Report from Dagstuhl Perspectives Workshop 19072

The Role of Non-monotonic Reasoning in Future Development of Artificial Intelligence

Edited by

Anthony Hunter[1], Gabriele Kern-Isberner[2], Thomas Meyer[3], and Renata Wassermann[4]

1 **University College London, GB,** anthony.hunter@ucl.ac.uk
2 **TU Dortmund, DE,** gabriele.kern-isberner@cs.tu-dortmund.de
3 **University of Cape Town, ZA,** tmeyer@cs.uct.ac.za
4 **University of Sao Paulo, BR,** renata@ime.usp.br

Abstract

This report documents the program and the outcomes of Dagstuhl Perspectives Workshop 19072 "The Role of Non-monotonic Reasoning in Future Development of Artificial Intelligence". The workshop brought together researchers both from core topics and peripheral areas of non-monotonic reasoning (NMR), but also attracted researchers from other scientific domains in which recent developments have shown an increased relevance of NMR topics. The overall goal of this workshop was to reshape NMR as a core methodology for artificial intelligence being able to meet present and future challenges. Participants of this workshop discussed in what shape NMR would be useful for future AI, and how NMR can be developed for those requirements. The workshop started with brief survey talks and had some technical talks on central topics of NMR afterwards. These were followed by working groups on core aspects of NMR and potential links with learning. On the last day of the seminar, each working group presented their ideas and future plans. The workshop closed with a plenary discussion on the future of NMR.

Seminar February 10–15, 2019 – http://www.dagstuhl.de/19072

2012 ACM Subject Classification Computing methodologies → Artificial intelligence, Computing methodologies → Knowledge representation and reasoning, Computing methodologies → Nonmonotonic, default reasoning and belief revision, Computing methodologies → Probabilistic reasoning, Computing methodologies → Logic programming and answer set programming, Computing methodologies → Ontology engineering, Computing methodologies → Cognitive science, Computing methodologies → Machine learning

Keywords and phrases Artificial intelligence, Knowledge representation and reasoning, Non-monotonic, default reasoning and belief revision, Probabilistic reasoning, Logic programming and answer set programming, Ontology engineering, Cognitive science, Machine learning

Digital Object Identifier 10.4230/DagRep.9.2.73

Edited in cooperation with Ivan José Varzinczak

1 Executive Summary

Anthony Hunter (University College London, GB)
Gabriele Kern-Isberner (TU Dortmund, DE)
Thomas Meyer (University of Cape Town, ZA)
Renata Wassermann (University of Sao Paulo, BR)

Nonmonotonic reasoning (NMR) addresses a fundamental problem that classical logic methods in computer science encounter when modelling real-world problems: New information may not only extend previously held knowledge (this would correspond to a monotonic extension) but can drastically change knowledge in that conclusions turn out to be wrong and need to be withdrawn. Nonmonotonic phenomena are present in all areas of our everyday lives mostly due to uncertain and incomplete information, but also due to humans reasoning with restricted ressources; on the other hand, humans do very well in determining relevant contexts of reasoning, so reasoning from incomplete information only may well be on purpose and for sake of efficiency. Nowadays, with computer systems taking on increasingly sophisticated roles in our lives, the need for computational intelligence to be able to also reason in a nonmonotonic way becomes increasingly urgent.

The international Nonmononotonic Reasoning (NMR) workshops have provided a premier specialized forum for researchers in non-monotonic reasoning and related areas since 1984. Over the years, NMR topics and results have been developed in areas such as answer set programming, computational models of argument, and description logics for ontologies. However, research on core topics of NMR has been scattered into different subcommunities that no longer collaborate in depth on a regular basis. As a consequence, much time and effort for solving specific, but in principle similar problems is wasted, general relevance of proposed solutions is overlooked, and general methodological competence is no longer developed to the same degree as ten years ago.

This Perspectives Seminar brought together researchers both from core topics and peripheral areas of NMR, but also attracted researchers from other scientific domains in which recent developments have shown an increased relevance of NMR topics. More precisely, researchers from various subcommunities within computer science and engineering (e.g., artificial intelligence, classical and non-classical logics, machine learning, agent and multiagent systems) met in Dagstuhl, but also researchers from other disciplines like philosophy and psychology contributed to the seminar. The overall goal of this seminar was to reshape NMR as a core methodology for artificial intelligence being able to meet present and future challenges. For AI to progress from pattern recognition and machine learning to broader cognitive reasoning, it needs to have commonsense reasoning, and this in turn calls for a deeper understanding of NMR. So participants of this workshop discussed in what shape NMR would be useful for future AI, and how NMR can be developed for those requirements. We started the seminar with brief survey talks on answer set programming, belief revision, argumentation, argument mining, machine learning, conditional reasoning, description logics, as well as NMR and cognition, and had some technical talks on central topics of NMR afterwards. For the rest of the week, we had working groups on NMR and learning, NMR and cognition, engineering NMR, and commonsense reasoning. We let people freely choose which working groups they wanted to attend each day, which resulted in vivid discussions and a particularly dynamic exchange of ideas. On the last day of the seminar, each working group presented their ideas and future plans, and we closed this seminar with a plenary discussion on the future of NMR. This report shows brief summaries of the presentations and of the results of the working groups.

