are added to one of the Setting timeframes and locations and are treated as an integral part of the narrative space. Violations of the Setting were not allowed in this experiment, and the user got a message requesting him/her to resolve the conflict before saving the scene (in the third experiment users could save scenes that violated the Setting, but users were warned before).

4.2.3 Experiment II in Detail

20 subjects of similar characteristics as those in the first experiment took part. Two groups of 10 were created randomly. The control group used the tool to read and contribute to the existing narrative, and the experimental group had some consistency constraints based on the Setting. The use was sequential, as each user found the story in the situation left by the previous one. No time limits were provided and the average time of the users was 20 minutes. A CrossTale prototype was created with an initial set of scenes describing the start of a traditional fairy tale. The 8 initial scenes introduced were almost identical to the ones used in the previous experiment, introducing a Princess, her kidnapping by a witch and the Prince trying to rescue her. Each subject was asked to read the story which was composed by the initial scenes provided plus the contributions made by previous subjects. No specific method
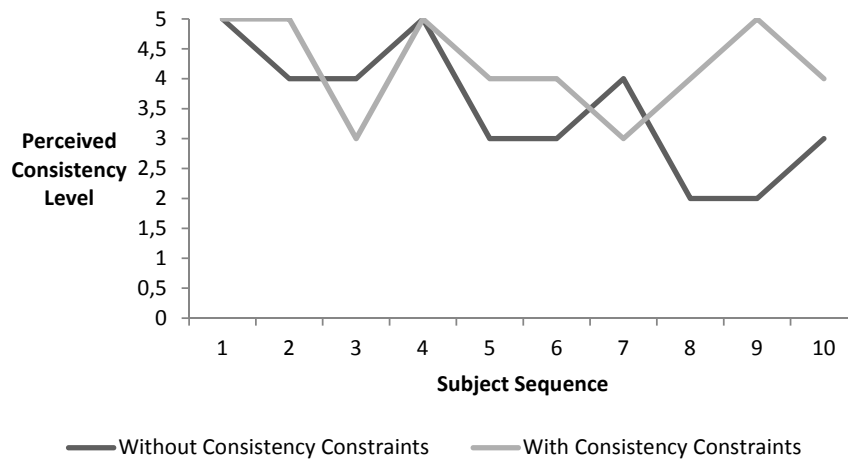
■ **Figure 4** CrossTale reader-perceived consistency level with and without consistency constraints.

was imposed. The story could be read completely or not. Reading character storylines could be a strategy amongst others. A brief interview was conducted to understand how users read and understood the whole scene set regarding the storylines. Then they were asked to add one or more scenes to the existing ones. After they were done, a second interview was conducted to understand what kind of additions and modifications they had made, their motivations, the intended influence on the previous state of the story, and any other relevant details of the interaction between the subjects and the story. The whole experiment was recorded for further coding and observations. The subjects were aware of the collaborative nature of the tool, but did not have contact with the rest of the subjects before, during or after the experiment.

4.2.4 Results

Results were analyzed independently for each group. It appears that subjects were not very concerned with reading the whole narrative before interacting with it. Users only read a fraction of the existing content. No user read the whole story. The most common interaction recorded during the reading phase involved the user selecting one or two storylines and reading its content before moving on to the contribution phase. The perceived consistency (figure 4) was steadily rated high in both groups, with a slight tendency to decay towards the end in the group without constraints. The difference did not seem very significant. Both groups ended up with a story composed of 28 scenes and 10 storylines. The average scene contribution was 2 scenes per user. Most users placed their scenes inside one and only one storyline. No user modified scenes created by other authors. The rating of the user experience was positive (average 4.4 out of 5) as well as of the application design (average 4 out of 5). We asked subjects if they would use CrossTale regularly with an average 3.6 out of 5 and if they would like to have a similar tool to create and share narratives in the context of a social network, with an average 3.7 out of 5.

The focus of the experiment was to observe if the introduction of consistency constraints derived from the Setting caused any interesting effects. The most remarkable observation was that the perceived consistency seemed to decay more quickly over each contribution for the group without constraints, although the resulting data isn't very significant. This could mean that enforcing certain notions of time and space through the scenes tends to produce more consistent results, supporting our initial hypothesis. A larger subject group

in future experiments could validate or refuse this claim. Adding the constraints seems to have an annoying effect on the experience of users who felt limited all the time (as seen during the video codification, where they complained almost every time a constraint blocking message popped up). This might be caused by the way messages themselves are displayed in CrossTale. It could be an interesting line for future research. Joining the data from both groups also revealed some interesting facts. The use of a computer program to conduct the experiment might have affected the user experience, limiting the user's freedom when compared to the previous experiment. The story in this experiment was read on a screen and embedded inside a software program instead of being on a glass wall. Subjects were less inclined to interact with the existing scenes; no user modified scenes created by other authors. Subjects spent less time interacting with the narrative (the decreased time could either be an indicator of a less pronounced learning curve, a good interaction design or a decrease in the motivation of subjects). Also, according to the interviews, they were less concerned by narrative inconsistencies. As previously mentioned, the story was now stored in a computer program. We believe this might have caused users to be less aware of the story as a whole and therefore less concerned with its global consistency. In fact, the reader-perceived consistency of the narrative was larger for both groups of users compared to the previous experiment. This might also be related to the fact that users never read the whole story. Users aren't concerned with the consistency of scenes they haven't read. We chose to follow the same cumulative mechanism as in the first experiment on both groups. This was done to gain some insight on the evolution and scalability of the story while comparing the results with the previous experience. We are aware that this decision prevents us from comparing subjects' individual performance in terms of consistency. The following section describes an experiment where this was done.

4.3 Experiment III: Measuring the Usage of Storylines

The third experiment explored the use of storylines further. Namely, we were interested in measuring certain aspects such as the number of storylines read by subjects, the degree of comprehension after reading, the performance when creating new storylines and their consistency. Moreover, we wanted to cross measures of the reading and contributing phases and find any significant correlations.

4.3.1 Experiment III in Detail

This experiment was fairly similar to the previous one. The main difference was that user contributions were not cumulative, every subject found the same initial set of scenes and there was only one group. Every subject started their contribution with the initial 12 scenes we provided. The story was the same fairy tale. The initial scenes introduced 3 main storylines that explained the events through the Prince, Princess and the witch's own viewpoints. CrossTale was used with the same rules derived from the Setting, the derived consistency constraints from the Setting were always active; its application was not enforced, only warning messages existed. There were minor usability refinements to CrossTale. We provided users with the ability to zoom in and out (using the mouse wheel) when viewing the scene grid. We also allowed users to scroll through the scene grid by dragging the mouse anywhere, not only the scrollbars. These additions were introduced to provide more visibility and accessibility to the existing scenes inside CrossTale. 16 subjects of similar characteristics as those involved in the previous experiments took part. An average time of 10 minutes of involvement with the system was measured. The experiment began with each subject

reading the story. CrossTale provided several mechanisms to do so: reading individual scenes, following specific storylines according to plot threads or characters. Users were free to read only a part if they wished. The interactions with the reading interface were registered, and a brief interview was conducted afterwards to analyze their reading experience. The next phase was the contribution. Every subject was asked to add more scenes to the same existing story if they wanted to. Their interaction was registered and a brief questionnaire was administered. This questionnaire was used to rate the user's general impression of the story when contributing to it. Subjects were asked to rate the warning messages, the story in terms of consistency and amusement through Likert scales, and also to propose one or more titles. In both phases the proceedings were run by a collaborator not directly involved in the research, who coded the interactions as well. Unlike the previous experiments, modifications to the scene set were not cumulative between subjects, so the consistency measurement was done through a 4 person jury evaluation of each subjects' contribution.

4.3.2   Results

Regarding the reading phase, most subjects read the existing scenes through the usage of storylines. Readers selected an average of 7.77 storylines to read. 83% of them were read from start to finish. 43.59% of the initial character storylines were read and 1.38% of the initial plot storylines were read. The average contribution per subject was 2.6 scenes. The number of scenes read seems to be correlated with the number of plot storylines used. There is a medium-high correlation between the number of titles for the story proposed by subjects and the number of characters mentioned in those titles. Also there's another medium-high correlation between the number of plot storylines referenced in the proposed titles and the amount of plot tags used later during the authoring phase. There's a positive correlation between the number scenes created, the number of storylines read and diversity of characters used in the created scenes. Very few message warnings about violations of the Setting rules were displayed (Warnings appeared in 24% of the composed scenes). Of these warnings, only 17% made the authors change the story. The resulting inconsistency level measured was an average of 1 inconsistency per contribution, or 0.46 inconsistencies per scene. Another interesting observation regarding consistency is the following; inconsistencies didn't increase in proportion to the number of scenes introduced inside a story.

Subjects seemed generally more inclined to use character storylines to read the provided story. There's a tendency towards a character-driven exploration of the story, possibly related to semiotics and some of its most popular theories. Nearly no subject read scenes without using storylines. We believe they proved to be a good mechanism to explore non-linear narratives such as the one we created in this experiment. Some users made extensive usage of the tool to create a large amount of scenes, which allowed us to briefly analyze the scalability of the system in terms of consistency. The number of inconsistencies remained stable during each user's session. In those cases, having the same author for all the contributions also ensured a more accessible and scalable development. We believe the small size of the initial narrative, along with the improvements and refinements to the CrossTale user experience were also instrumental for this to happen. This also could explain certain measurements, such as the average reduced time for each subject's interaction with the story. While these measurements might make it difficult to correlate the structural consistency of the narrative space with the consistency perceived by reader, the jury evaluation and our qualitative analysis of the stories suggest some major critical inconsistencies were avoided thanks to the warnings. Since we lack more evidence to sustain such a claim, we are already pursuing new experiments to provide more data in this direction. It is worth noting all elements tagged as

incoherent by the Setting's rules were not considered very incoherent by the jury evaluating the consistency of subject's resulting narrative.

5 Discussion and Conclusions

This research is about how people collaboratively write narratives and the role played by consistency in this writing. A medium term goal is to provide a useful tool to support it. In this section crucial issues emerging from the three experiments are discussed together with considering other interesting points for the near future research.

5.1 The Role of Consistency

Consistency appeared as a relevant factor during collaborative narrative composition, and it influences on the way stories are read and written in multi-authored scenarios. Let us recall that in the first experiment, authors introduced quite a few modifications to the overall story when it was necessary to maintain the consistency of the plot arc they were developing or to correct a discontinuity in the overall narrative consistency. Consistency provides stories with a sense of causality and makes them more accessible for new authors and enjoyable for readers.

We believe there is a certain cultural common knowledge of what is consistent and what represents a plot hole, defined by Ryan [17] as an inadvertent inconsistency in the logical and motivational texture of a story. In our model, a plot hole is a discontinuity in the cause-effect logic of the story discourse. Further experiments are needed to validate this hypothesis of the relevance of causal links.

However, in the second and third experiments authors were not as clearly concerned by consistency as in the first experiment. We believe this is due to the experimental settings, as the use of a more focused and constraining software prototype meant incoherencies were less visible to the users. The introduction of an underlying formal model with its own rules, and of reading mechanisms, which were absent on the first experiment, probably led to the reduced interest in providing consistency. CrossTale ensured consistency preservation in an effective way, and reduced the users' concerns.

However, consistency is not the only issue worth tracking when building stories collaboratively: the lack of visibility of scenes or the constraining effect of the model on creativity were not our focus in the experiments and should be further studied.

The distinction between the two types of consistency has been an effective way to formulate our research. The Setting provided an objective measure of consistency based on our model, and its impact in the perceived consistency level could be assessed.

5.2 Monitoring and Enhancing Consistency through the Setting

The Setting aimed at dealing with the user's concerns about consistency observed during the first experiment. These concerns seemed to mean that time and space limitations had to be enforced, and therefore, the Setting only deals with these aspects of stories. It established a framework for developing narrative collaboratively, with a clear interpretation of what is consistent and what is not. Forcing users to follow the Setting rules during the scene composition process was not a very popular design decision among authors, but the stories built under these conditions apparently provide better reading experiences. Therefore we illustrate an interesting situation; constraining scene composition under a Setting-like model may lead to more consistent results while hampering the authoring process. No

specific observations were made on creativity aspects, but we feel that the Setting could easily decrease the creativity of the stories it supports. This should be properly tested in subsequent experiments.

The Setting in the second experiment proved to be a double-bladed sword: authors were aware of some of the things they needed to take into account that might have ignored so far, but they also felt less able to express their creativity due to the constraining nature of the consistency rules. The implementation of the Setting in the third experiment is more successful; authors were always aware of violations to the Setting rules, but they could react in different ways. Some deliberately ignored the warnings, while others (the majority) prioritized such incoherencies and solved them before anything else. Ultimately, we believe there is no formal model valid and complete for all possible narratives. Our future attempts to provide support and guidance in building consistent multi-authored stories will probably involve the authors in the construction of their formal model. What might be consistent in one narrative space, such as involving magic characters, might be inconsistent in others, and there is no one better suited to establish these discriminations than the individuals who are creating the stories. Future experiments could even introduce inconsistency generators, based on approaches that generate events and situations, possibly reducing the user-perceived consistency but maybe providing some inspiration to the authors.

It is important to remark that the results coming from experiments where the modifications to the narrative persist and those where every subject deals with the same exact set of scenes are not directly comparable.

Another aspect of multi-authoring is group dynamics. The Setting essentially stated the game rules, which each author had to follow to enter into the game of story creation. On the other hand, each author introduced modifications to the narrative space that needed to be respected by subsequent authors, meaning that the learning time needed by the following author increased. A possible improvement could be to provide better communication amongst authors to support their coordination. This could improve cooperation during narrative composition and introduce specializations such as committing specific authors to preserve consistency by stating the fundamental consistency rules and reorganizing structured content.

5.3   Very Human and Causal Storylines

Human-generated stories within a narrative space, as those observed in the first experiment, are not random. Most contributions followed existing plots, commonly associated with a character or some abstract concept, such as a motivation or a specific theme. The introduction of formal storylines in the second and third experiments was meant to reinforce the sense of computational causality and continuity, trying to predict the authors' behavior to ultimately enhance the user experience. After analyzing their use during the experiments, it is safe to say that they meant a difference to the results. The reader has to follow the clear cause-to-effect relationship made explicit. The story exists in a specific region of the narrative space. Users embraced this storyline mechanism to explore and understand the narrative space, and in most cases avoided the free scene selection in favor of the sequential reading order provided. They also used this mechanism to link new scenes into existing storylines or even to start new storylines from scratch to propose new ways to read the content of their creations. This might have been one of the key reasons for the increase in the reader-perceived consistency measured in the experiments that used CrossTale.

We believe the use of storylines as tools to communicate stories is fundamental in the exchange between a storyteller and its audience. From the Setting computational point of view, storylines are not necessary for the narrative space to exist. However, without them, it

is rather information with no narrative quality. Even if storylines did not formally exist in the Setting information architecture, any story introduced by human beings would probably has cause-to-effect relationships.

Another interesting finding that we will probably introduce in future attempts to map a story to a formal model is that readers prefer storylines based on characters to those based on plots, as they chose the former almost always. Apparently, in the context of a non-linear story, users find more natural to follow specific characters instead of plots. One possible explanation is that in most of our stories (and in many stories found on contemporary media) a character only appears in one plot with a main role. While s/he could appear (seldom) in additional storylines, the character would then have a minor role. Some of the most popular Semiotic models [16] are built around characters and their roles, rarely depicting meaningful entities that display human-like behavior. We will explore this approach in the future.

