Report from Dagstuhl Seminar 11351

Computer Science & Problem Solving: New Foundations

Edited by

Iris van Rooij[1], Yll Haxhimusa[2], Zygmunt Pizlo[3], and Georg Gottlob[4]

1 **Radboud University Nijmegen, NL,** i.vanrooij@donders.ru.nl
2 **Vienna University of Technology, AT,** yll@prip.tuwien.ac.at
3 **Purdue University – West Lafayette, US,** pizlo@psych.purdue.edu
4 **University of Oxford, GB,** Georg.Gottlob@comlab.ox.ac.uk

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11351 "Computer Science & Problem Solving: New Foundations". This seminar was the first Dagstuhl seminar that brought together a balanced group of computer scientists and psychologists to exchange perspectives on problem solving. In the 1950s the seminal work of Allen Newell and Herbert Simon laid the theoretical foundations for problem solving research as we know it today, but the field had since become disconnected from contemporary computer science. The aim of this seminar was to promote theoretical progress in problem solving research by renewing the connection between psychology and computer science in this area.

Seminar 29. July–02. August, 2011 – www.dagstuhl.de/11351
1998 ACM Subject Classification F.1.1 Models of Computation, F.1.3 Complexity Measures and Classes, F.2.0 [Analysis of Algorithms and problem complexity] General, I.2.4 Knowledge Representation Formalisms and Methods, I.2.8 Problem Solving, Control Methods, and Search.
Keywords and phrases Problem solving, Cognitive psychology, Cognitive systems, Vision Representations, Computational complexity
Digital Object Identifier 10.4230/DagRep.1.8.96

1 Executive Summary

Iris van Rooij
Yll Haxhimusa
Zygmunt Pizlo
Georg Gottlob

This Dagstuhl seminar brought together a group of computer scientists and psychologists to discuss their perspectives on problem solving. The seminar was inspired by two previous Problem Solving workshops in 2005 and 2008 at Purdue University, USA. These workshops, organized primarily by psychologists, laid bare some fundamental theoretical questions in problem solving research. The organizers believed that research on these questions could benefit from more involvement of computer scientists in the area of problem solving. This motivated the organization on this seminar, which aimed to bring together computer scientists and psychologists to help build new formal foundations for problem solving research.

Of the 36 participants at the seminar about half were computer scientists and the other half were psychologists, though many identified as interdisciplinary researchers (e.g., cognitive scientists). To facilitate cross-disciplinary perspectives, computer science and psychology talks were alternated in the program of the seminar. There were 7 longer featured talks and 15 shorter talks.

On Day 1 of the seminar, Iris van Rooij opened the seminar by explaining its history and motivation. She discussed how computational complexity theory gives a formal framework for quantifying the difficulty of solving different types of search problems (e.g., the Traveling Salesman Problem and Minimum Spanning Tree), but that no analogous formal framework exists yet for quantifying the difficulty of solving so-called 'insight' problems (e.g, the Nine-dot problem), or more generally, quantifying the difficulty of representing a problem in the right way. The question of how one could develop such a formal framework was an overarching theme of the seminar. All participants were invited to think about this question. The topic resurfaced in several talks and workgroup discussions.

The featured talks on Day 1 were by Todd Wareham and Bill Batchelder. Wareham presented novel ideas on how a formal theory of 'insight difficulty' may take shape. His analysis was based on existing ideas in the psychology literature, such as the Representational Change Theory of Knoblich and et al., and his formalisms were inspired by Gentner's Structure-Mapping Theory. Batchelder presented a list of 19 classic examples of insight problems (these problems can be found in Appendix 7.1. These problems served as illustrations of problems that are not 'search problems' in the sense of Newell and Simon, yet for which problem solving researchers should nevertheless like to be able to model and explain the processes involved. Four shorter talks on Day 1 were given by Georg Gottlob, Sarah Carruthers, Sashank Varma and Jakub Szymanik. Gottlob introduced conceptual tools from computational complexity theory, graph theory and probabilistic computation that could inspire new ways of thinking about problem solving. Carruthers presented novel experimental data on how humans solve the graph problem VERTEX COVER (see Appendix 7.3 for a definition). Varma presented a methodology for modeling the resource requirements of different brain areas invoked during problem solving. Szymanik presented a generalization of the Muddy Children Problem and explained how its solution can be modeled using logic.