2 Table of Contents

Executive Summary
Anthony Hunter, Gabriele Kern-Isberner, Thomas Meyer, and Renata Wassermann 74

Overview of Talks

Conditional Approaches in NMR
James P. Delgrande 76

Tutorial on Probabilistic Logic Programming
Luc De Raedt 77

The Role of Answer Set Programs for Non-Monotonic Reasoning
Thomas Eiter 78

Reasoning about Exceptions in DL Ontologies
Laura Giordano 80

Logics and Human Reasoning: What do we Know?
Marco Ragni 81

Argumentation Systems: A Brief Glimpse
Guillermo R. Simari 81

Argument Mining: from Non-Monotonic Reasoning to Natural Language Processing and Back
Serena Villata 83

Working groups

Working Group on Commonsense Reasoning
Anthony Hunter 84

Working Group on Integrating ML and KR, and the Relevance of Prototypical Reasoning
Gabriele Kern-Isberner 85

Working Group on Implications for NMR and Cognition
Thomas Meyer 87

Working Group on NMR Engineering
Renata Wassermann 88

Participants 90

3 Overview of Talks

3.1 Conditional Approaches in NMR

James P. Delgrande (Simon Fraser University – Burnaby, CA)

A *conditional approach* in NMR is one that is intended to *model* a *defeasible* statement such as "birds fly". This is in contrast with most other approaches (such as default logic or circumscription) where the goal is to provide a nonmonotonic inference *mechanism*, in which a default could be encoded. So in a conditonal approach, defaults are objects that one can reason about. There are two broad approaches, nonmonotonic consequence relations, in which Gentzen-style rules for an operator $\mid\!\sim$ are specified, and conditional logic, in which a *binary modal operator* $\Rightarrow$ is developed.

3.1.1 Nonmonotonic Consequence Relations

The overall goal in these approaches is to develop a nonmonotonic consequence relation $\mid\!\sim$. This might be in order to study the properties of other formalisms or to directly study nonmonotonc inference principle. The emphasis in either case is on syntactic considerations, and axiomatic systems are the main object of interest.

Two types of relations have received significant attention, *preferential* and *rational* systems. Preferential systems are characterised as follows:

Ref: $\phi \mid\!\sim \phi$
LLE: If $\models \phi \equiv \psi$ and $\phi \mid\!\sim \gamma$ then $\psi \mid\!\sim \gamma$.
RW: If $\models \psi \supset \gamma$ and $\phi \mid\!\sim \psi$ then $\phi \mid\!\sim \gamma$.
And: If $\phi \mid\!\sim \psi$ and $\phi \mid\!\sim \gamma$ then $\phi \mid\!\sim \psi \wedge \gamma$
OR: If $\phi \mid\!\sim \gamma$ and $\psi \mid\!\sim \gamma$ then $\phi \vee \psi \mid\!\sim \gamma$
CM: If $\phi \mid\!\sim \psi$ and $\phi \mid\!\sim \gamma$ then $\phi \wedge \psi \mid\!\sim \gamma$

For rational systems, the following is added:
RM $\phi \mid\!\sim \gamma$ and $\phi \not\mid\!\sim \neg\psi$ then $\phi \wedge \psi \mid\!\sim \gamma$

3.1.2 Conditional logics

A conditional logic is a modal logic based on a binay modal operator $\Rightarrow$. The idea is to define in a Kripke structure what's meant by one world being *at least as normal as* another. For this, one can denote a (universal) accessibility relation on worlds by $\leq$, where $w_1 \leq w_2$ just if w_1 is at least as normal as w_2. Then in the two main approaches, $\leq$ is required to be either a partial or a total preorder. In either case, satisfaction in a model M at a world w is defined by:

$$M, w \models \gamma \Rightarrow \phi \quad \text{iff} \qquad \text{if} \quad [\gamma] \neq \emptyset \;\; \text{then} \;\; \min_{\leq}([\gamma]) \subseteq [\phi]$$

The resulting logics can be axiomatised as consisting of propositional logic along with:

ID: $\phi \Rightarrow \phi$
CC: $(\phi \Rightarrow \psi \wedge \phi \Rightarrow \gamma) \supset (\phi \Rightarrow \psi \wedge \gamma)$
RT: $\phi \Rightarrow \psi \supset (\phi \wedge \psi \Rightarrow \gamma \;\supset\; \phi \Rightarrow \gamma)$
CC': $(\phi \Rightarrow \gamma \wedge \psi \Rightarrow \gamma) \supset (\phi \vee \psi \Rightarrow \gamma)$

CM: $(\phi \Rightarrow \psi) \supset (\phi \Rightarrow \gamma \ \supset \ \phi \wedge \psi \Rightarrow \gamma)$
RCM: From $\psi \supset \gamma$ infer $\phi \Rightarrow \psi \ \supset \ \phi \Rightarrow \gamma$.

or the above along with the axiom:

CV: $(\phi \Rightarrow \gamma) \supset (\neg(\phi \Rightarrow \neg\psi) \ \supset \ \phi \wedge \psi \Rightarrow \gamma)$

The logics are sound and complete with respect to a partial preorder or total preorder for $\leq$, respectively. As well, preferential and rational consequence relations are interdefinable with conditional logics. However, these approaches do not handle irrelevant properties, for example $B \Rightarrow F \not\vdash B \wedge G \Rightarrow F$. This led to the *rational closure*; intuitively the formulas in a theory are ranked under the assumption that they are as normal as possible. The idea is that if there is no reason for a property to be relevant to a conditional, it is assumed to be irrelevant. In the approach, one obtains for example that a green bird flys. However the rational closure has problematic properties. Essentially, if an individual is exceptional in some way, then *all* other normality assertions are also inapplicable, and so nothing more can be concluded about that individual. While some solutions have been proposed (like the *lexicographic closure*), none appears to be wholly adequate.

3.1.3 Other Issues

These approaches have been extended and explored, e.g. to give probabilistic conditionals and to explore theorem proving. To date, approaches have focussed almost exclusively on normality conditionals; it would seem that there could be alternative approaches, dealing with e.g. deontic, counterfactual, or likelihood conditionals. As well, the problem of nested conditionals remains unaddressed, and first-order reasoning remains a challenge.