5.4   Conclusions and Other Future Work

Narratives are highly subjective, as any product of an artistic discipline. There is an implicit notion of causality in any story. Our experiments are not exceptional. Scenarios involving cooperation between authors often suffer from discontinuities in their causal relationships, which produce less satisfactory stories for their readers. We believe consistency plays a fundamental role and we presented experimental data that supports our belief. Our approach introducing a formal model that imposes consistency constraints derived from the narrative space was tested; showing it was capable to monitor and increase the structural consistency of the multi-authored narrative space as intended. This apparently translated into stories with an enhanced reader-perceived consistency. However, the negative reaction from authors when facing constraints imposed by the model requires further exploration. We believe some media (such as TV, films, comics amongst others) have the difficulties of collaboration amongst multiple authors discussed throughout the paper, and we plan to extend to them the methods introduced.

There are also some possible paths for future work that deal with some secondary factors observed. Regarding creativity, subjects from all experiments seem to perceive scenes created by authors with a background in communication or arts as generally more creative but not necessarily more consistent. The relation between creativity and consistency is not clear at all in our observations. A more specific experimental design, possibly involving subjects with specific backgrounds and narrative expertise, could shed more light into the matter and maybe provide some details on the hypothetical correlation between creativity and consistency.

On the other hand little attention was paid to the interaction and aesthetic design of CrossTale. This is an interesting line of research that deals mainly with usability and user experience, potentially improving the CrossTale results.

─── **References** ───

1   Articy draft. `www.nevigo.com/en/articydraft/overview`. Accessed: May 2013.

2   David Benioff and D. B. Weiss. *Game of Thrones: The Complete First Season.* HBO Entertainment, 2012. 5 DVDs.

3   Kraig Blackwelder, Jacob Klunder, Matthew McFarland, and Will Hindmarch. *Vampire: The Requiem, Coteries.* White Wolf Pub Print, 2004.

4   Marc Cavazza, Fred Charles, and Steven J. Mead. Character-based interactive storytelling. *IEEE Intelligent Systems*, 17(4):17–24, 2002.

5 Celtx. `www.celtx.com`. Accessed: May 2013.

6 Joaquim Colàs, Alan Tapscott, Ayman Moghnieh, and Josep Blat. CrossTale: Shared narratives as a new interactive medium. *MMEDIA 2012, The Fourth International Conferences on Advances in Multimedia*, pages 106–111, 2012.

7 Chris Crawford. Storytron. interactive storytelling. `www.storytron.com`. Accessed: May 2013.

8 Fanfiction.net. `www.fanfiction.net`. Accessed: May 2013.

9 Gary Gygax. *Advanced Dungeons and Dragons: Monster Manual.* Wizards of the Coast Print, 2012.

10 David Jackson. Storyjacker. `www.storyjacker.net`. Accessed: May 2013.

11 Brenda Laurel. *Towards the Design of a Computer-Based Interactive Fantasy System.* PhD thesis, Ohio State University, 1986.

12 George R. R. Martin. *Wild Cards.* Gollancz Print, London, 2012.

13 James R. Meehan. TALE-SPIN, an interactive program that writes stories. In D. R. Reddy, editor, *Proceedings of the 5th International Joint Conference on Artificial Intelligence. Cambridge, MA, August 1977*, pages 91–98. William Kaufmann, 1977.

14 Gregory Newby. Metric multidimensional information space. In *The Fifth Text REtrieval Conference (TREC-6), The Fifth Text REtrieval Conference (TREC-5), Gaithersburg, Maryland, USA, November 20–22, 1996*, number 500-238 in NIST Special Publications, pages 521–536. National Institute of Standards and Technology, 1997.

15 David Pizzi, Fred Charles, Jean-Luc Lugrin, and Marc Cavazza. Interactive storytelling with literary feelings. In Ana Paiva, Rui Prada, and Rosalind W. Picard, editors, *Affective Computing and Intelligent Interaction, Second International Conference, ACII 2007, Lisbon, Portugal, September 12–14, 2007, Proceedings*, number 4738 in Lecture Notes in Computer Science, pages 630–641, 2007.

16 Vladimir Yakovlevich Propp. *Morphology of the Folktale.* University of Texas Press, Austin, TX, 2nd edition, 1968. transl. L. Scott.

17 Marie-Laure Ryan. Cheap plot trick, plot holes, and narrative design. *Narrative*, 17(1):56–75, 2009.

18 David Thue, Vadim Bulitko, Marcia Spetch, and Eric Wasylishen. Interactive storytelling: A player modelling approach. In Jonathan Schaeffer and Michael Mateas, editors, *Proceedings of the Third Artificial Intelligence and Interactive Digital Entertainment Conference, June 6–8, 2007, Stanford, California, USA*, pages 43–48. AAAI Press, 2007.

19 Westeros: The 'A Song of Ice and Fire' Domain. `http://www.westeros.org`. Accessed: May 2013.

Having one's cake and eating it too: Coherence of children's emergent narratives*

Mariët Theune[1], Thijs Alofs[1], Jeroen Linssen[1], and Ivo Swartjes[2]

1 **Human Media Interaction**
   University of Twente
   Enschede, The Netherlands
   m.theune@utwente.nl, t.alofs@gmail.com, j.m.linssen@utwente.nl
2 **Ranj Serious Games**
   Rotterdam, The Netherlands
   ivo@ranj.nl

─── **Abstract** ───

In the emergent narrative approach to Interactive Storytelling, narratives arise from the interactions between player- or computer-controlled characters in a simulated story world. This approach offers much freedom to the players, but this freedom may come at the cost of narrative structure. In this paper we study stories created by children using a storytelling system based on the emergent narrative approach. We investigate how coherent these stories actually are and which types of character actions contribute the most to story coherence, defined in terms of the causal connectedness of story elements. We find that although the children do produce goal-directed story lines, overall the stories are only partially coherent. This can be explained by the improvisational nature of the children's storytelling with our system, where the interactive experience of the players is more important than the production of a coherent narrative. We also observe that the communication between the children, external to the system, plays an important role in establishing coherence of the created stories.

1 Introduction

Stories are more than just temporal sequences of events. A coherent narrative must have a causal structure, with the story events being causally connected and related to some overarching goal that 'glues' them together [10]. According to Trabasso's influential model of story understanding, coherent stories consist of hierarchies of goals, actions and outcomes [19]. In this paper we use a variation of this model to investigate the coherence of children's stories that were created using an interactive storytelling system.

   The *Interactive Storyteller* [1, 2] is a system for interactive digital storytelling that allows users to control the actions of one or more characters in a simulated story world, where they interact with characters controlled by intelligent agents. Stories emerge from the actions of the player- and computer-controlled characters; no part of the storyline is scripted in

---

advance. This approach is called *emergent narrative* [4]. Emergent narrative allows players a lot of freedom in shaping the story they take part in. However, a potential drawback of this approach is that when players can do (more or less) whatever they like within the story world, without any narrative structure being imposed on them, the resulting stories may not exhibit much coherence. This trade-off between narrative structure and player freedom in interactive storytelling is called the narrative paradox [8].

In this paper we study the coherence of children's stories created with the Interactive Storyteller. In this system, narrative is fully emergent; the human users are only constrained in their actions by the limitations of the implemented story world. For example, certain actions can only be carried out on certain locations or after other actions have been carried out. Some of the available actions lend themselves well for the creation of coherent storylines, while others do not. Thus, the Interactive Storyteller enables, but does not force, the creation of coherent stories. It is hypothesised that players, when given the opportunity by the emergent narrative system, will be motivated not only by acting out their character roles, but also by seeking out narrative coherence, i.e., perform actions that contribute to the story as a whole [16].

Our two main questions are: (1) To what extent do children playing with the Interactive Storyteller spontaneously create coherent storylines? (2) Which of their selected actions contribute to coherence, in terms of logical connections between story elements, and which do not? By answering these questions, we try to shed some light on the issue to what extent emergent narrative can achieve not only player freedom but also narrative coherence. We also hope to gain some preliminary insights into which types of actions may help to optimise both player experience and narrative structure in emergent narrative systems.

We let four pairs of children play with the system and analysed the stories they produced, determining their coherence by looking at the story components selected by the children and the extent to which these components were causally connected. We included the children's system-external communication in the analysis, because our approach to interactive storytelling is closely related to children's improvised dramatic play, a narrative activity in which metacommunication plays a very important role [14].

First we briefly discuss related work (Section 2) and describe the Interactive Storyteller (Section 3). Then we describe the set-up of the experiment (Section 4) and present our story analysis (Section 5). We end with a discussion and conclusions (Sections 6 and 7).

2 Related work

Children's narratives have been the subject of many studies in the field of cognitive psychology, investigating the influence on narrative coherence of factors such as children's age. In most of these studies, data were gathered by having children narrate the events shown in picture sequences [15, 20] or (more rarely) other media such as TV shows [9]. Our work is inspired by these studies, in that we use similar methods to determine coherence. However, instead of studying narrations of given event sequences, we look at stories created by children using an interactive storytelling system.

In the past, several interactive storytelling systems aimed at children have been developed. FearNot! [5] is an educational system in which intelligent virtual characters autonomously act out bullying scenarios, with the children giving advice to the victim. TEATRIX [12] is a tool for the development of narrative competence through collaborative story creation. Children can take on the role of a character and interact with computer-controlled characters in a virtual environment to create stories. Ghostwriter [13] is a virtual role-playing environment for

children, who interact with one computer-controlled character and two characters controlled by a human role-play leader. Like the Interactive Storyteller, these systems all take an emergent narrative approach and support children's story improvisation through roleplay in interaction with virtual characters. However, except for an analysis of the GhostWriter role play logs for indications of characters' personalities and moods [13], we have found no analyses of the stories that have been created with these (or other) interactive storytelling systems.

Storytelling systems with tabletop interfaces, like the Interactive Storyteller, include Reactoon [3], TellTable [7] and StoryTable [21]. These aim at facilitating storytelling by children and do not incorporate intelligent agents as characters in the story. In most cases, there appears to have been no evaluation of the stories created with these systems, with the exception of the StoryTable system. A comparison of stories created with or without the StoryTable revealed no significant differences in structure and cohesion [21].

3   The Interactive Storyteller

The Interactive Storyteller is a system for interactive digital storytelling with a multi-user tabletop interface. It was built on the basis of the Virtual Storyteller (VST), a multi-agent system for story generation in which intelligent agents act out the role of characters in the story [16, 17, 18]. Before describing the Interactive Storyteller and the story domain used for our experiment, we first provide a brief explanation of the underlying VST system.

Story generation in the VST happens in two phases: simulation followed by presentation. In the simulation phase, autonomous agents play the role of characters in the story world. They pursue goals, reason about their perceptions, experience emotions and take actions in the world. The resulting event sequence is captured in a formal representation called the *fabula*, which is based on the causal network model of Trabasso et al. [19]. The fabula model defines causal relationships between story elements such as goals, actions, beliefs and emotions. In the presentation phase, a narrative is generated based on a fabula representation of all that transpired in the story world. This narrative is a fluent text in which discourse markers (*because, so, then, . . .*) are used to express the causal relations between the story elements. See [17] for more details on the fabula model and how it can be used to generate cohesive narrative texts.

One of the story domains that has been created for the VST is loosely based on the "Little Red Riding Hood" (LRRH) fairy tale [18]. The LRRH story world contains three characters: Little Red Riding Hood (Red), Grandmother (Granny), and the wolf (Wolf). It has five locations: Red's house, Granny's house, the clearing in the forest, the beach, and the lake. Possible character actions in the LRRH domain include moving between locations, talking to other characters, baking, eating and poisoning cakes, and location-specific actions such as diving into the lake. This is the domain we used in the Interactive Storyteller.

Whereas in the VST, stories are generated by intelligent agents without human intervention, in the Interactive Storyteller human users can control the actions of one or more characters through a graphical interface. For each character, control can be switched between an autonomous agent and a human user. The story world is represented visually on a multi-touch table, using a top-down map view, as shown in Figure 1. To prevent one side of the table from being optimal for perceiving the story world, characters are displayed in one direction and houses are projected the other way. The locations in the story domain are marked by blue circles on the map. The map also shows the paths between the locations, which need to be followed by the characters. For example, a character cannot go directly

Figure 1 The interface of the Interactive Storyteller (touch-only version). Location names have been added for the reader's benefit.

from Red's house to Granny's house, but only via the beach or the clearing in the forest.

Users can change the locations of characters by moving physical toys that represent the story characters across the surface of the multi-touch table, from one location to the next. The use of these tangibles is optional; we also developed a touch-only version of the interface where graphical representations of the characters can be moved around by dragging them across the table surface. Figure 1 shows the touch-only version of the interface. The tangible interface is the same, except that the characters are not represented by images but by physical objects, as shown in Figure 2 below.

There is a menu for the selection of non-move actions by users, shown in the centre of Figure 1. Actions are carried out in a turn-based fashion. When a character gets the turn, the system determines which actions are possible for that character given the current state of the story world, and displays them in the menu. To allow an equal view from all sides of the table, the menu is presented in two directions.

When a character performs an action, this action is expressed using a simple sentence and narrated using Loquendo text-to-speech. As we found in pilot tests with the Interactive Storyteller, the spoken narration provides valuable feedback to the children. The narrations are also displayed in the two scrolls that can be seen in Figure 1 (one for each side of the table). These scrolls serve as a time-independent source of information about what has happened in the story so far. For more details on the interface and the design motivations of the Interactive Storyteller, see [1, 2].

4   Experiment

In a small-scale user experiment, we let four pairs of 8–11 year old children play with the Interactive Storyteller. All participants were pairs of siblings or friends; see Table 1. One child controlled the character Red; the other controlled Granny. Wolf was always computer-controlled.

Table 1 The four pairs of participants in our user experiment.

	Participants (age)	Relationship	Session 1	Session 2
Pair 1:	female (9) + male (9)	friends	touch-only	tangibles
Pair 2:	female (10) + male (8)	siblings	tangibles	touch-only
Pair 3:	male (8) + male (10)	siblings	touch-only	tangibles
Pair 4:	female (10) + female (11)	friends	tangibles	touch-only

Figure 2 Children playing with the Interactive Storyteller (tangible version).

The experiment was originally set up to test, among other things, whether the use of tangibles in our system setup is preferred over a touch-only approach. Therefore, each pair of children carried out two play sessions, one with and one without tangibles (in a counterbalanced order). The results of this comparison are reported in another paper [2]; here we only focus on the stories that were created by the children during the experiment. Because system version (tangible or touch-only) appeared to have no influence on the stories that were created, we ignore this factor in the remainder of this paper.

At the start of each run of the experiment, the children were told that they could use the system to create a story, and the basics of the user interface were explained to them. Because we were interested in their spontaneous behaviour while interacting with the system, we did not give them a specific goal to achieve, nor did we ask them to talk or think aloud. The entire experiment was recorded on video; see Figure 2. On average, the play sessions lasted 12 minutes (min. 9, max. 14 minutes). All story actions taken through the Interactive Storyteller interface were logged, and the children's communication during the play sessions was transcribed.