On Day 2 of the seminar there were 3 featured talks. In the first featured talk, Niels Taatgen presented the ACT-R modeling architecture and illustrated how it could model the development of more general problem solving skills as a re-combination of more basic skills. Rina Dechter and Ken Forbus each gave a different AI perspective on problem solving in their featured talks. Dechter presented several sophisticated algorithmic techniques for solving NP-hard problems, such as Constraint Satisfaction and Bayesian Inference. Forbus proposed to consider 'analogy' as a new foundation for problem solving research and illustrated his perspective using the Companion framework. There were 3 short talks on this day. The two short talks by Johan Kwisthout and Marco Ragni (like Wareham's talk on Day 1) touched clearly on the theme of the seminar. Kwisthout proposed a formal framework for capturing the notion of 'relevance' when it comes to finding a suitable problem representation, and Ragni proposed a framework for quantifying the *a priori* difficulty of problem items on an IQ test based the notion of 'representational transformation'. In the third short talk, Jelle van Dijk gave a designers' perspective on problem solving. Van Dijk made the case that much real-world problem solving is probably best studied from an embodied embedded cognitive perspective. The official program for this day was closed with a working group discussion on meanings of common terms used throughout the talks (see Section 4.1).

Day 3 opened with a featured talk by Dedre Gentner. The talk by Gentner complemented the talk by Forbus on Day 2 as she laid out the experimental evidence for the idea that analogical thinking (comparison and matching) lies at the foundation of human learning and reasoning. The featured talk was followed by two short talks, one by Liane Gabora and one by Daniel Reichman. Gabora presented a perspective on problem solving that is quite unlike the traditional view of problem solving as search through a well-defined space for a well-defined solution. Her perspective is that (creative) problem solving can perhaps best be seen as the recognition and actualization of a solution that before only existed in a state of potentiality. Reichman presented a theoretical computer science perspective on the well-known phenomenon of speed-accuracy tradeoffs in psychology. He proposed that algorithmic techniques from computer science can help predict what shape curves describing speed-accuracy tradeoffs will have in a variety of experimental conditions. In the afternoon of Day 3 there was no official program, and instead participants enjoyed the surroundings of Schloss Dagstuhl and/or went for a hike on one of the hills near the Schloss.

Day 4 started with the featured talk by Yun Chu, who gave an overview of the psychological research on insight problem solving (the talk had originally been scheduled for Day 1, but due to unforeseen circumstances Chu could not arrive at the seminar earlier). The talk by Chu helped build further common ground between the computer scientists and psychologists as it explained in more detail common paradigms and concepts used in psychological research on problem solving. The rest of the day consisted of two workshop sessions aimed at stimulating the formation of new interdisciplinary perspectives and collaborations (for details see Sections 4.2 and 4.3) and several short talks. Ute Schmid presented a framework for what Chomsky called a 'competence level model' of learning to problem solve, based on analytical inductive functional programming. Ulrike Stege gave a survey of typical computer science problems that are or could be used to investigate human problem solving strategies and pointed out some research challenges. Among them is the problem that researchers may think their participants are solving the problem that they posed, but the participants may in fact be solving an altogether different problem which the participants *think* the researchers have posed. Brendan Juba presented a new formal framework for heuristic rules based on PAC semantics. Nysret Musliu presented the concept of a (hyper)tree decomposition, a concept that can be utilized in algorithmic techniques for solving NP-hard problems efficiently. Jered Vroon presented a non-standard formalism in which problem solving is regarded as producing a solution rather than as a search through a search space. Last, Zyg Pizlo presented new algorithmic ideas for modeling human performance on the Traveling Salesman Problem based on the notion of multiresolution-multiscale pyramids. The day closed with a session in which participants brainstormed about novel interdisciplinary collaborations and open problems in the field. Some of these ideas were presented the same day, others were presented in the morning session of Day 5.