References

1 C. Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence Journal*, 68(1):87–154, 1994.
2 James Delgrande. What's in a default? In *Nonmonotonic Reasoning: Essays Celebrating its 30th Anniversary*. College Publications, 2011.
3 S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence Journal*, 44(1-2):167–207, 1990.
4 D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence Journal*, 55(1):1–60, 1992.
5 D. Makinson. General theory of cumulative inference. In *Proc. of the Second International Workshop on Non-Monotonic Reasoning*, volume 346 of *Lecture Notes in Artificial Intelligence*, pages 1–18. Springer Verlag, 1989.

3.2 Tutorial on Probabilistic Logic Programming

Luc De Raedt (KU Leuven, BE)

Joint work of Luc De Raedt, Angelika Kimmig

Probabilistic programming is emerging as a new paradigm for programming in which one combines the power of probabilistic modeling and learning with that of programming languages. The talk presented an introduction to this topic with a focus on probabilistic logic programming languages [2]. Probabilistic logic programming is closely related and extends probabilistic databases, and as such also fits the statistical relational artificial intelligence paradigm [1]. Throughout the tutorial we used the language ProbLog [3] to illustrate the key concepts. ProbLog is based on Sato's distribution semantics [5], which essentially assigns probabilities to ground facts, which then, together with the logic programs, defines a probability distribution over possible worlds. Various applications of probabilistic logic programming were presented, and also the newest results on the integration of ProbLog with neural networks, resulting in a framework for neuro-symbolic computation [4].

This tutorial was based on joint work with Angelika Kimmig.

References

1 Luc De Raedt, Kristian Kersting, Sriraam Natarajan, and David Poole. Statistical relational artificial intelligence: Logic, probability, and computation. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 10(2):1–189, 2016.
2 Luc De Raedt and Angelika Kimmig. Probabilistic (logic) programming concepts. *Machine Learning*, 100(1):5–47, 2015.
3 Daan Fierens, Guy Van den Broeck, Joris Renkens, Dimitar Shterionov, Bernd Gutmann, Ingo Thon, Gerda Janssens, and Luc De Raedt. Inference and learning in probabilistic logic programs using weighted Boolean formulas. *Theory and Practice of Logic Programming*, 15(3):358–401, 2015.
4 Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Deepproblog: Neural probabilistic logic programming. In *NeurIPS*, 2018.
5 T. Sato. A statistical learning method for logic programs with distribution semantics. In L. Sterling, editor, *Proceedings of the 12th International Conference on Logic Programming*, pages 715–729. MIT Press, 1995.

3.3 The Role of Answer Set Programs for Non-Monotonic Reasoning

Thomas Eiter (TU Wien, AT)

Answer Set Programming (ASP) is a well-known problem solving approach which has roots in Knowledge Representation, Databases & Logic Programming, and Non-Monotonic Reasoning (NMR). The term was coined by Lifschitz [13, 14] and the paradigm proposed by others at about the same time [15, 16]. At an abstract level, ASP relates to Satisfiability Solving (SAT) and Constraint Programming (CP), but has some distinctive features. A rich literature is available, with books, [1, 7, 9], handbook articles [8, 4], broad surveys [2], and special issues of journals, e.g. *AI Magazine* [3] and *KI Zeitschrift* [17], on ASP. Furthermore, numerous talks[1] and tutorials on ASP, e.g. [5, 10, 12, 6] are available.[2] In this talk, we review possible reasons why ASP has gained popularity (if at all), and which factors have supported this and which are still impediments. We then consider the role of applications and implementations for the development of ASP, starting out with a look at the history of ASP. We then present some lessons learned, among them

- that theory use might differ from the original intention;

[1] Historical reflections on ASP, relevant here: http://www.cs.uky.edu/~mirek/stuff/kr-2018-gm.pdf
[2] See also the Potsdam ASP course: http://potassco.sourceforge.net/teaching.html

- that implementations and systems are a chicken-egg problem, yet even draft and imperfect implementations are vital to push research;
- that theory and applications are the two legs on which any area in the computing sciences has to stand upon;
- and that community building is essential.

We then discuss whether ASP could be a role model for NMR, where we critically review the current recognition of NMR.

Finally, we consider possible usages of ASP for NMR in the future, and we also address the issue of the role of NMR for ASP in its future development.