5 Story analysis

In this section we analyse the stories generated in the experiment to see to what extent the children made use of the possibilities to create coherent stories offered by the Interactive Storyteller. Figure 3 shows excerpts of two of the stories that were created. The fragments

---

Story fragment 1 (Pair 4, session 1)
1. [At Granny's house] Granny bakes the cheese cake.
2. Wolf sneaks to Red's house. [From the clearing in the forest]
3. [At Red's house] "Hello," Red says to Wolf.
4. With a little bit of cyanide, Granny poisons the cheese cake.
5. "Hello back to you," Wolf says to Red.
6. Red does the dishes.
7. Granny shuffles to the beach.
8. "Give me that!" says Wolf and takes the birthday cake from Red.
9. "Give that back, it is mine!" says Red, and takes the birthday cake back from Wolf.
10. Granny shuffles to Red's house.
11. "Give that back!" says Wolf and takes the birthday cake back from Red.
12. "Oh Granny," says Red, "Wolf stole the birthday cake from me!"
13. "Oh, I feel sorry for you," Granny says to Red.
14. Wolf eats the birthday cake.
15. "This is for you," says Granny, and gives the cheese cake to Wolf.
16. Red bursts out laughing.

Story fragment 2 (Pair 1, session 2)
1. [At the lake] Red treads water for one minute.
2. [At the lake] Granny farts under water.
3. Red skips to the beach.
4. Granny bursts out in tears.
5. [At the lake] Wolf gets pretty hungry.
6. Red skips to Red's house.
7. Granny shuffles to the beach.
8. Red is angry at Granny.
9. Granny shuffles to Granny's house.
10. Red skips to the clearing in the forest.

Figure 3 Two fragments of stories created by the children in our experiment.

are shown as they were narrated in the story scrolls,[1] with each line describing the action of a character, controlled either by the children (playing the roles of Red and Granny) or an autonomous agent (Wolf). For convenience, we have numbered the lines and added information about the characters' initial locations between brackets. Because the actions are described one by one as they happen during the interaction, the narrations are not very fluent, unlike the fabula-based narratives that could be generated by the VST after simulation is finished. However, in our story analysis we focus on the created story content and not how it is narrated. We look at the coherence of the stories from three different angles: the type of story components selected by the children (Section 5.1), the causal connectedness of these components (Section 5.2) and the coherence from the children's perspective, based on the transcripts of their communication (Section 5.3).

5.1   Story components

First, we analyse the stories created by the children in terms of their components: the character actions the children selected during their interactive storytelling sessions, as captured in the system logs.

---

[1] The texts are translated from Dutch.

Previous studies of story coherence have looked at the presence of narrative components such as setting information, character descriptions, actions, dialogue, internal responses, obstacles and repairs [9, 15]. Of these component types, only three are available for selection by the players in the Interactive Storyteller: dialogue, internal responses (i.e., emotions), and physical actions. Setting and character information is given by the system at the start of each session: *"Once upon a time, there was a little girl wearing a little red riding hood. She wanted to bring a birthday cake to her grandmother. . . "*. Note that although this introductory text mentions the goal of the character Red, the human player controlling Red is not in any way forced to act on this goal.[2] Players can set their own goals, but these cannot be passed on to the system. As a consequence, our story logs contain no goal information.

For our current analysis we divide the action component into different subtypes. Inspired by Trabasso's distinction between causal chain and dead-end events [19], we distinguish three types of actions: causal chain actions, dead-end actions, and move actions. In Trabasso's causal network model of stories, causal chain events are causally connected to other events in a goal-action-outcome sequence, while dead-end events have no follow-up and do not lead to goal satisfaction or failure [19]. Note that where Trabasso's division between causal chain and dead-end events is based on a causal network analysis of actual stories, our action type classification is an *a priori* one, based on our estimation of the likelihood that actions of a certain type will be causally connected to other actions. It does not tell us whether the actions will actually end up as part of a causal chain or a dead ends in a story. We see move actions as a third, separate action type, because it is difficult to predict the likelihood of moves being part of a causal chain or not. Move actions can go either way, and in that respect they are somewhere in between causal chain and dead-end actions.

Summarising, we distinguish five types of story components that can be selected by the children: causal chain actions, dead-end actions, move actions, dialogue and emotions. We describe these types in more detail below, after which we report on how often the children used them in their stories.

- *Causal chain actions* such as baking and giving away cakes have the strongest potential for creating coherent stories. They establish the preconditions for other actions to be performed, and thus may be part of goal-directed action sequences: characters can plan a series of such actions to try to achieve some goal. For example, a plan to poison Wolf might involve first baking a cake, then poisoning it with cyanide, and finally giving it to Wolf (expecting him to eat it). Each action in such a sequence is a prerequisite for the next.

- *Dead-end actions* are the opposite of causal chain actions in that they have the least potential for contributing to story coherence. These are actions such as watching birds or diving into the lake, which (at least in our implementation) do not change the story world and thus cannot serve as prerequisites for any other actions: no new actions become available after a dead-end action has been carried out. This makes these actions dead ends by definition.

- *Move actions* may or may not be causally connected to other actions. They may be carried out more or less at random, as we see in Fragment 2, where Red is skipping around without any clear in-story purpose. But they can also be used to establish preconditions for other actions. For example, a character can only give a cake to another character if

---

[2] In fact, there was only one story in which Red actually gave the birthday cake to Granny, who ate it. In most other stories the cake was either eaten by Wolf or by Red herself.

■ **Table 2** Story components in the LRRH domain. The last column shows how often each type was used by the children in the experiment. The actions of Wolf, the computer-controlled character, are not included.

Component type	LRRH actions	#	Freq.
Causal chain action	Bake(Cake), Poison(Cake), GiveTo(Cake,Character), TakeFrom(Cake,Character), Eat(Cake)	71	20%
Dead-end action	WashDishes, DustCabinet, DiveIntoLake, TreadWater, FartUnderWater, EnjoyTheSun, RollInSand, TakeNap, BuildSandCastle, WatchBirds, WatchClouds, LeaveAlone(Character)	71	20%
Move action	MoveFromTo(CurrentLocation,NewLocation)	111	32%
Emotion	Laugh, Cry, BeAngryAt(Character)	67	20%
Dialogue	Greet(Character), TellAboutCakeTaken(Character), ExpressSympathyFor(Character), Thank(Character) AskWhatToDo(Character), TellAboutPlan(Character)	27	8%
Total		347	100%

both are at the same location, and to achieve this, a sequence of move actions may be necessary. We see this in Fragment 1, where Granny goes to Red's house via the beach as part of her plan to poison Wolf, who is at Red's house.

■ *Emotions* in the LRRH domain are limited to anger, happiness (expressed through laughter) and sadness (expressed through crying). They can be causally connected to other actions in two ways. First, emotions can cause goals to be adopted; for example, being angry at Wolf may cause Red or Granny to adopt the goal of poisoning him. Second, emotions may be internal responses to story events; for example, Red laughing when Wolf is given the poisoned cake in Fragment 1.

■ *Dialogue* includes actions such as characters greeting each other or telling each other about events in the story world. One dialogue action often triggers another, leading to brief exchanges like the one in Fragment 1, where Red complains to Granny about Wolf having taken her cake, and Granny reacts by expressing her sympathy. The boundaries between dialogue and emotion are not very clear-cut: expressing sympathy through dialogue could also be seen as an indication of emotion (feeling sorry for the other), and the same could be said of saying thanks as an indication of gratefulness. However, we feel the dialogue aspect is more prominent in these cases.[3]

Based on the system logs from the experiment, we determined the frequency with which the different story components were included by the children in the stories they created. Table 2 shows the results.

In total, the children carried out 347 character actions over all play sessions; an average of 43 actions per play session. Not included in this count are the actions by Wolf, the computer-controlled character. Wolf carried out 69 actions over all sessions, which amounts to only 17% of all character actions. The reason for this relatively low percentage is that Wolf only carried out goal-directed actions (in his case, chasing after cakes to eat) and frequently skipped a turn when unable to make a plan to achieve his goal of eating cake.[4]

---

[3] As can be seen in Fragment 1, some character dialogue is included in the narration of GiveTo and TakeFrom actions. However, these actions are not presented as dialogue actions to the players; the dialogue is only added to spice up the narration.

[4] In the LRRH domain only Granny and Red can bake cakes, so Wolf could not bake his own.

Looking at the actions selected by the children, we see that causal chain actions are just as frequent as dead-ends. This is remarkable, because in the LRRH world causal chain events can only be selected when certain preconditions are met. For example, poisoning and giving away a cake is only possible if it has been baked first, and baking a cake can only be done at Granny's house. In contrast, the other four locations all had at least two dead-end actions available with no other restrictions than the character being at one of those locations. The fact that the children did not carry out more dead-end actions, suggests they may have had a preference for actions that were more likely to be part of a causal chain. However, as noted above, the causal chain potential of the selected actions need not have been fulfilled in the stories that were produced: without follow-up, a causal chain action would be a dead end in practice. In other words, only looking at the presence of story components is not a sufficient indication of coherence; we also need to look at their causal connectedness in the story [9].

5.2 Causal connectedness

To see whether the selected story components were really part of causal chains or not, we examined the causal relations between them. To this end we looked at the stories produced by the children (as shown in Figure 3), and used common sense reasoning to find logical connections between the story components.[5] To count as causally connected, story components either had to be part of a goal-directed causal chain in the sense of Trabasso et al. [19], thus providing global coherence, or be connected to an immediate cause or consequence, thus providing local coherence [20]. Simple enablement relations (e.g., a move to some location enabling a character to do something at that location) were only counted if they were clearly part of a plan. Unlike [9], in our analysis we did not limit ourselves to determining causal relationships between adjacent story components. In the Interactive Storyteller, characters take turns to carry out actions in the story world. This means that in the narration, components of an individual character's plan are usually not adjacent but intersected by other characters' actions.

As an illustration, consider the two story fragments in Figure 3. Fragment 1 holds a clear example of a goal-action-outcome sequence with Granny carrying out a plan to poison Wolf, baking a cheese cake as the first step (1), poisoning it (4), and going to Wolf (7, 10) to give it to him (15).[6] Red, meanwhile, carries out more locally connected actions: she reacts to Wolf's arrival at her house by greeting him (3), and to his repeatedly taking her birthday cake, the first time by taking it back (9) and the second time by complaining to Granny (12). In her turn, Granny briefly interrupts her poisoning mission to commiserate with Red (13). Finally, Red reacts with laughter (15) to Granny giving the poisoned cake to Wolf. Overall, Fragment 1 is quite coherent. The only action that is not causally connected is Red doing the dishes (6), a dead-end action with no apparent cause or purpose within the story. Fragment 2, on the other hand, is incoherent: the characters seem to be carrying out actions at random, and when they display emotions it is not clear what these are in response to.

All stories produced in the experiment were annotated by two of the authors, who determined for each character action, dialogue or emotion whether it was coherent (i.e., causally connected to another story component) or not. Annotation took place in three

---

[5] The VST, the non-interactive version of our system, logs all causal connections between character actions in the fabula model. However, in the Interactive Storyteller these causal connections are only recorded for computer-controlled characters. For player-controlled characters, the system only records which actions they took, because it has no knowledge of their plans and motivations.

[6] The eventual outcome, Wolf eating the poisoned cake, is not included in the excerpt.

rounds. In the first round, half of the stories were annotated by both annotators. The interannotator agreement for this first round was fairly low (Cohen's κ 0.60). After discussing and finetuning our definition of connectedness, the remaining stories were annotated, this time achieving a satisfactory level of agreement (Cohen's κ 0.73). Finally, in the third round all differences between the annotators were resolved by discussion. The results are given in Tables 3 and 4.

Table 3 shows per story component type how many instances of that type were causally connected to some other component, or not. It also shows how many story components were causally connected overall, regardless of their type. Taking together all components of all stories, 40% of these components were causally connected (i.e., coherent) while 60% were not. Table 4 shows the proportion of all connected or disconnected story components that belonged to each of the component types, thus showing the relative contribution of all types to the overall coherence of the created stories.[7]

Table 3 Causal connectedness of story components per type.

Component type	Causally connected	Not causally connected
Causal chain action	75%	25%
Dead-end action	3%	97%
Move action	32%	68%
Emotion	34%	66%
Dialogue	93%	7%
All types	40%	60%

Table 4 Contribution to story coherence of each component type.

	Causal chain	Dead-end	Move	Emotion	Dialogue
Coherence	38%	1%	25%	17%	18%
Incoherence	9%	33%	36%	21%	1%

Unsurprisingly, most causal chain actions turned out to be causally connected, and causal chain actions contributed the most to story coherence, with the opposite holding for dead-end actions. We found two dead-end actions that actually seemed to have a causal connection to another story component: one was Granny rolling in the sand at the beach, triggering an emotional response by Red (laughing), and the other was Red leaving Wolf alone after having given him a cake to eat. Only 36% of the move actions turned out to be causally connected, either because they were part of a plan or because they triggered an immediate reaction such as a greeting. Due to their frequency, moves still provided the second largest contribution to story coherence. Across all stories, we counted 18 goal plans, i.e., goal-directed action sequences. These causal chains had an average length of four story components.

Of the emotions, only 34% were annotated as being causally connected, usually because they were clearly a response to another story component. In contrast, dialogue connectedness was 93%. Dialogue thus provided a relatively large contribution to story coherence in spite of being the least frequently selected component type. The most frequent dialogue actions

---

[7] Due to rounding, the percentages for coherence in Table 4 do not add up to 100%.

were exchanges of greetings, which (similar to emotions) provided local rather than global coherence.

A mixed picture emerges from this analysis. On the one hand, a large portion of the created stories was found to be coherent (in terms of the number of causally connected components), displaying both global and local coherence. On the other hand, an even larger portion of the stories was incoherent, consisting of seemingly unconnected story components. However, this analysis presents a somewhat misleading image of the stories' coherence, because not all that makes the stories coherent has been captured in the logs of the Interactive Storyteller. Specifically, what is missing are the verbal and nonverbal contributions the children made to the story while they were interacting with the system and with each other. We address these in the next section.

5.3   Children's communication

The emergent narrative approach to interactive stoytelling is closely related to improvisation and children's dramatic play. As argued by Sawyer [14], transforming an improvisation into a fixed text removes its most essential aspect, which is that it is a social, collaborative process, in the form of a dialogue between the players. Therefore, it is important not to overlook the children's communication when analysing the stories they created.

In the Interactive Storyteller, what happens in the story world is automatically narrated through synthetic speech and shown in the graphical interface. Thus, in principle, stories can be created without any spoken communication between the human players. In practice, we found that most of the children did spontaneously communicate with each other during the experiment, thereby adding an extra layer to the emergent story. We inspected the transcripts of children's communication to get some insight in their goals and motivations for selecting certain story components. Because many of the selected actions went undiscussed by the children (with one pair of children being too shy to say much at all), the following observations are anecdotal in nature.

First, we examined the communication transcripts to see if they provided any evidence for or against the in-story goals and motivations we had inferred as part of our story analysis in Section 5.2. For certain character actions, this inspection revealed external (out-of-story) motivations or unexpected in-story motivations were revealed. For example, one child announced *"I am going to bake, that is fun!"* indicating that this causal chain action was not initially selected as part of a plan, while another child explained *"I am giving the* [poisoned] *chocolate cake to Wolf. He didn't say anything, so I had to poison him."* Other examples are the following exchanges:

(1)   **Pair 2, session 2**
      Action (by Child 1): Red eats the birthday cake.
      Child 2: *Actually, Red should have given the birthday cake to Granny.*
      Child 1: *I already ate it. Otherwise the wolf would have eaten it.*

(2)   **Pair 1, session 2**
      Action (by Child 2): Granny bakes the cheese cake.
      Child 2 (in character): *"Especially for you, Red."*

Not all goals mentioned by the children were actually followed up by actions. An example is a child (controlling Red) exclaiming *"I want to kill Granny"* after having been angered by the child playing Granny. However, after an alternative proposal by the other child, it was jointly decided to poison Wolf instead.