The morning of the last day of the seminar, Day 5, was reserved for short presentations of new collaborative ideas that the participants came up with, as well as the presentation of new ideas and open problems (see Sections 4.3 and 5 for details). The seminar closed with a wrap-up session in which participants reflected on the process and outcomes of the seminar (see Section 4.4 for a summary). To conclude, the seminar was successful in several ways: (1) It has resulted in a renewed awareness of how computer science and psychology can complement each other in the study of problem solving; (2) it has created a new impetus for more involvement of computer scientists in contemporary problem solving research; (3) it has created more common ground between computer science and psychologist in the study of problem solving; (4) it has produced several novel ideas on how to conceptualize 'problem solving' and, in particular, 'problem solving by insight;' (5) it has produced several novel ideas on how to formalize these new conceptualizations; (6) it has produced concrete suggestions for new experimental paradigms for studying problem solving in the lab; (7) it has inspired new cross-disciplinary collaborative research projects; and last but not least (8) it has provided the groundwork on which follow-up Dagstuhl seminars can build in the future. With this seminar, the organizers hope to have contributed to an increased and sustained collaborative research effort between computer science and psychology in the domain of problem solving.

2 Table of Contents

3 Overview of Talks

3.1 Some Issues in Developing a Theory of Human Problem Representation

William H. Batchelder (University of California, US)

First I will describe some of my personal experiences interacting with computer scientists concerning the games of Chess and Go. There once was a strong belief that human knowledge could aid computational approaches to such games as these; however, approaches based on that notion ultimately failed, and instead brute force appears to have won the day. Next I will discuss some of the barriers that I see as standing in the way of developing a satisfactory formal theory of human problem solving. They include: (1) The lack of general experimental paradigms to study problem solving as found, for example, in other areas of cognitive psychology such as human memory, human attention, or visual perception. (2) The lack of a rich behavioral base that accompanies the act of solving a problem. (3) The lack of a formal theory of how problem solvers initially represent and possibly re-represent problems during solution efforts. Finally, in the last half of the talk, I will discuss variations on a set of twelve or so problems drawn from the folklore of brain teasers. I selected these problems because they are not move problems in the sense of Newell and Simon, but instead they require creative problem representations. For each of these problems, once a good representation is achieved, the solution follows pretty easily. I will organize the problems and variations on them around aspects human cognition that must be utilized or overcome to achieve a productive problem representation. These aspects of cognition include the nature of the human senses, imagery, memory, cognitive biases, and reasoning processes. In particular, I will use the problems to draw out issues that may need to be handled in constructing a formal theory of human problem representation. I will file a list of the brain teasers for your possible interest the week before the workshop starts.

3.2 Vertex Cover and Human Problem Solving

Sarah Carruthers (Univ. of Victoria, CA)

In this seminar, we will look at preliminary results from a study of human solutions to Vertex Cover problems. The purpose of this pilot study is to identify: select strategies employed by participants, and features of instances which may impact performance. We will also discuss what measures of performance are of interest, as well as related problems which may be of interest.

3.3 Human Performance on Insight Problem Solving: A Review

Yun Chu (Univ. of Hawaii at Manoa, US)

A review of recent research on insight problem-solving performance. We discuss what insight problems are, the different types of classic and newer insight problems, and how we can classify them. We also explain some of the other aspects that affect insight performance, such as hints, analogs, training, thinking aloud, and individual differences. In addition, we describe some of the main theoretical explanations that have been offered. Finally, we present some measures of insight and relevant neuroscience contributions to the area over the last decade.

3.4 Advanced Reasoning in Graphical models

Rina Dechter (Univ. California – Irvine, US)

Graphical models, including constraint networks, belief networks, Markov random fields and influence diagrams, are knowledge representation schemes that capture independencies in the knowledge base and support efficient, graph-based algorithms for a variety of reasoning tasks, including scheduling, planning, diagnosis and situation assessment, design, and hardware and software verification. Algorithms for processing graphical models are of two primary types: inference-based and search-based. Inference-based algorithms (e.g., variable-elimination, join-tree clustering) are time and space exponentially bounded by the tree-width of the problem's graph. Search-based algorithms can be executed in linear space and often outperform their worst-case predictions. The thrust of advanced schemes is in combining inference and search yielding a spectrum of memory-sensitive algorithms universally applicable across graphical models. The talk will provide an overview of principles of reasoning with graphical models developed in the last decade in constraints and probabilistic reasoning.