References

1 Baral, C.: Knowledge Representation, Reasoning and Declarative Problem Solving. Cambridge University Press (2002)
2 Brewka, G., Eiter, T., Truszczyński, M.: Answer set programming at a glance. Communications of the ACM 54(12), 92–103 (2011), http://dx.doi.org/10.1145/2043174.2043195
3 Brewka, G., Eiter, T., Truszczynski, M.: Answer set programming: An introduction to the special issue. AI Magazine 37(3), 5–6 (2016), http://www.aaai.org/ojs/index.php/aimagazine/article/view/2669
4 Dovier, A., Formisano, A., Pontelli, E.: Parallel answer set programming. In: Hamadi, Y., Sais, L. (eds.) Handbook of Parallel Constraint Reasoning., pp. 237–282. Springer (2018), https://doi.org/10.1007/978-3-319-63516-3_7
5 Eiter, T., Ianni, G., Krennwallner, T.: Answer set programming: A primer. In: Tessaris, S., Franconi, E., Eiter, T., Gutiérrez, C., Handschuh, S., Rousset, M., Schmidt, R.A. (eds.) Reasoning Web. Semantic Technologies for Information Systems, 5th International Summer School 2009, Brixen-Bressanone, Italy, August 30 – September 4, 2009, Tutorial Lectures. LNCS 5689, pp. 40–110. Springer (2009), https://doi.org/10.1007/978-3-642-03754-2_2
6 Eiter, T., Kaminski, T., Redl, C., Schüller, P., Weinzierl, A.: Answer set programming with external source access. In: Ianni et al. [11], pp. 204–275, https://doi.org/10.1007/978-3-319-61033-7_7
7 Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: Answer Set Solving in Practice. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers (2012), http://dx.doi.org/10.2200/S00457ED1V01Y201211AIM019
8 Gelfond, M.: Answer sets. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) Handbook of Knowledge Representation, chap. 7, pp. 285–316. Elsevier (2007)
9 Gelfond, M., Kahl, Y. (eds.): Knowledge Representation, Reasoning, and the Design of Intelligent Agents: The Answer-Set Programming Approach. Cambridge University Press (2014), http://dx.doi.org/10.1007/978-0-387-30164-8
10 Hölldobler, S., Schweizer, L.: Answer set programming and CLASP – A tutorial. In: Hölldobler, S., Malikov, A., Wernhard, C. (eds.) Proc. Young Scientists' International Workshop on Trends in Information Processing (YSIP). CEUR Workshop Proc. 1145, pp. 77–95. CEUR-WS.org (2014), http://ceur-ws.org/Vol-1145/tutorial1.pdf
11 Ianni, G., Lembo, D., Bertossi, L.E., Faber, W., Glimm, B., Gottlob, G., Staab, S. (eds.): Reasoning Web. Semantic Interoperability on the Web – 13th International Summer School 2017, London, UK, July 7-11, 2017, Tutorial Lectures, LNCS 10370. Springer (2017), https://doi.org/10.1007/978-3-319-61033-7
12 Kaminski, R., Schaub, T., Wanko, P.: A tutorial on hybrid answer set solving with clingo. In: Ianni et al. [11], pp. 167–203, https://doi.org/10.1007/978-3-319-61033-7_6
13 Lifschitz, V.: Answer set planning. In: ICLP. pp. 23–37 (1999)
14 Lifschitz, V.: Answer Set Programming and Plan Generation. Artificial Intelligence 138, 39–54 (2002)

15 Marek, V. W., Truszczyński, M.: Stable Models and an Alternative Logic Programming Paradigm. In: Apt, K., Marek, V. W., Truszczyński, M., Warren, D.S. (eds.) The Logic Programming Paradigm – A 25-Year Perspective, pp. 375–398. Springer (1999)
16 Niemelä, I.: Logic Programming with Stable Model Semantics as Constraint Programming Paradigm. Annals of Mathematics and Artificial Intelligence 25(3–4), 241–273 (1999)
17 Schaub, T., Woltran, S.: Special issue on answer set programming. KI 32(2-3), 101–103 (2018), https://doi.org/10.1007/s13218-018-0554-8

3.4 Reasoning about Exceptions in DL Ontologies

Laura Giordano (University of Eastern Piedmont – Alessandria, IT)

The study of nonmonotonic extensions of Description Logics is motivated by a problem in standard ontology languages (and, specifically, in OWL [1]) where a class inherits the properties of its superclasses, and where the treatment of exceptions is required in many application domains, from those concerning laws and regulations (where new laws override old ones) to medical ontologies.

The presentation is a short survey of the main approaches to reasoning about exceptions in Description Logics (DLs). Many non-monotonic extensions of DLs have been developed incorporating non-monotonic features from most of the non-monotonic formalisms in the literature, from default and autoepistemic logics, to circumscription and preferential logics, including also the approaches based on Answer Set Programming and, in general, on rule languages.

The landscape is very rich and the complexity of the various approaches has been studied, both for low complexity description logics and for high complexity description logics. The case of Description logic is an interesting case study for non-monotonic reasoning, which encompasses a limited treatment of non-monotonic reasoning in first order logic, namely, the treatment of the decidable fragment including only unary and binary predicates.

For highly expressive description logics, tractable constructions are especially important. Among the approaches, the rational closure has a polynomial construction, and it has been adapted to DLs [3], but it suffers from the well known problem called by Pearl [4] “the blocking of property inheritance problem". Refinements of his construction, avoiding this problem, have been presented in the literature, the prominent one being the lexicographic closure. The presentation compares these approaches, through examples, with other proposals, such as the logic of overriding [2], which has been defined in the setting of DLs, pointing at the achievements and open issues.

References

1 P. F. Patel-Schneider, P. H. Hayes, and I. Horrocks. OWL Web Ontology Language; Semantics and Abstract Syntax. In http://www.w3.org/TR/owl-semantics/, 2002.
2 P. A. Bonatti, M. Faella, I. Petrova, and L. Sauro. A new semantics for overriding in description logics. *Artif. Intell.*, 222:1–48, 2015.
3 G. Casini and U. Straccia. Rational Closure for Defeasible Description Logics. In T. Janhunen and I. Niemelä, editors, *Proceedings of the 12th European Conference on Logics in Artificial Intelligence (JELIA 2010)*, volume 6341 of *Lecture Notes in Artificial Intelligence*, pages 77–90, Helsinki, Finland, September 2010. Springer.

4 J. Pearl. System Z: A natural ordering of defaults with tractable applications to non-monotonic reasoning. In R. Parikh, editor, *TARK (3rd Conference on Theoretical Aspects of Reasoning about Knowledge)*, pages 121–135, Pacific Grove, CA, USA, 1990. Morgan Kaufmann.

3.5 Logics and Human Reasoning: What do we Know?

Marco Ragni (Universität Freiburg, DE)

Classical monotonic logics such as propositional or first-order logic have been considered the normative framework in psychology and cognitive science for evaluating human inferences. Core psychological results like the Wason Selection Task or the Suppression Task demonstrate that human common sense reasoning systematically deviates from classical logical reasoning. This has recently lead to a shift towards probabilistic or heuristic modeling approaches.