The presence of goal plans was often revealed through the children discussing what to do next. Examples are the following.

(3)   **Pair 3, session 2**
      Child 1: *Bake another cake.*
      Child 2: *That can only be done in Granny's house.*
      Action (by Child 2): Granny shuffles to Granny's house.

(4)   **Pair 4, session 2**
      Action (by Wolf): "Hello!" Wolf says to Red.
      Child 1: *You have to give him the cake!*
      Child 2: *It has to be poisoned first.*
      Child 1: *(. . .) I will do that, alright?*
      Child 2: *OK.*

(5)   **Pair 1, session 1**
      Action (by Child 2): Granny poisons the chocolate cake.
      Child 1: *I am going to tell Wolf that you took my chocolate cake!*
      Child 2: *No, I am poisoning it.*
      Child 1: *Yes, but then I'll give it to Wolf.*

When the children explicitly mentioned making plans for the next actions of their characters, this usually happened late in the first interaction session, or in the second, after the children had gained some insight in the available actions in the story world and the causal connections between them. Overall, it seems that many causal chain actions were initially selected just like all other actions, as part of an exploration of the story world, to see what would happen, or simply because they were thought to be 'fun'. However, most children quickly discovered that these actions could be used to achieve goals such as giving away cakes and poisoning other characters (usually Wolf, but sometimes also their playmate's character), and then they invariably started making plans to do so. Those plans were sometimes abandoned prematurely, or the causal chains were disrupted by other actions, for example if the children were distracted from their plan by location-specific dead-end actions. The children did not appear to perceive this as incoherent. In line with Sawyer's observations of children's dramatic role play [14], they were caught up in the moment-to-moment contingency of the improvisation process and did not seem overly concerned with the global coherence of the emerging story. This is similar to adult role players, who focus on a local rather than a global story level [8]. We noted one possible exception, where the children did not want to end their play session until Wolf had eaten the cake they had poisoned for him. However, it seems likely that these children simply wanted to see the achievement of their goal, and did not consciously aim for narrative closure.

The children's dialogue also showed that several actions that seemed incoherent within the story narration were not incoherent at all in the context of the improvisation. For example, it turns out that Granny's crying in Fragment 2 of Figure 3 did in fact have a causal connection with the preceding action (Red skipping to the beach), as the child controlling Granny provided the in-character motivation *"Because you are leaving me!"*. Another seemingly incoherent emotion in the same fragment is Red getting angry at Granny for no clear reason. Inspection of the children's communication transcripts revealed that Red's emotion was actually a reaction to an event outside the story:

(6) **Pair 1, session 2**

Child 2 (holding tangibles up to camera): *Isn't this a nice Granny?*
Isn't this an ugly Red Riding Hood?
Action (by Child 1): Red is angry at Granny.

Another example of a child providing a motivation for a seemingly incoherent emotion is the following:

(7) **Pair 3, session 2**

Action (by Child 2): Granny bakes the cheese cake.
Action (by Child 1): Red bursts out in tears.
Child 1: *Why can't I bake the cake?*

In this fragment it is not entirely clear whether the child was speaking in character or out of character, but it was most probably the latter. This means that like in example (6), Child 1 used the character to express his own emotion. These examples show that the children did not always distinguish between what was part of the "story proper" and what was not; they were busy playing, not story-making.

Through their play communication, the children did not only motivate and negotiate the character actions, but also enriched the emerging story in various ways. In a few cases, instead of selecting dialogue actions for their characters through the system, the children carried out the character dialogue in person, for example thanking the other or expressing sympathy for the loss of a cake. The fact that dialogue could be more easily expressed via direct communication than via the system may explain the relatively low frequency of dialogue among the story components.

In most cases, the children's play-acting did not replace the system's narrations, but augmented them. The children frequently played out what happened in the story world, expressing their character's emotions through sounds and facial expressions, and miming character actions such as shuffling, diving and eating, often while repeating the system's narrations. In other cases, the children brought in new elements that were outside the possibilities of the system. They added their own narrations (*"Much later… Wolf is still trying to think of a plan"*), expressed character emotions that were not available in the system (*"Ooh, I'm scared!"*), invented imaginary props (a cake-mold to bake a cake), and even made up entirely new actions (Red kicking Granny out of the water at the lake).

The observations we have presented in this section show that the actions that were logged and narrated by the Interactive Storyteller do not tell the whole story. The children's communication is an integral part of the improvised narratives, supplying causal connections that may not be obvious at first sight, and enriching the story in various ways.

6 Discussion

When looking at the components of the stories created in our experiment, we see that the children used more actions that had a strong potential for being causally connected ('causal chain actions'), and fewer actions with little such potential ('dead-end actions') than might have been expected based on the availability of these action types. Inspection of the causal links between these components, as inferred from the narrations, showed that 40% of all story components were causally connected. Moreover, when seen in the context of the children's play most stories turned out to be more coherent than they seemed to be when judging only from what was logged by the system. The children's system-external communication confirmed that they were acting in a goal-directed fashion at least part of the time.

At the same time, it is clear that the children did not use all available opportunities for achieving story coherence, in the sense of maximising the causal connections between the story events. Unlike the computer-controlled character Wolf, they did not always plan ahead, and if they did, they did not always follow through with their plans, allowing causal chains to be disrupted by unrelated actions. One possible reason for this is that at first, the children were not yet aware of all available actions and their effects. They only gradually discovered the things they could do within the story world, and this prevented them from exploiting all opportunities for goal-directed action right away. Our data support this explanation: 4 goal plans were carried out in the first play sessions, against 14 in the second. The proportion of causally connected story components was 36% for the first sessions, and 43% for the second.

Another possible reason for the suboptimal coherence was that, although the children had been told they could use the system to create stories, they had not been given the explicit goal to do so. Especially in the first sessions, much of the time the children were simply exploring the story world and trying out different actions without any concern for creating a coherent narrative. In general, the children's story making was done in a moment-to-moment fashion, with the children reacting to each other's and Wolf's actions as they happened and thus achieving local rather than global coherence. According to Sawyer [14], this is typical for children's dramatic role play: children who are engaged in such play do not attempt to create an overall coherent narrative, but instead react to each other turn by turn, which leads to 'pockets of local coherence' rather than global coherence in Trabasso's sense.

Overall, the children in our experiment were not overtly concerned with producing coherent narratives. Nevertheless, they spontaneously carried out many character actions that led to narrative coherence. Their motivation for creating globally coherent action sequences may have been simply that they thought (for example) poisoning another character would be a fun thing to do. In these cases, the players' goal of having a satisfying experience happened to coincide with the requirements of a coherent narrative.

Similar to role play, in emergent narrative the process (the players' experience) is more important than the outcome (the resulting story as seen from an observer's perspective) [8]. To support this process, all kinds of actions are needed. Causal chain actions allow the players to achieve goals; emotions and dialogue allow them to express themselves and communicate with computer-controlled characters; and dead-end actions can simply provide fun things to do in the story world. From our small-scale experiment, we cannot draw any strong conclusions as to which actions are most important to achieve both player satisfaction and story coherence. However, we suspect that providing more causal chain actions, as suggested by [18], may be most beneficial for this purpose, as it will allow the users to have a major influence on the course of the story (global agency, [11]), and thus presumably increase their enjoyment. As a 'side-effect', the generated narratives will also be more coherent.

To test this hypothesis, larger and more focused user experiments will be necessary. In such experiments, the players should first be given the opportunity to thoroughly explore the story world and acquaint themselves with the system, to make sure that any lack of coherence in the created stories cannot be attributed to the players' lack of awareness of the possibilities. In addition, it might be useful if the players could not only select character actions but also character goals through the system's interface. Finally, the players could be asked to think aloud to provide more insight in their reasoning and motivations than could be gained from their spontaneous communication. A drawback of such an instruction is that it is likely to interfere with the children's improvisation. The alternative 'talk aloud' approach (where children are instructed to talk about what they are doing instead of what they are thinking, [6]) is supposed to be easier for children, but may still suffer the same problem.

On the other hand, unintrusive methods such as retrospection and post-task interviews are expected to be less informative.

Future experiments should be carried out not only with children, but also with adult players, using an appropriate story domain. Previous studies have shown that at age 8, children "seem to be on the verge of a more sophisticated understanding of goals" [10, p. 327] and at age 9, they perform at a similar level as adults with regard to narrating stories in terms of goal plans [20]. Given their age group, the children who participated in our experiment should therefore have been capable of generating equally coherent stories as adults. However, children tend to be more playful and have a shorter attention span than adults, and this may have influenced the results. Although we cannot generalise our current findings to adults, we do expect that adult players will have similar preferences as children in regard to story component types and goal-directed planning.

7   Conclusion

The emergent narrative approach to interactive storytelling offers the players much freedom, but this comes at the potential expense of narrative structure. To investigate the extent of this trade-off (also called the narrative paradox), we analysed the stories created by children using an emergent narrative system. We determined the coherence of these stories, defined in terms of the causal connectedness of their components, and examined which types of character actions contributed the most to story coherence.

We found that the stories created by the children exhibited both global coherence, in the form of goal-action-outcome sequences [19] and local coherence in the form of immediate responses to other character's actions [14]. We also observed that the communication between the children, external to the system, played an important role in establishing coherence of the created stories. However, the children's stories were only partially coherent. This can be explained by the improvisational nature of emergent narrative, where the interactive experience of the players is more important than the production of a coherent narrative. Given that the children in our experiment spontaneously selected relatively many actions that had a strong potential for being causally connected, and used them to create goal plans, we surmise that focusing on such actions in the creation of story worlds for emergent narrative may improve both the players' experience and the coherence of the created stories (cf. [18]).

Returning to the narrative paradox, what it basically says is that in interactive storytelling, one cannot have one's cake (maximise the player's influence on the story) and eat it too (have a good narrative structure). However, the results of our study suggest that when taking an emergent narrative approach, it is possible to have the cake and eat a large piece of it too.

───── **References** ───

1   T. Alofs, M. Theune, and I. Swartjes. A tabletop board game interface for multi-user inter-
     action with a storytelling system. In A. Camurri and C. Costa, editors, *Intelligent Tech-*
     nologies for Interactive Entertainment—4th International ICST Conference, INTETAIN
     2011, Genova, Italy, May 25–27, 2011, Revised Selected Papers, number 78 in Lecture
     Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications
     Engineering, pages 123–128. Springer, 2011.

2 T. Alofs, M. Theune, and I. Swartjes. A tabletop interactive storytelling system: Designing for social interaction. *International Journal of Arts and Technology*, To appear.

3 A. Alves, R. Lopes, P. Matos, L. Velho, and D. Silva. Reactoon: Storytelling in a tangible environment. In G. Biswas, D. Carr, Y.S. Chee, and W.-H. Hwang, editors, *2010 Third IEEE International Conference on Digital Game and Intelligent Toy Enhanced Learning, DIGITEL 2010, Kaohsiung, Taiwan, April 12–16, 2010*, pages 161–165. IEEE, 2010.

4 R. Aylett. Narrative in virtual environments – towards emergent narrative. In M. Mateas and P. Sengers, editors, *Narrative Intelligence: Papers from the AAAI Fall Symposium*, number FS-99-01 in AAAI Technical Reports, pages 83–86, 1999.

5 R.S. Aylett, S. Louchart, J. Ferreira Dias, A. Paiva, and M. Vala. FearNot! – an experiment in emergent narrative. In T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, and T. Rist, editors, *Intelligent Virtual Agents, 5th International Working Conference, IVA 2005, Kos, Greece, September 12–14, 2005, Proceedings*, number 3661 in Lecture Notes in Computer Science, pages 305–316. Springer, 2005.

6 A. Donker and P. Reitsma. Usability testing with young children. In A. Druin, editor, *Proceedings of the 2004 Conference on Interaction Design and Children: Building a Community*, pages 43–48, New York, NY, USA, 2004. Association for Computing Machinery.

7 J. Helmes, X. Cao, S. Lindley, and A. Sellen. Developing the story: Designing an interactive storytelling application. In G. Morrison, S. Subramanian, M. Sheelagh, T. Carpendale, M. Haller, and S.D. Scott, editors, *ACM International Conference on Interactive Tabletops and Surfaces, ITS 2009, Banff / Calgary, Alberta, Canada, November 23–25, 2009*, pages 49–52, 2009.

8 S. Louchart and R. Aylett. Solving the narrative paradox in VEs – lessons from RPGs. In Th. Rist, R.S. Aylett, D. Ballin, and J. Rickel, editors, *Intelligent Virtual Agents, 4th International Workshop, IVA 2003, Kloster Irsee, Germany, September 15-17, 2003, Proceedings*, number 2792 in Lecture Notes in Computer Science, pages 244–248. Springer, 2003.

9 J. Low and K. Durkin. Structure and causal connections in children's on-line television narratives: What develops? *Cognitive Development*, 13:201–225, 1998.

10 J. S. Lynch and P. van den Broek. Understanding the glue of narrative structure: Children's on- and off-line inferences about characters' goals. *Cognitive Development*, 22:323–340, 2007.

11 M. Mateas and A. Stern. Structuring content in the Façade interactive drama architecture. In R.M. Young and J.E. Laird, editors, *Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference, June 1–5, 2005, Marina del Rey, California, USA*, pages 93–98. AAAI Press, Menlo Park, CA, 2005.

12 R. Prada, A. Paiva, I. Machado, and C. Gouveia. "You cannot use my broom! I'm the witch, you're the prince": Collaboration in a virtual dramatic game. In S. A. Cerri, G. Gouardères, and F. Paraguaçu, editors, *Intelligent Tutoring Systems, 6th International Conference, ITS 2002, Biarritz, France and San Sebastian, Spain, June 2–7, 2002, Proceedings*, number 2363 in Lecture Notes in Computer Science, pages 913–922. Springer, 2002.

13 J. Robertson and J. Oberlander. Ghostwriter: Educational drama and presence in a virtual environment. *Journal of Computer Mediated Communication*, 8(1), 2002.

14 R. K. Sawyer. Improvisation and narrative. *Narrative Inquiry*, 12(2):319–349, 2002.

15 L. R. Shapiro and J. A. Hudson. Tell me a make-believe story: Coherence and cohesion in young children's picture-elicited narratives. *Developmental Psychology*, 27(6):960–974, 1991.

16 I. Swartjes. *Whose Story Is It Anyway? How Improve Informs Agency and Authorship of Emergent Narrative.* PhD thesis, University of Twente, 2010.

17 I. Swartjes and M. Theune. A fabula model for emergent narrative. In S. Göbel, R. Malke-witz, and I.A. Iurgel, editors, *Technologies for Interactive Digital Storytelling and Entertainment, Third International Conference, TIDSE 2006, Darmstadt, Germany, December 4–6, 2006*, number 4326 in Lecture Notes in Computer Science, pages 49–60. Springer, 2006.

18 I. Swartjes and M. Theune. Iterative authoring using story generation feedback: Debugging or co-creation? In I.A. Iurgel, N. Zagalo, and P. Petta, editors, *Interactive Storytelling, Second Joint International Conference on Interactive Digital Storytelling, ICIDS 2009, Guimarães, Portugal, December 9–11, 2009. Proceedings*, number 5915 in Lecture Notes in Computer Science, pages 62–73. Springer, 2009.

19 T. Trabasso, T. Secco, and P. van den Broek. Causal cohesion and story coherence. In H. Mandl, N.L. Stein, and T. Trabasso, editors, *Learning and Comprehension of Text*, pages 223–25, Hillsdale, NJ, 1982. Erlbaum.