3.5 Analogy as a Computational Foundation for Problem-solving and Learning

Kenneth D. Forbus (Northwestern University – Evanston, US)

Classical approaches to problem-solving focus on first-principles reasoning (e.g. logic), using quantified representations and chaining. Analogy and similarity, to the extent they are considered at all, are viewed as rare operations which can safely be ignored, at least to first order. This talk, which describes joint work with Dedre Gentner, argues that the opposite is true: That analogy and similarity should be viewed as primary reasoning operations, with logic and chaining being used in support of them. We start by using a series of examples (including visual problem solving, physics problem solving, counterterrorism, and moral decision-making) to introduce the primitive operations of analogical processing: Matching, retrieval, and generalization. The importance of qualitative representations, which facilitate analogical reasoning and learning, will be outlined. Finally, we describe some of the new issues that this framework raises for computational models of problem solving, including experience, learning encoding, rerepresentation, and skolem resolution.

3.6 Problem Solving as the Recognition and Actualization of Potentiality

Liane Gabora (University of British Columbia – Vancouver, CA)

Theories of creativity, such as Finke, Ward and Smith's Geneplore model and Simonton's Darwinian model, generally assume that creative problem solving involves searching through memory, selecting well-defined candidate ideas, and then tweaking them in response to problem constraints. I would like to discuss an alternative: that creative individuals wrestle with issues or ideas that are, for them, not well-defined, or in a state of potentiality. Over time, as these ideas are considered from different perspectives, they come to assume a form that is more fully actualized, or well-defined. This suggests that (in accordance with Einstein's assertion that finding the right question is a more significant step than answering it) the challenge is to construct one's mental model of reality in such a careful and precise way that one is led directly to the frontiers of what is known, and thus able to recognize and engage in informed speculation about what lies beyond these frontiers, i.e. what is currently in a state of potentiality for everyone. Problem solving is thus viewed as redefining the unknown in terms of what is known. However, the 'unknown' can take the form of not just gaps in knowledge but experiences that are so traumatic or unusual that we have not fully come to terms with them. This leads to a related question that I would like to explore: can artistic endeavors be understood within a problem-solving framework? I suggest that artistic tasks are those for which the topology of the fitness landscape is determined not by objective, agreed-upon aspects of the world, but by personal experiences and the emotions surrounding them. The problem-solving task for the artist is to translate information patterns underlying the dynamics of, for example, neurotransmitter release that rendered the particular emotional impact of an experience or situation into the constraints of the artistic form; for example, the tragic experience of losing a family member might be translated into the constraints of music. In so doing, one re-frames the experience in terms of what one has experienced before, and thus incorporates it into the fabric of one's understanding of the world, and comes to terms with it.

3.7 The Analogical Mind

Dedre Gentner (Northwestern University – Evanston, US)

Analogical processes are central in human learning and reasoning. Analogical comparison engages a process of structural alignment and mapping that fosters learning and reasoning in at least three distinct ways: it highlights common relational systems; it promotes inferences; and t calls attention to potentially important differences between situations. It can also lead to re-representing the situations in ways that reveal new facets.

An important outcome of analogical comparison is that the common relational structure becomes more salient and more available for transfer in short, a portable abstraction is formed. Thus, structure-mapping processes bootstrap much of human learning.

The power of analogy is amplified by language learning. Hearing a common label invites comparison between the referents, and this structure-mapping process yields insight into the meaning of the term. The mutual facilitation of analogical processing and relational language contributes to the power and flexibility of human learning.

Finally, although analogy is sometimes thought of as a clever, somewhat effortful process, in fact it is pervasive in human processing, In this talk, I present psychological studies showing the role of analogy processes in human learning and reasoning. I will try to convey not only the power and importance of analogical processes but also their ubiquity in human cognition.