From a computational perspective, we can observe that recent cognitive theories have neither been formalized nor systematically analyzed or optimized. From this starting point, I will first introduce core findings and theories from cognitive science. Second, by applying methods from the AI field knowledge representation and reasoning and mathematical psychology I will analyze existing cognitive theories. As an example, I will present a reanalysis of the most prominent cognitive theories for syllogistic reasoning and show that any existing cognitive theory (including the probabilistic and heuristic theories) significantly deviates from the empirical data from psychological experiments. Hence, cognitive science still needs better theories in the sense of a better fitting (and predicting) empirical data. As a consequence, I will demonstrate that methods based on AI approaches can contribute to develop better cognitive theories. A discussion of the important role of non-monotonic logics concludes the presentation.

3.6 Argumentation Systems: A Brief Glimpse

Guillermo R. Simari (National University of the South – Bahía Blanca, AR)

Giving a knowledge repository, a natural way of finding a secure footing for the conclusions that could be obtained from it is arguing; particularly, when conflicting outcomes are reached arguing about which is to be supported is a rational way of handling such dispute. The process of arguing, and the very nature of an argument, have been the point of in-depth analysis in Philosophy since ancient times (see for instance [8, 12]) or more recently by proposing concrete models for arguments [16, 17, 14, 18]. Furthermore, the very nature of the discipline of Logic comes from the effort to clarify the presentation and exchange of arguments [15, 9]. Recently, in the field of Artificial Intelligence where many disciplines participate in the task of elucidating the essence of reasoning, research on computational argumentation has been expanding, giving birth to a field that is both exciting and fecund. In this presentation, we provide a glance at some of the ideas that are important in the field. We start with a description of the mechanics of argumentation as a procedure that could be regarded as a confrontation of postures regarding a claim. This activity can take the form, in a dialogical view, of an exchange of arguments between two participants where an arbiter decides over the initial claim considered in the debate, or in a monological description, as an internal debate where an agent performs all the roles itself, including that of the arbiter.

Describing an argumentation system requires the specification of different parts [11]. The starting point is to introduce a repository of beliefs that are represented in a formal language constituting the initial component of the system. Then, it becomes necessary to provide explicit rules for the construction from the belief base of an argument for a claim. An inference mechanism associated with the belief base usually will provide the reasoning that links the claim and the beliefs that act as premises. The next three stages address the problem of how arguments interact, formally defining the conflict between arguments, the comparison criteria between arguments, and how defeat is decided between conflicting arguments.

As the extensive range of topics explored in the literature precludes a full exploration of today's computational argumentation, we will limit ourselves to provide a concise foundation for further discussion, giving a short presentation of abstract argumentation frameworks [7] and structured argumentation systems [4]. For the interested reader, several references to general works are provided below. Finally, we will succinctly describe the international meetings, and some of the initiatives carried out in the field.

References

1 Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Christopher Reed, Guillermo R. Simari, Matthias Thimm, and Serena Villata. Towards artificial argumentation. *AI Magazine*, 38(3):25–36, 2017.
2 Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.
3 Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors. *Handbook of Formal Argumentation.* College Publications, 2018.
4 Philippe Besnard, Alejandro Javier García, Anthony Hunter, Sanjay Modgil, Henry Prakken, Guillermo Ricardo Simari, and Francesca Toni. Introduction to structured argumentation. *Argument & Computation*, 5(1):1–4, 2014.
5 Philippe Besnard and Anthony Hunter. Argumentation based on classical logic. In Iyad Rahwan and Guillermo Ricardo Simari, editors, *Argumentation in Artificial Intelligence*, pages 133–152. Springer, 2009.
6 Carlos I. Chesñevar, Ana G. Maguitman, and Ronald P. Loui. Logical models of argument. *ACM Computing Surveys*, 32(4):337–383, 2000.
7 Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.
8 Charles Griswold. Plato on rhetoric and poetry. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy.* Fall 2009 edition, 2009.
9 Dov M. Gabbay and John Woods, editors. *Greek, Indian and Arabic Logic*, volume 1 of *Handbook of the History of Logic.* Elsevier, The Netherlands, 2004.
10 Sanjay Modgil, Francesca Toni, Floris Bex, Ivan Bratko, Carlos I. Chesñevar, Wolfgang Dvořák, Marcelo A. Falappa, Xiuyi Fan, Sarah A. Gaggl, Alejandro J García, et al. The added value of argumentation. In *Agreement Technologies*, pages 357–403. Springer, 2013.
11 Henry Prakken and Gerard Vreeswijk. Logics for defeasible argumentation. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 4, pages 218–319. Kluwer Academic Pub., 2002.

12 Christof Rapp. Aristotle's rhetoric. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2010 edition, 2010.
13 Iyad Rahwan and Guillermo Ricardo Simari. *Argumentation in Artificial Intelligence*. Springer, 2009.
14 Michael Scriven. *Reasoning*. McGraw-Hill Book Company, 1976.
15 Stewart Shapiro and Teresa Kouri Kissel. Classical logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2018 edition, 2018.
16 Stephen E. Toulmin. *The Uses of Argument*. Cambridge University Press, 1958.
17 Stephen E. Toulmin. *The Uses of Argument. Updated Edition*. Cambridge University Press, 2003.
18 F. H. van Eemeren and R. Grootendorst. *Argumentation, communication, and fallacies: A pragma-dialectical perspective*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1992.

3.7 Argument Mining: from Non-Monotonic Reasoning to Natural Language Processing and Back

Serena Villata (Laboratoire I3S – Sophia Antipolis, FR)

Argument(ation) Mining (AM) [1, 2] is a recent research area in Artificial Intelligence (AI), mainly across the standard areas of Knowledge Representation and Reasoning (KRR) on the one side, and Natural Language Processing (NLP), on the other side. Few approaches to what is now called argument mining started to appear around 2010, when the first methods to mine (different connotations of) *arguments* from natural language documents were proposed: [4] introduced the definition of argumentative zoning for scientific articles, and [3] proposed a way to detect arguments from legal texts. Since these seminal approaches, the need for automated methods to mine arguments and the relations among them from natural language text was brought to light, but it was only briefly touched upon. The parallel advances, from the formal point of view in the research field of computational models of argument, and from the point of view of the computational techniques for learning and understanding human language content in the NLP and the Machine Learning fields, boosted the almost contemporary organization of two events in 2014 targeting open discussions about the challenge of mining arguments from text.[3]. Since then, AM has became a topic in major AI and NLP conferences.