20 T. Trabasso, N.L. Stein, P.C. Rodkin, M.P. Munger, and C.R. Baughn. Knowledge of goals and plans in the on-line narration of events. *Cognitive Development*, 7:133–170, 1992.

21 M. Zancanaro, F. Pianesi, O. Stock, P. Venuti, A. Cappelletti, G. Iandolo, M. Prete, and F. Rossi. Children in the museum: an environment for collaborative storytelling. In O. Stock and M. Zancanaro, editors, *PEACH – Intelligent Interfaces for Museum Visits*, pages 165–184, Berlin Heidelberg, 2007. Springer.

Emotional Expression in Oral History Narratives: Comparing Results of Automated Verbal and Nonverbal Analyses*

Khiet P. Truong[1], Gerben J. Westerhof[2], Sanne M.A. Lamers[2], Franciska de Jong[1,3], and Anneke Sools[2]

1   Human Media Interaction
       University of Twente
       Enschede, The Netherlands
       {k.p.truong,f.m.g.dejong}@utwente.nl
2   Department of Psychology, Health and Technology
       University of Twente
       Enschede, The Netherlands
       {g.j.westerhof,s.m.a.lamers,a.m.sools}@utwente.nl
3   Erasmus Studio
       Erasmus University Rotterdam
       Rotterdam, The Netherlands

── Abstract ──

Audiovisual collections of narratives about war-traumas are rich in descriptions of personal and emotional experiences which can be expressed through verbal and nonverbal means. We complement a commonly used verbal analysis with a nonverbal one to study emotional developments in narratives. Using automatic text, vocal, and facial expression analysis we found that verbal emotional expressions do not correspond much to nonverbal ones. This observation may have important implications for the way narratives traditionally are being studied. We aim to understand how different modes of narrative expression relate to each other, and to enrich digital audiovisual interview collections with emotion-oriented tags.

1   Introduction

In narrative psychology, life stories are generally seen as cognitive devices that people use to attribute meaning to their lives by telling and structuring their life stories [11, 16]. Numerous studies have shown that people who have suffered traumas can improve their physical and mental health by talking or writing about their experiences [6]. The linguistic expression of emotions is believed to serve as the cognitive encoding of analogous emotional experiences that makes it possible to further attribute meaning to them [13]. Resources that psychologists usually use in their analyses are texts about personal experiences, either written down by

the subject or transcribed from the subject's speech by a human transcriber. The style of writing and the words used in these written narratives reflect how people express their emotional experiences, cognitively processing and structuring them. For example, people use more positive emotion words to write about positive events and more negative content words to write about negative events [8]. In a similar vein, the use of pronouns can provide information about how people are cognitively processing the situation: more first-person singular pronouns are used by people who are experiencing physical or emotional pain since they tend to have their attention drawn to themselves [14]. Based on these findings, narrative psychologists consider written words and language (i.e., verbal behaviors) the medium for assessing and understanding the cognitive processes related to emotional expression and meaning construction.

An intriguing question is how verbal expressions in written narratives relate to non-verbal behaviors in the production of narratives. It is known that elements of non-verbal behavior such as the way people speak and their accompanying facial and bodily expressions reflect cognitive and emotional aspects of the narrator's state of mind, e.g., [4, 9]. The frequency of disfluencies (e.g., filled pauses, hesitations) for example is often linked to planning effort and cognitive load during speaking [7]. Consequently, it seems important to take both modalities, verbal and non-verbal expressions, into account when analyzing emotional experiences in narratives. With the increased availability of multimedia content, in particular oral history interview collections, it becomes possible to study both the verbal and non-verbal content of narratives. In our study, we provide a first investigation into how analyses of verbal and non-verbal emotional content in narratives relate to each other.

We apply our analyses to a large collection of audiovisual testimonies on war-related experiences in Croatia's past (CroMe [3]). For the verbal analysis, we use a computerized text analysis tool, the Linguistic Inquiry and Word Count (LIWC) [12] that is commonly employed by psychologists. For the non-verbal analysis, we automatically extract prosodic features from the narrator's speech. In addition, tools for automatic emotion detection in vocal and facial expressions are applied. These verbal and non-verbal measurements are compared to each other through correlation analysis. Our long-term goal is to develop computational models of narrative structure (based on verbal and non-verbal information) that are informed and validated by psychologists to make meaningful browsing and searching through large collections of narratives possible. More fundamental research goals are to understand how different modes of narrative expression relate to each other and how convergence and divergence between the modes is informative for understanding the possibly different cognitive processes and structures involved in verbal and nonverbal expression.

2   Material

We use audiovisual recordings of interviews made in the Croatian Memories project [3]. In this project, citizens of Croatia are interviewed about their personal experiences with war and trauma. The project is covering three timeframes: WWII, the Yugoslav period, and the war of the nineties. Guided by the interviewer's questions and responses, the interviewee (while being continuously on camera) tells personal stories about his/her experiences. Hence, these interviews have narrative characteristics and contain emotional episodes. Each interview is approximately between 40 and 70 minutes long. Currently, 50 interviews have been transcribed and translated from Croatian into English. Metadata of the interviewees (e.g., profession, religion) and summaries of the interviews are available, including English (time-aligned) subtitles. A collection of around 400 interviews will be released in the Fall of 2013, and a similar but smaller dataset is being created in Bosnia and Herzegovina.

3   Method

Verbal and non-verbal features were extracted from the transcriptions, audio and video. A manual labelling of who is speaking, the interviewer or interviewee, was performed.

3.1   Verbal analysis

The verbal analysis was carried out using the LIWC program [12]. This program automatically counts the percentage of words that have linguistic and/or psychological relevance. For this study, we calculated the percentage of positive emotions words and negative emotion words.

3.2   Non-verbal analysis

The non-verbal analysis focused on the vocal and facial modality. For the vocal analysis, we automatically measured F0, intensity and pausing behavior with Praat [2] from the interviewee's speech. These measures are known to be related to emotional and cognitive processes in speaking, e.g., [1, 15]. First, speech activity detection was performed (by thresholding intensity) in order to obtain so-called talkspurts. Talkspurts are defined as stretches of continuous speech with minimal durations of 300 ms, bounded by silences with minimal durations of 200 ms. F0 and intensity were extracted each 10 ms, and averaged over each talkspurt. Pauses are defined as within-speaker silences. In addition to these prosodic measurements, we also obtained emotion-oriented labels(and their corresponding confidence scores) from automatic emotion classification tools such as openSMILE [5] for vocal expressions and CERT [10] for facial expressions. Table 1 gives a summary of all features extracted.

Table 1 Non-verbal measures automatically extracted from speech and face.

Praat [2]	mean and standard deviation pitch, mean and standard deviation intensity, mean duration talkspurt, max duration talkspurt, mean duration pause, max duration pause, $\frac{\text{mean duration pause}}{\text{mean duration talkspurt}}$, $\frac{\text{number of pauses}}{\text{mean duration talkspurt}}$
openSMILE [5]	positive emotion, negative emotion, arousal, valence, anger, boredom, disgust, fear, happiness, neutral, sadness, aggressive, cheerful, intoxicated, nervous, tired, interest
CERT [10]	positive emotion, negative emotion, smile, agner, contempt, disgust, fear, joy, sad, surprise, neutral

4   Experimental setup

Based on the LIWC analysis, we selected 6 interviews varying in their overall level of emotional expressivity as well as their ratio between positive versus negative emotion words. Each interview was divided into 1-minute segments and analyzed as described in the methods section above. All the verbal and non-verbal measures were averaged over these 1-minute segments. We used Pearson correlations to compute the association of the percentage of positive and negative emotion words with all nonverbal markers over whole interviews.

5 Results

Fig. 1 provides the minute-by-minute graphs of some of the verbal and nonverbal markers of one interview. Such figures provide the basis description per interview. In the top panel figure, the fluctuation of positive and negative emotion words is visualized. It can be seen how the interviewee uses more or less emotion words throughout the interview and how the interviewee switches from a predominant use of positive emotion words during the first fifteen minutes to a predominance of negative emotion words during the second thirty minutes to a more mixed level in the remainder of the interview. This increased use of negative emotion words seems to be associated with a lower talkspurt duration which is reflected in a significant moderate correlation of $-0.409, p < 0.005$ (Bonferroni correction was applied). In a similar way, we also computed correlations between the expression of positive and negative emotions words with all other non-verbal indicators. However, in general, most of the correlations were low and not signficant.


Figure 1 Examples of verbal and non-verbal measurements of 1 interview. Top pane: LIWC positive and negative emotion scores. Bottom pane: average duration talkspurt (s).

6 Discussion

We conclude that the verbal and nonverbal characteristics were unrelated in our study. This might be due to limitations in our analyses. This applies for the verbal analysis, the nonverbal analysis as well as the correlational analysis. In general, we used a one minute interval, whereas longer intervals that are characterized by meaningful story sequences about particular experiences might be more sensitive. Further research is envisaged to check this assumption.

For the verbal analysis, we used the translated transcripts of the originally Croatian narratives. Due to the translation and its purpose for subtitling the interviews, important information on the verbal expression of emotions and its timing might have been lost.

With regard to the nonverbal analysis, we used automatic emotion classification tools that were trained on specific material that not necessarily matches the realistic and noisy audiovisual material we used for our analysis. For example, the facial expression software CERT appeared to work suboptimally as the software is trained on posed facial expressions and requires a frontal face view. In general, we need to analyze more interviews to investigate how idiosyncratic certain speech and facial behaviors are in these specific narratives. In addition, we plan to compare the use of other affect recognition and text mining tools, besides the LIWC.

With regard to the relation between verbal and nonverbal characteristics, we only performed Pearson correlation analysis that does not take the complex nature of the data

into account. In the future, we will use mixed models that can take the dependency and timing within the data of an individual as well as inter-individual differences into account. For these analyses, more interviews need to be analyzed in order to obtain reliable results.

Despite these limitations, the overall lack of relations between the verbal and nonverbal characteristics may have important consequences for theory development as well. It might point to differences in the cognitive processes and structures involved in the verbal and nonverbal expression of emotions. Against the overall lack of relations, the convergence between verbal and nonverbal emotional expression might have a distinct meaning.

References

1 R. Banse and K. R. Scherer. Acoustic profiles in vocal emotion expression. *Journal of Personality and Social Psychology*, 70:614–636, 1996.

2 P. Boersma. Praat: a system for doing phonetics by computer. *Glot International*, 5(9/10):341–345, 2001.

3 Project Croatian memories (CroMe), 2013. `http://balkanvoices.org/?page_id=23`, accessed 24 February 2013.

4 C. Darwin. *The expression of the emotions in man and animals*. D. Appleton and Co., New York, 1872.

5 F. Eyben, M. Wollmer, and B. Schuller. openSMILE – the Munich versatile and fast open-source audio feature extractor. In Alberto Del Bimbo, Shih-Fu Chang, and Arnold W. M. Smeulders, editors, *Proceedings of the 18th International Conference on Multimedia 2010, Firenze, Italy, October 25–29, 2010*, pages 1459–1462, 2010.

6 J. Frattaroli. Experimental disclosure and its moderators: A meta-analysis. *Psychological Bulletin*, 132:823–865, 2006.

7 F. Goldman-Eisler. *Psycholinguistics: Experiments in spontaneous speech*. Academic Press, New York, 1968.

8 J. H. Kahn, R. M. Tobin, A. E. Massey, and J. A. Anderson. Measuring emotional expression with the linguistic inquire and word count. *American Journal of Psychology*, 120:263–286, 2007.

9 M. L. Knapp and J. A. Hall. *Nonverbal Communication in Human Interaction*. Wadsworth Publishing, Boston, MA, USA, 2009.

10 G. Littlewort, J. Whitehill, T. Wu, I. R. Fasel, M. G. Frank, J. R. Movellan, and M. S. Bartlett. The computer expression recognition toolbox (CERT). In *Ninth IEEE International Conference on Automatic Face and Gesture Recognition (FG 2011), Santa Barbara, CA, USA, 21–25 March 2011*, pages 298–305, 2011.

11 D. P. McAdams. Personal narratives and the life story. In O. John, R. Robins, and L. A. Pervin, editors, *Handbook of personality: Theory and research*, pages 241–261. Guilford Press, 2008.

12 J. W. Pennebaker, R. E. Booth, and M. E. Francis. Linguistic inquiry and word count (liwc), 2007. `http://www.LIWC.net`.

13 J. W. Pennebaker and C. K. Chung. Expressive writing and its links to mental and physical health. In H. S. Friedman, editor, *Oxford Handbook of Health Psychology*, pages 417–437, New York, NY, USA, 2011. Oxford University Press.

14 S. Rude, E. M. Gortner, and J. Pennebajer. Language use of depressed and depression-vulnerable college students. *Cognition & Emotion*, 18:1121–1133, 2004.

15 E. E. Shriberg. *Preliminaries to a Theory of Speech Disfluencies*. PhD thesis, University of California at Berkeley, 1994.

16 G. J. Westerhof and E. Bohlmeijer. Life stories and mental health: The role of identifcation processes in theory and interventions. *Narrative Works*, 2(1):106–128, 2012.

Representing and Evaluating Legal Narratives with Subscenarios in a Bayesian Network*

Charlotte S. Vlek[1], Henry Prakken[2, 3], Silja Renooij[3], and Bart Verheij[1]

1 **Institute of Artificial Intelligence**
University of Groningen
Groningen, The Netherlands
`c.s.vlek@rug.nl, b.verheij@ai.rug.nl`
2 **Faculty of Law,**
University of Groningen
Groningen, The Netherlands
3 **Department of Information and Computing Sciences,**
Utrecht University
Utrecht, The Netherlands
`{h.prakken,s.renooij}@uu.nl`

⎯ Abstract ⎯

In legal cases, stories or scenarios can serve as the context for a crime when reasoning with evidence. In order to develop a scientifically founded technique for evidential reasoning, a method is required for the representation and evaluation of various scenarios in a case. In this paper the probabilistic technique of Bayesian networks is proposed as a method for modeling narrative, and it is shown how this can be used to capture a number of narrative properties.

Bayesian networks quantify how the variables in a case interact. Recent research on Bayesian networks applied to legal cases includes the development of a list of legal idioms: recurring substructures in legal Bayesian networks. Scenarios are coherent presentations of a collection of states and events, and qualitative in nature. A method combining the quantitative, probabilistic approach with the narrative approach would strengthen the tools to represent and evaluate scenarios.

In a previous paper, the development of a design method for modeling multiple scenarios in a Bayesian network was initiated. The design method includes two narrative idioms: the scenario idiom and the merged scenarios idiom. In this current paper, the method of [34] is extended with a subscenario idiom and it is shown how the method can be used to represent characteristic features of narrative.

1 Introduction

In a criminal trial, narrative can provide a context for what happened. A story or scenario gives a coherent presentation of the states and events around a crime and (implicitly) how

---

these led to the evidence that was found. It is then up to a judge or juror to evaluate whether there is a scenario sufficiently supported by evidence to believe that this is in fact what happened. In order to develop a scientifically founded method for reasoning with evidence in a legal case, we propose to represent and evaluate crime scenarios in a Bayesian network.

Narrative has been one of three dominant approaches in the literature on reasoning with evidence, next to argumentation and probability [17]. Recently, Verheij [33] proposed to integrate the three approaches. In a previous project [5, 4, 6, 8], a hybrid theory for stories and arguments in legal cases was developed. Currently, we are working on the connection between probabilistic models and narrative. We build upon the recent application of Bayesian networks to legal cases.