3.8 Living with Computational Complexity

Georg Gottlob (University of Oxford, GB)

Many computational problems that are important in real life are intractable. There are different ways of coping with intractability. This talk will focus, in particular, on methods of recognizing large 'islands of tractability' for NP-hard problem, i.e., large tractable subclasses. We will illustrate the use of graph-theoretic concepts such as tree-width and hyper-tree width in order to obtain large polynomial classes of intractable problems. In addition, we will mention a number of other ways of coping with complexity and illustrate how both computer programs and fruit flies (rather than bees) can solve certain 'complex' problems. The talk will end with some thoughts of more philosophical nature about computational complexity.

3.9 PAC Semantics: A Framework for Heuristic Rules

Brendan Juba (MIT – Cambridge, US)

Valiant's PAC Semantics [1] provides a clean standard that captures the utility of 'rules of thumb' about a given domain that may, for example, be derived from a sample of typical experiences in the domain; we suggest that it may be useful for the analysis of the acquisition and use of heuristic rules. In support of this suggestion, we show that PAC Semantics also features some natural tractable cases for inference. We describe a simple and efficient algorithm that tests the validity of candidate assertions given access to 'partially obscured' samples from the domain, correctly classifying all assertions except for good rules of thumb that cannot be established by a 'simple proof' using typical obscured examples.

References

1 Leslie G. Valiant. Robust logics. *Artificial Intelligence*, 117:231–253, 2000.

3.10 Relevant Representations

Johan Kwisthout (Radboud University Nijmegen, NL)

When computer scientists discuss the computational complexity of, e.g., finding the shortest path from A to B, their starting point typically is a formal description of the problem at hand, e.g., a graph with weights on every edge.

Given such a formal description, either tractability or intractability of the problem is established, by proving that the problem enjoys a polynomial time algorithm, respectively is NP-hard. However, this problem description is in fact an abstraction of the actual problem of being in A and desiring to go to B: it focuses on the relevant aspects of the problem (e.g., distances between landmarks and crossings) and leaves out a lot of irrelevant details.

This abstraction step is often overlooked, but may well contribute to the overall complexity of solving the problem at hand. For example, it appears that 'going from A to B' is easier to abstract: it is fairly clear that the distance between A and the next crossing is relevant, and that the color of the roof of B is typically not. However, when the problem to be solved is 'make X love me', where the current state is (assumed to be) 'X does not love me', it is hard to agree on all the relevant aspects of this problem.

In this talk, I will propose a framework for capturing the notion of relevance when it comes to finding a suitable problem representation, with the ultimate goal of formally separating 'hard to represent' and 'easy to represent' problem instances.

3.11 Algorithms for Computing (Hyper)tree Decompositions

Nysret Musliu (TU Wien, AT)

Constructing decompositions of small width is crucial to solve efficiently problems based on their (hyper)tree decomposition. In recent years, several methods have been proposed to generate good (hyper)tree decompositions. Such methods include exact methods that are used to find optimal decompositions for small problems, and (meta)heuristic algorithms that find (hyper)tree width upper bounds for larger problems. In this talk, we will first give a survey of existing techniques for constructing (hyper)tree decompositions and compare these algorithms on benchmark problems from the literature. Further, we will discuss the following open questions: (1) Can we find more efficient methods to compute upper bounds for (hyper)tree width? (2) Can the existing techniques be easily adapted to generate (hyper)tree decompositions of small width that fulfill other specific conditions?

3.12 Multiresolution-multiscale Pyramids and the Traveling Salesman Problem

Zygmunt Pizlo (Purdue University – West Lafayette, US)

After presenting representative results from experiments on how human subjects produce near-optimal tours, I will describe the main aspects of pyramid models of the human visual system. One of the two main operations in pyramid models of TSP is hierarchical clustering. The second operation is a top-down sequence of approximations of a TSP tour, where centers of clusters are used in lieu of cities. The tour is produced sequentially by moving the model's attention from one city to another. Decisions on finer representations are guided by coarse representations. The model stores in its memory minimal amount of information related to the currently analyzed part of the problem. When additional information is needed, the model 'looks' at the TSP problem again. The errors and memory requirements will be presented and discussed. I will conclude by conjecturing that such a pyramid algorithm is a plausible model of human problem solving, in general.