Argument mining involves several research areas from the AI panorama: NLP provides the methods to process natural language text, to identify the arguments and their components (i.e., premises and claims) in texts and to predict the relations among such arguments, KRR contributes with the reasoning capabilities upon the retrieved arguments and relations so that, for instance, fallacies and inconsistencies can be automatically identified in such texts, and Human-Computer Interaction guides the design of good human-computer digital argument-based supportive tools. The argument mining pipeline is composed of two main steps: first, the argument components are identified in the text (i.e., premises and claims),

[3] The workshop on Argument Mining (https://goo.gl/kF4Eep) co-located with ACL, and the workshop on Frontiers and Connections between Argumentation Theory and Natural Language Processing (https://goo.gl/ttVUZk)

and second, the relations between arguments (e.g., support and attack) are predicted. Usually supervised learning methods are used to face these tasks, leading to the need of defining beforehand annotated datasets for the specific task and application scenario.

In addition, AM is strongly connected with hot topics in AI, as deep learning (heavily used in AM), fact checking (the prediction of the attacks between arguments is a building block for fake news detection), explanations of machine decisions (AM can disclose how the information on which the machine relies to make its own decisions is retrieved), medicine (where AM can detect information needed to reason upon randomised clinical trials), politics (where AM can provide the means to automatically identify fallacies and unfair propaganda), and for cyberbullism prevention (where AM can support the detection of repeated attacks against a person).[4]

References

1. Elena Cabrio and Serena Villata. Five years of argument mining: a data-driven analysis. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, pages 5427–5433. ijcai.org, 2018.
2. Marco Lippi and Paolo Torroni. Argumentation mining: State of the art and emerging trends. *ACM Trans. Internet Techn.*, 16(2):10, 2016.
3. Rachele Mochales and Marie-Francine Moens. Argumentation mining. *Artificial Intelligence and Law*, 19(1):1–22, 2011.
4. Simone Teufel, Advaith Siddharthan, and Colin Batchelor. Towards domain-independent argumentative zoning: Evidence from chemistry and computational linguistics. In *EMNLP*, pages 1493–1502, 2009.

4 Working groups

4.1 Working Group on Commonsense Reasoning

Anthony Hunter (University College London, GB)

Commonsense is an innate ability of humans, and appears to be critically important for us to operate individually and collectively in the world. Wikipedia defines commonsense as, “the basic ability to perceive, understand, and judge things that are shared by (‘common to’) nearly all people and can reasonably be expected of nearly all people without need for debate”. To help us understand the ubiquity and complexity of commonsense reasoning, we can consider three spheres of human activity where commonsense reasoning is important: Physics – General understanding of how the physical world works (naive physics), e.g. explaining why a vase breaks when dropped on the floor; Psychology – Basic understanding of human motives and behaviors (i.e., a theory of mind), e.g. explaining emotions of a colleague who is unhappy when they get a paper rejected; Society – Basic knowledge about how people can operate in societies, e.g. process to pay for a meal, how to get a credit card, etc. If we are to develop more capable and robust artificial intelligence, then we need to incorporate commonsense reasoning. Current challenging problems in artificial intelligence research such

[4] http://creep-project.eu

as explainability in machine learning, and deeper semantic understanding in natural language processing, are just two examples where commonsense reasoning is required. Commonsense reasoning has been a key driver of research in nonmonotonic reasoning, and going forward, it will be valuable to further develop and apply this research in leading edge problems in areas such as machine learning, machine vision, planning, and natural language understanding.

4.2 Working Group on Integrating ML and KR, and the Relevance of Prototypical Reasoning

Gabriele Kern-Isberner (TU Dortmund, DE)

Participants: Zied Bouraoui, Giovanni Casini, Célia da Costa Pereira, Luc de Raedt, Eduardo Fermé, Gabriele Kern-Isberner, Ken Satoh, Serena Villata

The discussions in this working group focussed on two issues:

- How can the fields of argument mining, NLP (natural language processing), machine learning (ML), and NMR can benefit from one another?
- Links between ML and NMR (or knowledge representation and reasoning (KR) in general) have long been discussed, but obviously, we are still far from successful overall integration. What could be the next steps to promote collaborations? What challenges are interesting and rewarding for both areas?

First, we present plans for combining argument mining, NLP, and NMR, and then set up a roadmap for promoting the integration of ML and NMR.

4.2.1 Plans for Combining Argument Mining, NLP, and NMR

Argument mining extracts arguments from text the basic building blocks of which are often rules. Besides feeding these arguments into argumentative systems, the extracted rules could also be used to build up knowledge bases for NMR systems. NMR inferences can then be compared to inferences obtained via argumentation. Actually, knowledge bases are of crucial importance for real-world applications of NMR/KR, therefore argument mining groups could expand its scope to collaborate with NMR groups, and NMR groups (in particular, those focussing on inductive reasoning from knowledge bases) could apply their approaches to larger examples. Benchmarks may arise from this. On the more theoretical side, more material for elaborating on the differences/connections between argumentation and NMR would be available.

On the argument mining/NLP side, textual inference is an important topic. It seems obvious that NMR methods can be useful to formalize and compute such inferences. However, first the principles underlying textual inferences have to be investigated and better understood on the NMR side.