A Bayesian network represents a joint probability distribution over a collection of variables and consists of a graph, expressing the connections between variables in the domain, and underlying probability tables for these variables. As such, a Bayesian network can be used to represent the information that is available about a case, including evidence, hypothetical events and their connections.

In the legal field, the construction of a Bayesian network for a case is not at all straightforward. The network should include variables relevant to the case, but these differ from case to case. This contrasts with, for example, the medical field, where Bayesian networks have been used successfully to determine the most probable disease given a set of symptoms. For a given set of symptoms, a doctor knows in advance which tests will be relevant to perform and what their possible outcomes are, and a preconstructed Bayesian network can be selected. Unlike this relatively closed world of medical diagnoses, the legal field deals with a quite open world, where unpredictable relevant variables may turn up. For example, a yellow car passing by the evening before a burglary can seem irrelevant, until it turns out that it was driven by one of the suspects, inspecting the property before breaking in.

The application of Bayesian networks to legal cases has received quite some attention in recent research. Keppens [18] studied the combination of arguments and Bayesian networks in the context of law. Hepler, Dawid and Leucari [16] proposed the idea of often recurring substructures in the graph of a Bayesian network, which is also the basis of work by Fenton, Neil and Lagnado [14, 13, 20]. Fenton et al. developed a list of legal idioms, substructures that often occur in Bayesian networks for legal cases. Such idioms can be regarded as building blocks for a network, representing basic patterns in evidential reasoning. We intend to develop narrative idioms for representing scenarios in a legal case.

In our previous paper, we proposed to represent crime scenarios in a Bayesian network with the use of two narrative idioms: the scenario idiom for modeling scenarios, and the merged scenarios idiom for modeling multiple scenarios in one Bayesian network. Furthermore, we provided an initial sketch of a procedure for constructing such a Bayesian network. In this paper, we add a third idiom, the subscenario idiom, and we discuss how the scenario idiom and the subscenario idiom can be used to capture a number of narrative properties as they have been discussed in the literature. The contributions of this paper are threefold: (1) we give an analysis of characteristic features of narrative following recent work in the emerging field of computational narrative; (2) we extend the design method with a subscenario idiom; and (3) we show how the extended method can be used to represent the characteristic features from the analysis of narrative: structure, coherence, plausibility and the use of commonsense knowledge.

This paper is organized as follows: Section 2 is a literature study on narrative and its properties. In Section 3 some preliminaries on Bayesian networks are presented. In Section

(a) Substories from the Anchored Narratives Theory ([9] as adapted by Verheij [32]).

(b) An example of a story with substories from [7]. Arrows are pointed up since they represent inferences made from evidence and common knowledge.

Figure 1 Substories.

4 our design method and its development so far is presented. We extend this method with the new subscenario idiom in Section 5. Section 6 returns to the narrative properties, and the way in which these can be modeled with our extended design method. The paper is concluded in Section 7.

2 Narrative and its properties: a literature study

Literature on narrative spans a broad range of interests, from folk tales [27] to computer games [36] and TV-shows [22], and from determining the underlying plot of a story [21] to parsing and understanding a text [23]. This paper is concerned with legal stories, which we call scenarios: coherent collections of states and events, describing what can have happened around a supposed crime. Typically, there are multiple scenarios describing various accounts of what happened, and it is up to a judge or juror to find out which is true, based on the available and admissible evidence. Recent research on narrative in legal cases includes the development of a hybrid theory for stories and arguments [5] and simulating or animating a specific scenario with agents [31]. We focus on representing and evaluating various scenarios for a case.

The value of narrative in legal applications has been investigated in terms of various properties of narrative, such as a narrative's coherence, plausibility and the fact that it builds upon commonsense knowledge of the world. The sections below treat several properties of narrative that can be found in the literature, where reports on the 2009 workshop on the computational modeling of narrative [29, 15] served as a starting point.

2.1 Narrative structure

Three common denominators amongst representations of narrative are [15, 29]:
1. narratives have to do with sequences of events;
2. narratives have a hierarchical structure;
3. narratives are (eventually) grounded in commonsense knowledge of the world.

Following items (1) and (2), we take a scenario to be a collection of states and events with some coherent structure, either sequential or built up from subscenarios. Item (3), concerning commonsense knowledge, will be discussed in Section 2.4.

In a sequential structure, narrative is viewed as 'a succession of happenings' [29]. Schank and Abelson in their famous Script Theory [30] claim that for a proper text, a 'causal chain' can be constructed to represent it. Our case study [34] was done with this perspective in mind.

Alternatively, stories can be viewed as built up from substories. The idea of substories is prominently present in the Anchored Narratives Theory by Wagenaar et al. [35], see Figure 1a for their schematic representation of a story. There, the main story is ultimately anchored in commonsense generalizations. An example of a story with substories, taken from [7], is the following (see Figure 1b, where the anchoring is in evidence rather than in commonsense generalizations): Julius and Peter had a fight. This led to Julius firing a gun at Peter, who died as a result of this gunshot. This story consists of substories about Julius and Peter having a fight, Julius firing a gun at Peter and Peter dying of the gunshot.

2.2 Coherence

In the previous section, a scenario was said to be a collection of states and events with some *coherent structure*. In this section we further explore this notion of coherence, in light of scripts or story schemes (Section 2.2.1) and the transfer of evidential support as a consequence of narrative coherence (Section 2.2.2)

2.2.1 Scripts or story schemes

In an attempt to elucidate what makes a story or scenario coherent, one can study it from the perspective of Schank and Abelson's scripts [30]. Their scripts are used to explain how a listener can understand a story, and fill up the gaps that were left out when the story was told. Or as they say: "the meaning of a text is more than the sum of the meanings of individual sentences."

Schank and Abelson famously illustrated their theory with the example of a restaurant script: when a story is told about someone having dinner in a restaurant, the listener recognizes these events because he or she has a script of a typical restaurant visit in mind. This makes it possible for the listener to infer details that were omitted in the story. For example, when the story includes 'after ordering the food, he ate it', the listener will infer that between those two events, the waitress brought him his food.

A script is much like a 'template' for what elements a story (about a restaurant, for example), can or should contain. On the one hand, a listener uses this to make small inferences and fill up the gaps in a story. On the other hand, and this is not so much emphasized by Schank and Abelson [30], the storyteller makes sure that his story is perceived as *coherent* by adhering to a script. This idea of a template for what makes a complete story, can also be found in Pennington and Hastie's ideas on completeness [26] of a story, which led to Bex's story schemes [5].

Pennington and Hastie [25, 26] divide the coherence of narrative into three factors: consistency, completeness and plausibility. For them, the consistency of a story means that there should be no contradictions within the story. Plausibility is used to describe how well the story fits in with our knowledge of the real world. This will be discussed more elaborately in Section 2.3 below. Finally, completeness is 'the extent to which a story has all of its parts' [26] and can be regarded as a measure of how well a story follows a script.

In previous work by Bex and colleagues [5, 4, 6, 8], Pennington and Hastie's ideas on narrative coherence resurfaced in the formal setting of a hybrid theory for arguments and stories dealing with evidence in legal cases. There, a story is taken to be a coherent sequence of states and events. Arguments can be used to reason about the quality of the story: built upon evidential data available in the case, arguments can support states and events or causal connections in the story. Finally, arguments can also be used to reason about how well a story fits in and completes a so-called story scheme. The hybrid theory thereby implements the concept of a story's completeness.

Due to story schemes, a story for a legal case usually involves more states or events than what can be inferred directly from the available evidence. This is a valuable property in legal applications: it can lead to the finding of new evidence. In the hybrid theory, the notions of evidential gaps and story consequences are introduced. These refer to states or events that remain unsupported by evidence (evidential gaps) and new evidence that is found by trying to fill up these evidential gaps (story consequences).

2.2.2   Transfer of evidential support

A story or scenario is more than the sum of its parts. Separately, each state or event may seem uninteresting, or irrelevant to the case. By putting the states and events into a coherent whole and providing evidence for some of them, the scenario can be strong enough to make us believe in an event for which there is no direct evidence. This is illustrated by the following example:

▶ **Example 1.** We consider a famous Dutch case (known as 'De Deventer Moordzaak')[1], in which a widow was murdered. Her accountant was convicted for the murder, but according to some legal experts this was unjust. One of their arguments [11] presents an alternative scenario consisting of a number of small observations of the crime scene as it was found after the murder. In the original scenario for which the accountant was convicted, the suspect called the widow on the phone at 20:36 and drove to her home to kill her. In the alternative scenario, the killer must have been in the house much earlier than the accountant could have been given the phone call. In this scenario, the widow was doing the dishes and hadn't finished writing her shopping list when she was interrupted by the killer. Due to her strict routine this must have been shortly after the end of the eight o'clock news. The ingredients for this alternative scenario are small observations of seemingly unrelated details, such as an open notebook and pen on the table (she hadn't finished her shopping list) and the widow's apron on a chair in the conservatory (she was doing the dishes when she was interrupted by the doorbell). The neighbors testified that the widow always had a very strict routine, and together with the aforementioned details this leads to a coherent alternate scenario.

In this murder case example, the factor of interest is what time the killer entered the house. Given the time of the phone call, it would have taken the accountant quite some time to drive to the widow's house, giving her the time to finish her dishes and her shopping list. In the alternative scenario, someone else must have been the killer. There was no direct evidence for this specific event, but by presenting a coherent scenario with events for which evidence is available (such as, the widow was disturbed while she was doing the dishes), it can still become believable. It is this manifestation of coherence, which we shall refer to as *transfer of evidential support*, that we want to capture in our models.

---

[1] Information about this case can be found on www.rechtspraak.nl with code LJN BA 1024.

2.3 Plausibility

Plausibility of a story is a term often discussed in literature on narrative and law (for example, in [26, 35, 5]. In a criminal trial, scenarios are often highly unlikely, such as an alibi that just seems hard to believe. But even a very implausible scenario can become probable when there is enough evidence to support it. It is then up to the judge or juror to take this evidence into consideration and decide which scenario is probable enough to assume that this is what happened.

Pennington and Hastie describe the plausibility of a story as "the extent to which the story is consistent which knowledge of real or imagined events in the real world" [26]. Bex formalizes this by expressing a story's plausibility in terms of how many elements of the story are supported by commonsense knowledge. The key idea of plausibility is that a story is plausible when as a whole, it seems credible to us given our knowledge of (and experience with) the real world.

2.4 Commonsense knowledge

In order to understand narrative, a listener needs commonsense knowledge about the world. This was already a factor in the idea of scripts to understand stories, and in the concept of plausibility as described above. The use of commonsense knowledge was mentioned as one of the three common denominators of narrative. According to Bex and Verheij [8, 7], commonsense knowledge can be captured in either story schemes or argument schemes.

Commonsense knowledge plays an important role in the Anchored Narratives Theory by Wagenaar et al. [35]. There, a story in a criminal trial should be firmly anchored in commonsense knowledge in the form of generalizations such as 'an expert witness usually speaks the truth'. Bex's hybrid theory [5] is centered around the idea that a story should be supported by evidence and commonsense knowledge.

2.5 Summarizing: properties of narrative

To summarize, the following properties of narrative have been discussed in this section:
1. Narrative structure: sequential or built up from substories;
2. Coherence, manifested in three key features:
   a. A script or story scheme that serves as a template for a story;
   b. Evidential gaps and story consequences: events unsupported by evidence (evidential gaps) and the finding of new evidence (story consequences) as a result of these gaps;
   c. Transfer of evidential support: evidence for one element of the story can increase the belief in the entire story and thereby all elements of the story;
3. Plausibility: the extent to which a story seems credible to us given our knowledge of (and experience with) the real world;
4. Commonsense knowledge: the basic knowledge needed to understand the story.
In Section 6, these properties will be further discussed, including how they are captured in the design method presented in this paper.

3 Bayesian networks in legal cases

A Bayesian network consists of a graph (such as in Figure 2) and probability tables (such as Tables 2a and 2b). The nodes in the graph represent variables in the domain: a Bayesian network for a legal case typically contains hypotheses (such as `Fingerprints X`, abbreviated as `FP X`, describing that X's fingerprints were found at the crime scene), intermediate nodes

Figure 2 Example of a Bayesian network.

Table 1 Examples of probability tables. `Fingerprints X` was abbreviated to `FP X`.

FP X = y	FP X = n
0.2	0.8

(a) Prior probability $P(\texttt{FP X})$.

	FP X = y	FP X = n
Fingerprint match = y	0.9	0.01
Fingerprint match = n	0.1	0.99

(b) CPT for $P(\texttt{Fingerprint match} \mid \texttt{FP X})$

and evidential nodes (in Figure 2, `Fingerprint match` describes that the police found a match between the fingerprints found at the crime scene and suspect X).

With the arrows between nodes, dependencies and independencies between variables are shown in the graph. These arrows are often thought of as representing causality, which can be helpful when constructing a Bayesian network [24]. However, the arrows represent correlation, not causality [10]. The conditional probability tables (CPT's) contain the probabilities for a node conditional on its predecessors (Table 2b shows the conditional probability table for `Fingerprint match`). A node with no predecessors contains the (unconditional) probabilities for each value of the variable (see Table 2a). Such probabilities are called *prior probabilities*.

A Bayesian network is a representation of a joint probability distribution (JPD) [19]. The JPD contains the probabilities for each combination of values of variables, such as $P(\texttt{Fingerprint match = n}, \texttt{Fingerprints X = y})$. From a Bayesian network, the numbers in the joint probability distribution can be retrieved, as well as any prior or posterior probability of interest.

After constructing the Bayesian network, the evidence nodes can be *instantiated* in the network: the probability of the appropriate value of the evidential nodes is set to 1, and this information is propagated through the network, leading to updated (posterior) probabilities for the other nodes. There are tools available for such calculations, such as GeNIe 2.0.[2]

Bayesian networks are often used as a compact representation of a joint probability distribution. An advantage of a Bayesian network is the insight that the graphical structure provides into the connections between the variables. Though a Bayesian network requires less numbers to be made explicit, both a JPD and a Bayesian network require full information about the probabilities in the domain. Eliciting these numbers is a known issue for Bayesian networks. A number of methods for finding these numbers, or guiding experts to find these numbers, are available [28].

There is an ongoing debate about the use of Bayesian methods in court. There has been a ruling by the Court of Appeal in the UK in 2010, stating that Bayes' theorem should not be used in evaluating evidence, except for DNA and 'possibly other areas where there is a firm statistical base' [12]. On the other hand, a member of the Supreme Court in the Netherlands, together with the Netherlands Forensic Institute (Nederlands Forensisch Instituut, NFI), recently advocated the use of Bayesian thinking [1, 2, 3].

---

[2] GeNIe 2.0 is available for free on genie.sis.pitt.edu.

(a) The scenario idiom

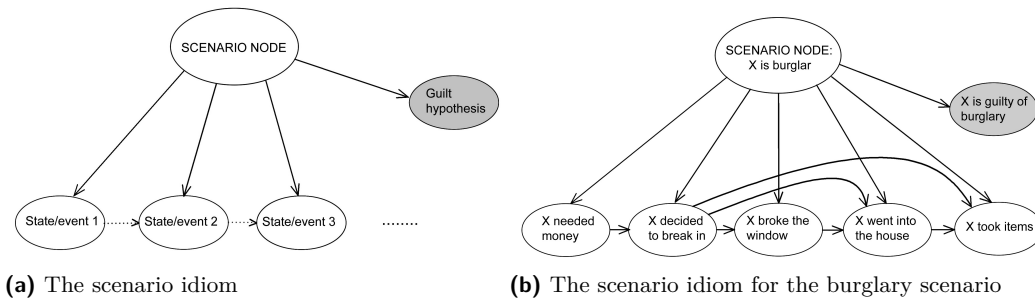**(b)** The scenario idiom for the burglary scenario

■ **Figure 3** The scenario idiom. Dotted lines denote possible connections between states and events.