3.13 In Search of a Cognitive Complexity Measure for Matrix Reasoning Problems

Marco Ragni (Universität Freiburg, DE)

Reasoning difficulty for items in IQ-tests is generally determined empirically: The item difficulty is measured by the number of reasoners who are able to solve the problem. Although this method has proven successful (nearly all IQ-Tests are designed this way) it is desirable to have an inherent formal measure reflecting the reasoning complexity involved. This talk will present some geometrical analogy reasoning problems and based on the types of functions necessary to solve these problems, a difficulty measure is introduced. This is finally compared to the empirical difficulty ranking as determined by Cattell's Culture Fair Test, Evans Analogy problems, and an own experiment.

3.14 Speed-Accuracy Tradeoffs: A computational Perspective

Daniel Reichman (Weizmann Institute – Rehovot, Israel)

Speed-Accuracy questions are central in several subfields of cognitive psychology such as problem solving, decision making and perception. We address several algorithmic techniques (e.g., property testing, stochastic optimization) as well as hardness results in addressing how will speed-accuracy curves look like when handling chalenging problems in psychological contexts.

3.15 Learning Productive Rules from Problem Solving Experience

Ute Schmid (Universität Bamberg, DE)

Joint work of Schmid, Ute; Kitzelmann, Emanuel
Main reference Ute Schmid, Emanuel Kitzelmann, "Inductive Rule Learning on the Knowledge Level," in: Cognitive Systems Research 12 (2011), Nr. 3, pp. 237–248.
URL http://dx.doi.org/10.1016/j.cogsys.2010.12.002

One specific characteristic of human autonomous learning is that humans are able to extract productive rule sets from experience which often is a stream of only positive examples. Following Chomsky, a productive rule set allows a person to produce systematic behavior in situations of arbitrary complexity, for example being able to build towers of sorted blocks for an arbitrary number of blocks. Such productive rules represent the competence of a person – in contrast to a person's performance which is open to unsystematic variations and errors. Furthermore, productive rule sets often are verbalizable, that is, a person can explain a general solution procedure to another person. I propose to use an approach to analytical inductive functional programming to model this type of high-level learning. Analytical inductive programming provides algorithms with clearly defined restriction and preference biases for learning recursive rule sets from small sets of positive examples.

3.16 Human Problem Solving of (hard) Computational Problems-A Computer Scientist's Thoughts and Interests

Ulrike Stege (University of Victoria, CA)

We survey typical computer science problems that are or could be used to investigate human problem solving strategies, such as the Traveling Salesperson problem and the Minimum Spanning Tree problem, as well as other graph problems. We discuss research questions and approaches that are investigated, highlight possible difficulties with current approaches and pose a set of questions that we believe are realistic to investigate in the near future.

3.17 Generalizing Muddy Children Puzzle

Jakub Szymanik (University of Groningen, NL)

Joint work of Gierasimczuk, Nina, Szymanik, Jakub;
Main reference Nina Gierasimczuk, Jakub Szymanik, "Invariance Properties of Quantifiers and Multiagent Information Exchange," Proceedings of the 12th Meeting on Mathematics of Language, Lecture Notes in Artificial Intelligence 6878, M. Kanazawa et al. (Eds.), pp. 72–89, 2011.
URL http://dx.doi.org/10.1007/978-3-642-23211-4_5

We study a generalization of the Muddy Children puzzle by allowing public announcements with arbitrary generalized quantifiers [1, 2]. We propose a new concise logical modeling of the puzzle based on the number triangle representation of quantifiers. Our general aim is to discuss the possibility of epistemic modeling that is cut for specific informational dynamics.

Moreover, we show that the puzzle is solvable for any number of agents if and only if the quantifier in the announcement is positively active (satisfies a form of variety).
Slides can be found at

http://prezi.com/96_wd3mgx_d1/a-generalization-of-the-muddy-children-puzzle/.