NMR and NLP need a better interface between them. A joint Dagstuhl-like workshop would be a good starting point for this. The plan is to set up an application for a Dagstuhl workshop, or to apply somewhere else (e.g., Madeira, Cape Town, Bertinoro, Leibniz-Center in NL).

4.2.2 ML and NMR – a Possible Roadmap

The working group discussed crucial challenges and innovative views that are interesting and beneficial both for ML and NMR, and apt to promote integrative approaches:

- Given that KR needs knowledge (bases) for real-world applications (but also for good toy and benchmark examples), extracting (qualitative) knowledge bases from data seems to be a major challenge. Inductive logic programming (ILP) and its extensions to answer set programming is a promising approach, but more algorithms to mine default rules from data that can be used in other NMR approaches are urgently needed. Only few approaches to that exist to date.
- ML algorithms often mine too much knowledge from data so that users are swamped by (e.g.) too many rules (when mining association rules from data). Here, formal NMR methods can help to compactify the set of mined rules so that a useful knowledge base results in the end.
- Explainable AI is broadly discussed currently. Can ML algorithms be explained in terms of NMR? Can we enhance ML methods in general by integrating (NM) reasoning? The other way round, is NMR a suitable context for ML? Which kind of reasoning/NMR approach should we use for that? Will reasoning improve the interpretability of ML results beyond explainability? How can strict rules be used to increase performance of ML algorithms? How can domain and contextual knowledge help interpreting the prediction of an ML model? Benchmark examples for comparing pure ML techniques and KR-enhanced techniques would be very useful here.
- ML usually makes use of statistical measures to validate their results and base decisions upon. Can the concept of plausibility help us to find better decisions?
- Each ML approach needs an inductive bias, how can this be related to NMR inductive reasoning?
- The role of prototypes (stereotypes) and their dynamics should be better explored in NMR and belief revision (BR), and relevance of prototypical reasoning for ML should be investigated.

A high-level vision on the connection between ML and NMR/KR is sketched in the picture below.

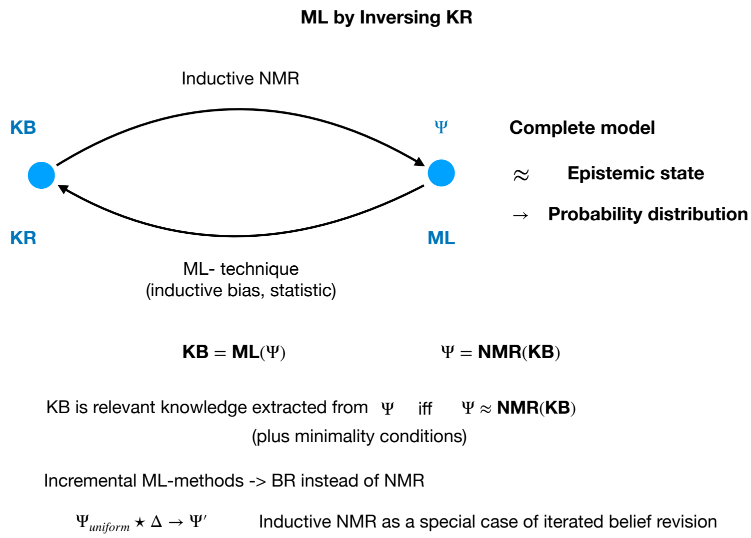

4.3 Working Group on Implications for NMR and Cognition

Thomas Meyer (University of Cape Town, ZA)

Participants on the first day of discussion: Arina Britz, Thomas Meyer, Abhaya Nayak, Marco Ragni, Hans Rott.

The number of participants went up significantly on the second day of the working group discussion, indicating substantial interest in this topic.

4.3.1 Topics of Discussion

In the working group we focused on two sub-topics:
1. Implications of NMR for Philosophy of Mind/Philosophy of Science (and vice versa)
2. Implications of NMR for cognition (and vice versa)

- **Philosophy of Mind:** The consensus was that this falls under Cognition, and should therefore be included under the discussions related to Cognition.
- **Philosophy of Science:** Current approaches in this field are mostly Bayesian in nature. There is a need to emphasise the advantages of the qualitative approaches in NMR. Any work on establishing links with NMR at present will depend heavily on Philosophers of Science recognising the need for this. At present it is not clear that there is such a need from this community.
- **Cognition:** The situation is different when it comes to links between NMR and Cognition. In this case there is clear interest in investigating the connections between the two areas from both communities. The primary benefits for the Cognition community relate to the need for formal systems to provide clarity. For the NMR community, the main benefits relate to the availability of a suite of benchmarks against which to test formal theories, as well as the potential for running experiments: a methodology that is well understood in the Cognition community, but not often employed in the NMR community.

4.3.2 Discussion Points

Because of the interest in Cognition, the more detailed points of discussion centred exclusively on this sub-topic.
- **Descriptive vs. Normative:** There is general agreement that it is important to investigate and model both descriptive and normative aspects of NMR with respect to Cognition. That is, we should investigate how people ought to reason in a non-monotonic context, as well as how they actually reason in such contexts. Building on that, there is interest in teasing out the difference between errors in reasoning on the one hand, and (sound) reasoning patterns that simply differ from what is expected, on the other hand.
- **Common Forms of Mistakes:** The goal here is to focus on systems that reason correctly (plausibly), that can point out mistakes by humans and correct them, and that can predict and explain mistakes made by humans. The question arises whether these processes can be automated. One possible avenue of investigation is whether this type of automation could benefit from machine learning.
- **Benchmark problems:** There is a clear need to update the 1990 list of NMR benchmark problems set up by Vladimir Lifschitz. It was also noted that, while the existing benchmark

problems seem to be well-designed for concrete instances of rules and reasoning patterns, they are not suitable for more abstract ones, such as those encountered in deontic reasoning, for example.