Arguments against the use of Bayesian networks in court concern the known problem of eliciting the numbers in the probability tables, and it is yet to be investigated how well the methods mentioned above can help in the specific area of legal applications. Another argument is that, even when these numbers are known, it is questionable whether they can be used to make decisions about one individual. Being aware of these arguments, we intend the Bayesian networks resulting from our method to be used as a tool to compare scenarios and their evidential support, rather than to calculate absolute numbers. Furthermore, we explicitly do not intend to make the decision for a judge or juror; the resulting network is meant to advise and provide insight into the case.

4   Representing scenarios in a Bayesian network

In this section we review our design method from [34] for modeling crime scenarios in a Bayesian network. The goal of this method is to represent multiple scenarios concerning a crime in one network. We focus on constructing the graph for the Bayesian network, modeling the relevant variables of a scenario in a coherent structure. Our design method as developed so far has two narrative idioms: the scenario idiom and the merged scenarios idiom. In Section 5, the design method will be extended with a third idiom.

The procedure from [34] for constructing a Bayesian network consists of the following four steps (more elaborately discussed in [34]): (1) collect all relevant scenarios, (2) model each scenario using the scenario idiom (or the subscenario idiom), (3) merge these idioms into one large Bayesian network with the merged scenarios idiom, and (4) add the evidence to the network. We assume that the admissibility of the evidence has been established before constructing this model. Evidential nodes are modeled using Fenton, Neil and Lagnado's evidence accuracy idiom and their idioms about dependency between evidence [14].

In Section 4.1 we review the scenario idiom, which was used in [34] to model a scenario as a sequence of states and events in a Bayesian network. In Section 4.2 we treat the merged scenarios idiom, and how it can be used to merge multiple scenarios in one network. [34] gives further details about the design method, including a case study.

4.1   The scenario idiom

The scenario idiom is intended to model a scenario as a whole, capturing its coherence as described in Section 2.2. To do this, we model connections between states and events in the scenario, and we include a scenario node to model the underlying coherence.

■ **Table 2** Probability tables for the scenario idiom. `Scenario Node` was abbreviated to `ScN`.

	ScN = y	ScN = n
Event = y	1	x
Event = n	0	1-x

(a) Conditional probability table for a node 'Event'

	ScN = y	ScN = n
Guilt hypothesis = y	1	0
Guilt hypothesis = n	0	1

(b) Conditional probability table for the guilt hypothesis
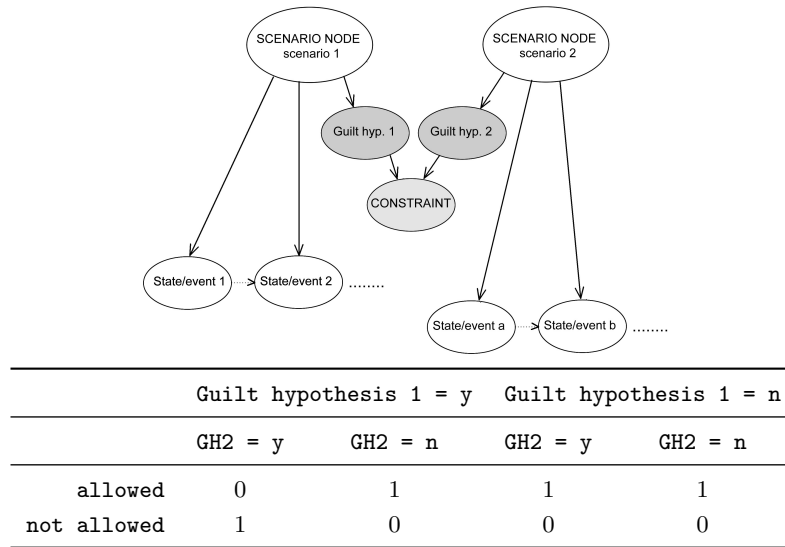
Consider the following scenario about a burglary:

Suspect X needed money, so X decided to break in. X broke the window of the house, went in and took some items from the house.

Figure 3b (from the case study in our previous paper [34]) shows how this scenario can be represented as a sequence of states and events. The figure illustrates the idea of the scenario idiom: the states and events are represented as nodes in the network, with connections between states and events drawn as arrows. The underlying probability tables express how certain these connections are. For example, a lack of money does not always lead to the decision to break in.

Furthermore, a scenario node is included, which is connected to all elements of the scenario. This scenario node is needed to model a scenario's coherence (see also Section 6.2). Finally, there is a node for the guilt hypothesis, describing who committed what crime: this node is included because this is ultimately the variable a judge or juror is interested in.

A general version of the scenario idiom is shown in Figure 3a. The idiom is constructed as follows:

1. Each state or event in the scenario is represented as a binary node with values 'yes' and 'no' in the network.
2. When there are connections between states or events in the scenario, arrows are drawn between the corresponding nodes (see [24] for what constitutes a connection between variables). Note that the connections between states and events within a scenario do not neccessarily need to form a sequence; one state or event can be connected to multiple elements of the scenario (not shown in Figure 3a). However, representing more complex connections within a scenario will be easier with the subscenario idiom (see Section 5).
3. A scenario node with values 'yes' and 'no' is included in the network. Arrows are drawn from this scenario node to each of the states or events in that scenario.
4. The probability table for the scenario node expresses the probability that the scenario is true without taking any of the evidence into account. This number corresponds to the plausibility of the scenario and is a subjective number that can be estimated by a judge or juror.
5. The conditional probability table for a state or event node depends on the connections of this state or event with the rest of the scenario. When the node is connected to other elements of the scenario, the numbers should be filled in accordingly. With no other connections, the probability table will look like Table 3a. The left column shows the logical relation that the event is an element in the scenario: when the scenario is true, all its elements must be true. The right column (when the scenario node is not true) is less straightforward. It expresses the probability that the event took place when the scenario as a whole is not true. These numbers in the right column are crucial for the evidential support: the higher the number for $P(\texttt{Event = yes} \mid \texttt{ScN = no})$ (x in the upper right of

	Guilt hypothesis 1 = y		Guilt hypothesis 1 = n	
	GH2 = y	GH2 = n	GH2 = y	GH2 = n
allowed	0	1	1	1
not allowed	1	0	0	0

Figure 4 The merged scenarios idiom and the conditional probability table for the constraint node. `Guilt hypothesis 2` was abbreviated to `GH2`.

the table), the *lower* the evidential support. This makes sense: when an event is quite likely to happen even when the scenario is not true, knowing that it happened has less of an influence on the probability of the entire scenario (see Section 6.2 for a more elaborate discussion of this point).

6. Finally, a guilt hypothesis is included, stating briefly what the scenario describes more elaborately: who committed what crime. Now that only one scenario is modeled (this will change when the merged scenarios idiom is used), the conditional probability table for the guilt hypothesis is straightforward: it follows logically from the scenario node. When the scenario node is true, the guilt hypothesis is true, and when the scenario node is not true, neither is the guilt hypothesis (see Table 3b).

4.2   The merged scenarios idiom

The merged scenarios idiom is used to combine multiple scenarios, in order to model them in one Bayesian network. The idiom puts a constraint on the guilt hypotheses, making sure that they cannot be true simultaneously. Therefore, a crucial step in the design method is to make sure that all guilt hypotheses are mutually exclusive or equal. Then when two scenarios are merged with the merged scenarios idiom, equal guilt hypotheses are represented by one node[3] and the constraint is put on mutually exclusive guilt hypotheses.

A constraint node is a common technique [19] to make sure that two or more nodes cannot be true simultaneously. It has values `allowed` and `not allowed` and there are arrows from the nodes that need to be constrained to the constraint node. The conditional probability table expresses logically that when more than one of the parent nodes is true, the constraint

---

[3] All scenario nodes that were connected to any of the original nodes will now be connected to the one new node. The probability table for this new node expresses that the guilt hypothesis is true when at least one of the scenarios connected to it is true.

(a) The subscenario idiom   **(b)** The subscenario idiom for the burglary scenario

Figure 5 The subscenario idiom. The dotted lines denote possible connections between states and events or between subscenarios.

node has value 'not allowed' (see the table in Figure 4). By instantiating the value to 'allowed', it can never be the case that more than one parent node is true at the same time.

When merging two scenarios, it may be the case that certain states or events in different scenarios contradict each other. In that case, a constraint needs to be put on these states or events; this constraint node is exactly like the constraint node from the merged scenarios idiom.

Furthermore, different scenarios may overlap, containing the same states or events, or even the same subscenario. When this happens, there will be separate nodes in different (sub)scenario idioms describing equal states, events or subscenarios. These will be replaced by one node describing this state, event or subscenario, which is then connected to all nodes that the original nodes were connected to.

The merged scenarios idiom in Figure 4 shows the merging of two scenario idioms. Analogously, the merged scenarios idiom can be used to merge two subscenario idioms or a combination of a subscenario idiom and a scenario idiom.

5 Extending the method: the subscenario idiom

In this section we introduce the scenario idiom in order to represent scenarios with a structure of subscenarios such as in Figure 1a.

5.1 The subscenario idiom

In the burglary example for the scenario idiom, the state 'X decided to break in' was connected to multiple states to express that it was the motive behind several actions of the burglar. With the subscenario idiom this can be modeled as in Figure 5b. There, the subscenario 'X needed money so X decided to break in' now serves as a motive for the subscenario 'X broke the window, went into the house and took some items'.

The ideas behind the subscenario idiom are closely related to the ideas behind the scenario idiom in the previous section. The coherence of an entire scenario is again modeled with use of a scenario node, which is connected to all elements of the scenario. In addition, there is now a level between the scenario node and the state or event nodes: the subscenario nodes. A general version of the subscenario idiom is shown in Figure 5a. It can be constructed in

■ **Table 3** Conditional probability table for a subscenario node. `Scenario node` was abbreviated to `ScN`.

	ScN = y	ScN = n
Subscenario node 1= y	1	0
Subscenario node 1= n	0	1

■ **Table 4** Conditional probability table for a subscenario node with a connection to another subscenario node. `Scenario node` was abbreviated to `ScN` and `Subscenario node 1` to `Sub1`.

	ScN = y		ScN = n	
	Sub1 = y	Sub1 = n	Sub1 = y	Sub1 = n
Subscenario node 2= y	1	0	0	0
Subscenario node 2= n	0	1	1	1

the same way as the scenario idiom, adding the following items to the list from Section 4.1 and changing items 3 and 5 from that list to read as below:

A. For a collection of states and events that form a subscenario, a subscenario node with values 'yes' and 'no' is included. An arrow is drawn from the subscenario to each state or event in the subscenario.

B. When there are connections between subscenarios, arrows are drawn between the corresponding subscenario nodes.

C. The conditional probability tables for the subscenario nodes express a logical relation: when the scenario node is true, all subscenario nodes must be true. When a scenario node is not true, none of the subscenario nodes are true. This leads to a probability table as in Table 3, or in the case of a connection between subscenarios, Table 4. In the case of the scenario node is false, we assume that a subscenario cannot be true in itself, but that it really needs the entire scenario to be true. This choice was made for technical reasons, to ensure that the subscenario idiom amounts to the same as the scenario idiom when all other connections are the same. See also Section 5.2.

3. (adapted) A scenario node with values 'yes' and 'no' is included in the network. Arrows are drawn from the scenario node to each of the subscenario nodes and any state or event in that scenario that is not connected to a subscenario node.

5. (adapted) The conditional probability table for a state or event node depends on the connections of this state or event with the rest of the scenario. When there is only a connection with the subscenario node, the numbers are filled in as follows: when the subscenario node is true, the state or event logically follows. When the subscenario node is not true, the numbers express the probability that this particular state or event would occur without the subscenario being true.

5.2 The subscenario idiom versus the scenario idiom

The scenario idiom and the subscenario idiom both represent states and events in a scenario as nodes, adding a scenario node to model the coherence of the scenario. However, the subscenario idiom includes another level in between the state and event nodes and the scenario node: the subscenario nodes. The subscenario nodes make it easily expressible that a collection of states or events (one subscenario) is connected to another collection in its

■ **Figure 6** A representation with the subscenario idiom (left) and one with the scenario idiom (right).

entirety: the decision to break in to a house can be connected to the entire subscenario of actually breaking in. Another advantage of using subscenarios is that when modeling multiple scenarios, the same subscenario can be connected to multiple scenario nodes.

When a scenario with no complicating connections between subscenario nodes (unlike the burglary example) is represented in the scenario idiom and in the subscenario idiom, there will be no difference in interpretation when calculating the probabilities. In particular, for a situation as shown in Figure 6, the probabilities $P(\texttt{Event1} \mid \texttt{Scenario node})$ for all values of the scenario node and Event1 will be the equal in both networks. For example, consider $P(\texttt{Event1=y} \mid \texttt{Scenario node=n})$ in both networks.

First note that by construction, the conditional probability tables for the state and event nodes in the subscenario idiom consist of the same numbers as in the scenario idiom (see Table 3a), now as probabilities conditional on the subscenario nodes. Furthermore, the probability tables for the subscenario idioms consist only of zeros and ones. Therefore, for $P(\texttt{Event1=y} \mid \texttt{Scenario node=n})$ in the left network, we have

$$
\begin{aligned}
&P(\texttt{Event1=y} \mid \texttt{Scenario node=n}) \\
=\ &P(\texttt{Event1=y} \mid \texttt{Subscenario node1=y}) \\
&\quad \cdot P(\texttt{Subscenario node1=y} \mid \texttt{Scenario node=n}) \\
&+ P(\texttt{Event1=y} \mid \texttt{Subscenario node1=n}) \\
&\quad \cdot P(\texttt{Subscenario node1=n} \mid \texttt{Scenario node=n}) \\
=\ &P(\texttt{Event1=y} \mid \texttt{Subscenario node1=y}) \cdot 0 + P(\texttt{Event1=y} \mid \texttt{Subscenario node1=n}) \cdot 1 \\
=\ &P(\texttt{Event1=y} \mid \texttt{Subscenario node1=n})
\end{aligned}
$$

and this was set equal to $P(\texttt{Event1=y} \mid \texttt{Scenario node=n})$ (in the right network) in the probability tables.

Note that this calculation will not hold for the burglary example, since the connections between state and event nodes are not the same in the scenario idiom and the subscenario idiom. The connection from the event 'X decided to break in' to the event 'X broke a window' has now moved to the level of the subscenario nodes, which changes the interpretation of the scenario and thereby the probability tables. The subscenario version of the burgarly example is therefore really different from the example with the scenario idiom; this shows that the subscenario idiom can model a different interpretation of the scenario.
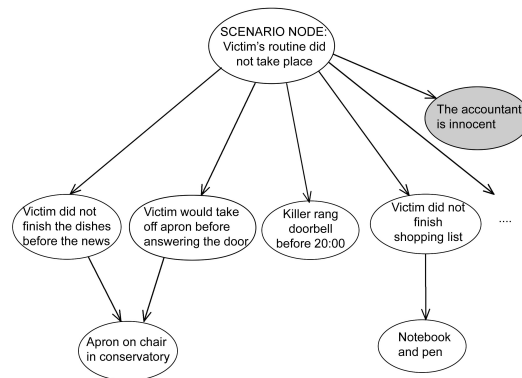
◾ **Figure 7** Partial scenario idiom for the transfer of evidential support example.