- **Explainability vs Rationality:** An important discussion point that was raised is the one of explainability vs. rationality. One of the driving forces, and indeed, one of the strongest selling points of NMR, is to design and build systems exhibiting rational behaviour that can be explicitly explained. But it needs to be borne in mind that explainability doesn't necessarily imply rationality. A simple example of this would be laws of nature, which are explainable, but not rational.

4.3.3 Next Steps

- **Follow-up seminar within a year:** There needs to be a follow-up seminar within a year to keep the momentum going. The focus of the follow-up seminar should be the link between NMR and Cognition. An important component of this seminar should be a series of tutorials (on NMR and Cognition) aimed at educating members of the two communities about the state of the art in the other community. Next steps after the follow-up seminar could be an event at NMR 2020 (in Patras, Greece) and a Dagstuhl Seminar in the longer run.
- **Collaborative Site:** There is an urgent need for a collaborative site to be established for this sub-topic. The immediate need is for a repository and a collaborative platform. In the short term it was proposed to use the Dagstuhl Wiki for this. In the longer term it should be hosted on an NMR site, which is to be established. The establishment of the NMR site is a point that should be taken up with the new chairs of the NMR Steering Committee, Gabriele Kern-Isberner and Renata Wassermann.

4.4 Working Group on NMR Engineering

Renata Wassermann (University of Sao Paulo, BR)

While there are still many interesting theoretical issues to be solved in the realm of NMR, we feel that experimenting with real applications and having real world problems in mind can benefit the area. In this working group, the participants discussed possible strategies for increasing the visibility of existing implementations and applications, as well as the need to encourage new contributions to this "engineering" part of NMR.

A first challenge reported by those who are involved in implementations is to gather reliable benchmarks for the different branches of NMR. For some sub-areas of NMR, there are system competitions that take place regularly and could serve as an inspiration for devising benchmarks for the other areas. For example, in Answer Set Programming (ASP), the ASP Challenge (https://sites.google.com/view/aspcomp2019) happens every two years since 2007. The International Conference on Computational Models of Argument (COMMA – http://comma.csc.liv.ac.uk), features a track on system demonstrations. During the meeting, we have listed a series of other sources for inspiration:

- The Competition on Legal Information Extraction and Entailment[5] (COLIEE), which provides examples for legal reasoning.
- The International General Game Playing Competition[6] (IGGPC).
- The SAT competition[7]
- Examples from the International Workshop on Principles of Diagnosis Series[8] (DX).
- The SATLIB[9] and TPTP[10] repositories provide propositional knowledge bases for SAT problems.
- Several biomedical ontologies and knowledge bases that can be found on the web, for example, in the BioPortal[11].

It is important to remark that in order to have good benchmarks, it is not enough to collect knowledge bases, we have to generate interesting problems using them. For example, in belief revision, we must also select input sentences which give rise to interesting contraction or revision problems.

Another topic of discussion was the visibility of existing tools and applications, both for researchers within NMR, that could benefit from what colleagues are doing as for the Artificial Intelligence community, that still sees NMR as a purely theoretical area.

As short term actions that were discussed in the meeting, we can mention:

- Create a web page and a repository, possibly linked to the page of the NMR Workshop series.
- Collect for the repository existing implementations and tools, with their description.
- Collect also examples of “real world” applications.
- Collect tutorials, overview papers, slides on various aspects of NMR.
- Edit the Wikipedia pages on NMR and related topics:
  - https://en.wikipedia.org/wiki/Non-monotonic_logic
  - https://en.wikipedia.org/wiki/Defeasible_reasoning
  - https://en.wikipedia.org/wiki/Belief_revision
- Create or link repository for Knowledge Bases in propositional logic and Description Logics.
- Define guidelines for creating reasoning problems from the knowledge bases.
- Define simple forms for researchers to submit new tools or benchmarks sets to the repository.

[5] https://sites.ualberta.ca/~miyoung2/COLIEE2018
[6] http://ggp.stanford.edu/iggpc/index.php
[7] http://www.satcompetition.org
[8] https://dx-workshop.org/dx-series
[9] https://www.cs.ubc.ca/~hoos/SATLIB/benchm.html
[10] http://www.tptp.org
[11] https://bioportal.bioontology.org

Participants

- Zied Bouraoui
Artois University – Lens, FR
- Arina Britz
University of Stellenbosch, ZA
- Giovanni Casini
University of Luxembourg, LU
- Célia da Costa Pereira
Laboratoire I3S – Sophia Antipolis, FR
- Luc De Raedt
KU Leuven, BE
- James P. Delgrande
Simon Fraser University – Burnaby, CA
- Thomas Eiter
TU Wien, AT
- Eduardo Fermé
University of Madeira – Funchal, PT
- Laura Giordano
University of Eastern Piedmont – Alessandria, IT
- Andreas Herzig
University of Toulouse, FR
- Anthony Hunter
University College London, GB
- Gabriele Kern-Isberner
TU Dortmund, DE
- Sebastien Konieczny
Artois University – Lens, FR
- Maria Vanina Martinez
University of Buenos Aires, AR
- Thomas Meyer
University of Cape Town, ZA
- Abhaya Nayak
Macquarie University – Sydney, AU
- Odile Papini
University of Marseille, FR
- Marco Ragni
Universität Freiburg, DE
- Gavin Rens
KU Leuven, BE
- Hans Rott
Universität Regensburg, DE
- Ken Satoh
National Institute of Informatics – Tokyo, JP
- Guillermo R. Simari
National University of the South – Bahía Blanca, AR
- Ivan José Varzinczak
Artois University – Lens, FR
- Serena Villata
Laboratoire I3S – Sophia Antipolis, FR
- Renata Wassermann
University of Sao Paulo, BR
- Emil Weydert
University of Luxembourg, LU