6   Representing narrative properties with the design method

Returning to the discussion of narrative and its properties from Section 2, in this section we explain how each of the properties listed in Section 2.5 can be treated using the representational techniques from Sections 4 and 5.

6.1   Narrative structure

With our design method, the collection of states and events of which a scenario consists serves as a basis for the domain of the Bayesian network: each state or event was represented in the network with a node. The structure of narrative (item **1** from the list in Section 2.5) is captured with either the scenario idiom (best suited for sequential narrative representation) or the subscenario idiom (for modeling subscenarios). It depends on the interpretation of the scenario which idiom can best be used to represent it.

6.2   Coherence

In the construction of the scenario idiom and the subscenario idiom, the scenario node was intended to capture the coherence of a scenario. In this section we discuss how the structures of these idioms relate to scripts or story schemes and the transfer of evidential support.

6.2.1   Scripts or story schemes

The concept of a script or story scheme (item **2a**) inspired the idea of the scenario idiom and the subscenario idiom. They can be regarded as templates for representing a scenario in a Bayesian network. The idioms as presented in this paper can be the basis for scripts or schemes for specific crimes, such as a scheme prescribing what elements a typical murder case scenario has. However, crimes in particular are out of the ordinary, so a corpus of typical crime schemes will always be only the starting point of a specific model for a case. Further research is needed on the use of scripts and schemes for legal cases.

The idea of evidential gaps and story consequences (item **2b**) fits particularly well with the technique of Bayesian networks. By convention, Bayesian networks are usually constructed such that the arrows are directed from cause to effect. Intuitively, the elements of the scenario seem to 'predict' the evidence. When a predicted piece of evidence is not available, the evidential node is nonetheless included in the network, but it is left uninstantiated: it

is not fixed to have value 'yes'. This means there is an evidential gap. In the investigative process of a crime, such obvious absence of evidence may lead to a search for this particular piece of evidence. When this evidence is then found, the evidential node can be instantiated to 'yes', and we have found a story consequence.

6.2.2   Transfer of evidential support

The scenario idiom and the subscenario idiom are constructed such that they can capture the transfer of evidential support (item **2c**), namely, via the scenario node. Since there is no direct evidence about the scenario node, it is never fixed on a value (it is never instantiated), leaving a connection between any two elements of the scenario connected to the scenario node through which they can influence each other (any pair of states or events in the scenario is d-connected (see [19]) via the scenario node).

In Section 2.2, the concept of transfer of evidential support was introduced with the example of the murder case involving a scenario with seemingly unconnected states or events. However, since these elements together formed a coherent scenario, a transfer of evidential support was possible: by providing evidence for some elements the belief in all elements of the scenario increases.

The example scenario described the victim not finishing her usual evening routine, doing the dishes, writing a shopping list, etcetera (more details were involved in the actual case). Instead, she was disturbed by the killer ringing her doorbell shortly after the eight o'clock news. Figure 7 shows part of the scenario idiom for this scenario (the dots suggest that the scenario actually involved more elements).

There is no direct evidence for the event that the killer came in soon after the news. However, there is evidence for the fact that the victim was interrupted from doing the dishes (the apron) and that she did not finish her shopping list (the notebook and pen). The corresponding nodes for these evidential data can be instantiated to have value 'yes', which leads to a higher probability for these particular events in the scenario. Since we have no direct knowledge about the scenario node itself (it is never instantiated), a higher probability for one of the events will lead to a higher probability of the scenario node being true, leading to a higher probability of the killer ringing the doorbell shortly after the eight o'clock news.

The transfer of evidential support thus proceeds via the scenario node. This has to do with the evidential support of a piece of evidence for the entire scenario: as the posterior probability for the scenario node changes, due to the evidential support for it, the posterior probability for all events in that scenario change simultaneously: they logically follow from the scenario node.

When a piece of evidence supports an element of a scenario, the posterior probability $P(\texttt{ScN=y} \mid \texttt{Event=y})$ that the scenario node is true given that the event is true will be different from the prior probability $P(\texttt{ScN=y})$. The more $P(\texttt{ScN=y} \mid \texttt{Event=y})$ differs from $P(\texttt{ScN=y})$, the stronger the evidential support is for the scenario. We propose to use the fraction $\frac{P(\texttt{ScN=y} \mid \texttt{Event=y})}{P(\texttt{ScN=y})}$ as a measure of the evidential support.

As mentioned in Section 4.1, the strength of the evidential support depends on the conditional probability tables for the events in the scenario (as in Table 3a), and in particular the number $P(\texttt{Event=y} \mid \texttt{ScN=n})$, denoted as x in the table. When this number is high, this means that the event has a high probability of taking place when the scenario node is not true. This leads to a lower evidential support. In particular, the probability of the scenario node taking place is changes *less* when the number x is closer to 1. To see this, we expand the probability $P(\texttt{ScN=y} \mid \texttt{Event=y})$ using Bayes' rule:

$$
\begin{aligned}
P(\texttt{ScN=y} \mid \texttt{Event=y}) \;\; &= \;\; \frac{P(\texttt{Event=y} \mid \texttt{ScN=y}) \cdot P(\texttt{ScN=y})}{P(\texttt{Event=y})} = \frac{1 \cdot P(\texttt{ScN=y})}{P(\texttt{Event=y})} \\[2mm]
&= \;\; \frac{P(\texttt{ScN=y})}{P(\texttt{Event=y}|\texttt{ScN=y}) \cdot P(\texttt{ScN=y}) + P(\texttt{Event=y}|\texttt{ScN=n}) \cdot P(\texttt{ScN=n})} \\[2mm]
&= \;\; \frac{P(\texttt{ScN=y})}{1 \cdot P(\texttt{ScN=y}) + x \cdot P(\texttt{ScN=n})} \\[2mm]
\frac{P(\texttt{ScN=y} \mid \texttt{Event=y})}{P(\texttt{ScN=y})} \;\; &= \;\; \frac{1}{P(\texttt{ScN=y}) + x \cdot P(\texttt{ScN=n})}.
\end{aligned}
$$

When the number x is close to 1, then the fraction on the right is close to 1 (since the prior probabilities of the scenario node add up to 1). This means that $P(\texttt{ScN=y} \mid \texttt{Event=y})$ is almost equal to $P(\texttt{ScN=y})$, so there is less evidential support. The smaller x is, the smaller the denominator, so the larger the fraction and thereby the evidential support. This is in line with our intuition: when an event can perfectly well take place without the scenario node being true, then knowing that the event took place contributes less to our belief in the scenario node.

6.3   Plausibility

Plausibility (item **3**) can be interpreted in terms of probability as the probability of the scenario without taking any evidence into account. In our representation, this is the prior probability for the scenario node having value 'yes'. This number is a subjective estimate, which can be provided by a judge or juror, and it is taken up in the probability table for the scenario node.

6.4   Commonsense knowledge

Commonsense knowledge (item **4**) in the form of generalizations underlying a connection between two states or events in a scenario is expressed numerically as a number in the probability table in our method: the more 'common' the connection is (such as, 'when a suspect breaks a window, he might leave fingerprints'), the higher the conditional probability connecting the two states or events. As a consequence, scenarios that are close to our commonsense knowledge will have stronger connections in the Bayesian network.

7   Conclusion

In this paper, a design method for modeling crime scenarios in a Bayesian network has been presented and extended, and connections with literature on narrative have been discussed. In our previous paper [34], the development of the design method was started. There, a procedure for constructing a Bayesian network based on scenarios was introduced, including two narrative idioms: the scenario idiom and the merged scenarios idiom. In this current paper, the method was extended with a subscenario idiom, such that both sequential scenarios and scenarios built up from subscenarios could be represented in the Bayesian network.

The notions of narrative coherence, plausibility and the use of commonsense knowledge have been interpreted in terms of the representational techniques developed. The transfer of evidential support (as a consequence of narrative coherence) is captured with the construction of the (sub)scenario idiom. The probabilities in the network reflect the plausibility and

the use of commonsense knowledge. By modeling these narrative properties, the Bayesian network can be used to evaluate and compare crime scenarios.

Further research is needed on the use of scripts and story schemes in the representation of a crime scenario. It deserves to be investigated how a corpus of schemes for crimes can serve as a starting point for the construction of specific models. Employing the concept of a scheme may help to systematize the construction of Bayesian networks for legal cases.

References

1 C. E. H. Berger and D. Aben. Bewijs en overtuiging: Een helder zicht op valkuilen. *Expertise en Recht*, 5/6:159–165, 2010.

2 C. E. H. Berger and D. Aben. Bewijs en overtuiging: Rationeel redeneren sinds Aristoteles. *Expertise en Recht*, 2:52–56, 2010.

3 C. E. H. Berger and D. Aben. Bewijs en overtuiging: Redeneren in de rechtszaal. *Expertise en Recht*, 3:86–90, 2010.

4 F. J. Bex. Analysing stories using schemes. In H. Kaptein, H. Prakken, and B. Verheij, editors, *Legal Evidence and Proof: Statistics, Stories, Logic*, pages 93–116, Aldershot, 2009. Ashgate Publishing.

5 F. J. Bex. *Arguments, Stories and Criminal Evidence, a Formal Hybrid Theory.* Springer, Dordrecht, 2011.

6 F. J. Bex, P. J. van Koppen, H. Prakken, and B. Verheij. A hybrid formal theory of arguments, stories and criminal evidence. *Artificial Intelligence and Law*, 18:123–152, 2010.

7 F. J. Bex and B. Verheij. Story schemes for argumentation about the facts of a crime. In Mark Alan Finlayson, editor, *Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium*, number FS-10-04 in AAAI Technical Reports, pages 6–13. Association for the Advancement of Artificial Intelligence, 2010.

8 F. J. Bex and B. Verheij. Solving a murder case by asking critical questions: An approach to fact-finding in terms of argumentation and story schemes. *Argumentation*, 26:325–353, 2012.

9 H. F. M. Crombag, P. J. van Koppen, and W. A. Wagenaar. *Dubieuze Zaken: De Psychologie van Strafrechtelijk Bewijs.* Uitgeverij Contact, Amsterdam, 1992.

10 A. P. Dawid. Beware of the DAG. In Isabelle Guyon, Dominik Janzing, and Bernhard Schölkopf, editors, *Causality: Objectives and Assessment (NIPS 2008 Workshop), December 12, 2008, Whistler, Canada*, number 6 in JMLR Workshop and Conference Proceedings, pages 59–86, 2009.

11 T. Derksen. *De Ware Toedracht, Praktische Wetenschapsfilosofie voor Waarheidszoekers.* Uitgeverij Veen Magazines BV, Diemen, 2010.

12 England and Wales Court of Appeal. EWCA Crim 2439: Regina v T, 2010. Decision 26 October 2010; published online by the British and Irish Legal Information Institute.

13 N. E. Fenton and M. Neil. Avoiding probabilistic reasoning fallacies in legal practice using Bayesian networks. *Australian Journal of Legal Philosophy*, 36:114–151, 2011.

14 N. E. Fenton, M. Neil, and D. A. Lagnado. A general structure for legal arguments using Bayesian networks. *Cognitive Science*, 37(1):61–102, 2013.

15 M. A. Finlayson, W. Richards, and P. H. Winston. Computational models of narrative: Review of a workshop. *AI Magazine*, 31(2):97–100, 2010.

16 A. B. Hepler, A. P. Dawid, and V. Leucari. Object-oriented graphical representations of complex patterns of evidence. *Law, Probability & Risk*, 6:275–293, 2004.

17 H. Kaptein, H. Prakken, and B. Verheij. *Legal Evidence and Proof: Statistics, Stories, Logic.* Ashgate Publishing Company, Aldershot, 2009.

18 J. Keppens. On extracting arguments from Bayesian network representations of evidential reasoning. In K. D. Ashley and T. M. van Engers, editors, *The 13th International Conference on Artificial Intelligence and Law, Proceedings of the Conference, June 6–10, 2011, Pittsburgh, PA, USA*, pages 141 – 150, New York, 2011. Association for Computing Machinery.

19 U. B. Kjaerulff and A. L. Madsen. *Bayesian Networks and Influence Diagrams.* Springer, New York, 2008.

20 D. A. Lagnado, N. E. Fenton, and M. Neil. Legal idioms: a framework for evidential reasoning. *Argument and Computation*, 4(1):46–63, 2013.

21 W. G. Lehnert. Plot units and narrative summarization. *Cognitive Science*, 5(4):293–331, 1981.

22 B. Löwe, E. Pacuit, and S. Saraf. Identifying the structure of a narrative via an agent-based logic of preferences and beliefs: Formalizations of episodes from CSI: Crime Scene Investigation™. In Michael Duvigneau and Daniel Moldt, editors, *Proceedings of the Fifth International Workshop on Modelling of Objects, Components and Agents. MOCA '09*, number 290/09 in FBI-HH-B, pages 45–63, 2009.

23 E. T. Mueller. Understanding script-based stories using commonsense reasoning. *Cognitive Systems Research*, 5:307–340, 2004.

24 J. Pearl. *Causality: Models, Reasoning and Inference.* Cambridge University Press, Cambridge, 2000.

25 N. Pennington and R. Hastie. Evidence evaluation in complex decision making. *Journal of Personality and Social Psychology*, 51(2):242–258, 1986.

26 N. Pennington and R. Hastie. Explaining the evidence: Tests of the story model for juror decision making. *Journal of Personality and Social Psychology*, 62(2):189–206, 1992.

27 Vladimir Yakovlevich Propp. *Morphology of the Folktale.* University of Texas Press, Austin, TX, 2nd edition, 1968. transl. L. Scott.

28 S. Renooij. Probability elicitation for belief networks: Issues to consider. *Knowledge Engineering Review*, 16(3):255–269, 2001.

29 W. Richards, M. A. Finlayson, and P. H. Winston. Advancing computational models of narrative. Technical report, Massachussetts Institute of Technology, 2009.

30 R. C. Schank and R. P. Abelson. *Scripts, Plans, Goals and Understanding, An Inquiry into Human Knowledge Structures.* Erlbaum, Hillsdale, 1977.

31 G. Sileno, A. Boer, and T. van Engers. Analysis of legal narratives: a conceptual framework. In B. Schäfer, editor, *Legal Knowledge and Information Systems – JURIX 2012: The Twenty-Fifth Annual Conference, University of Amsterdam, The Netherlands, 17–19 December 2012*, pages 143–146. IOS Press, 2012.

32 B. Verheij. Anchored narratives and dialectical argumentation, 2001. Presented at ICAIL-2001 Workshop on AI and Legal Evidence.

33 B. Verheij. Integrating argumentation, narrative and probability in legal evidence. In Mark Alan Finlayson, editor, *Proceedings of Computational Models of Narrative 2012*, pages 176–177, İstanbul, 2012.

34 C. S. Vlek, H. Prakken, S. Renooij, and B. Verheij. Modeling crime scenarios in a Bayesian network. In *The 14th International Conference on Artificial Intelligence and Law, Proceedings of the Conference*, 2013. In press.

35 W. A. Wagenaar, P. J. van Koppen, and H. F. M. Crombag. *Anchored Narratives: the Psychology of Criminal Evidence.* Harvester Wheatsheaf, Hemel Hempstead, 1993.

36 R. M. Young. Story and discourse: A bipartite model of narrative generation in virtual worlds. *Interaction Studies*, 8:177–208, 2007.