References

1. Nina Gierasimczuk and Jakub Szymanik. A Note on a Generalization of the Muddy Children Puzzle. *Proceeding of the 13th Conference on Theoretical Aspects of Rationality and Knowledge, K. Apt (Ed.), ACM Digital Library*, pp. 257– 264, 201. http://dx.doi.org/10.1145/2000378.2000409
2. Nina Gierasimczuk and Jakub Szymanik. Invariance Properties of Quantifiers and Multiagent Information Exchange. *Proceedings of the 12th Meeting on Mathematics of Language, Lecture Notes in Artificial Intelligence 6878, M. Kanazawa et al. (Eds.)*, pp. 72–89, 2011. http://dx.doi.org/10.1007/978-3-642-23211-4_5

3.18 Human Problem Solving: The Search for the Right Toolkit

Niels A. Taatgen (University of Groningen, NL)

Cognitive architectures have two approaches to modeling problem solving. One approach is to consider problem solving as a fundamental property of the architecture, and assume we approach new and unknown problems with set of weak methods that is the same for any individual. The Soar architecture (Newell, 1990) is an example of this approach. Other architectures, like the ACT-R architecture (Anderson, 2007), assume no architectural mechanisms for problem solving it all, but assume that problem solving consists a set of cognitive skills that have to be learned. However, models within such architectures typically encode the problem-solving strategy necessary for the task at hand, and therefore contribute little to a general account of human problem solving.

In my talk I will present a modeling framework that can serve as a starting point for a general theory of how human problem-solving skills develop within an architecture with no architectural problem-solving strategies. The idea is that the model starts with the most basic skills that are possible within a rule-based architecture, which is making single comparisons and simple elementary actions. Guided by declarative knowledge and the process of production compilation, these elementary skills can be combined to more complex skills. I will demonstrate this idea with models of cognitive transfer, in which knowledge learned in one task is used in for another.

3.19 Spatial Problem Solving: The Optimal Deployment of Cortical Resources

Sashank Varma (University of Minnesota, US)

Newell and Simon's analysis of problem solving as search through problem spaces is foundational for artificial intelligence and cognitive psychology. Because problem spaces are typically large, cognitive agents must deploy their limited resources judiciously, through planning and heuristic reasoning. The current research extends the classical conception of problem solving to the level of brain function. The cortex is understood as a set of centers, each possessing a finite supply of computational resources. Problem states, heuristics, and goals are mapped to different centers. In this view, problem solving is the optimal deployment of limited cortical resources across a network of collaborating centers. This is formalized as a linear programming problem that the brain is hypothesized to solve on a moment-by-moment basis. The resulting model provides a good account of the solution times, error rates, and brain activation fluctuations of normal adults and patients with lesions as they solve spatial problems. The implications of this research for artificial intelligence, cognitive psychology, and cognitive neuroscience are discussed.

3.20 Problem Solving as Producing a Solution

Jered Vroon (Radboud University Nijmegen, NL)

Joint work of Vroon, Jered; van Rooij, Iris; Wareham, Todd; Haselager, Pim

A new formalism for describing the structure and (associated) hardware of a system will be introduced. Within this formalism, problem solving is regarded as producing a solution rather than as a search through search space.

This formalism might be more limited than approaches that regard problem solving as a search through search space as it seems to require that a solution-producing structure is already in place. Nonetheless, even within this formalism we can distinguish between systems that require more or less structures and (associated) hardware. In this talk I will discuss these considerations and their relevance to the bigger field of problem solving.

3.21 What Does (and Doesn't) Make Problem Solving by Insight Easy? A Complexity-Theoretic Investigation

H. Todd Wareham (Memorial Univ. of Newfoundland, CA)

Problem solving is a very important and commonly-invoked cognitive ability. There are several recognized types of problem solving, e.g., by analogy, by search, by insight, and each is successful to various degrees in particular situations. Several information-processing theories have been proposed for these types. However, it is very difficult to use empirical studies to characterize the situations in which these types do and do not work, let alone link such situations to (and hence verify) the mechanisms proposed by these theories.

In this talk, we will describe an approach complementary to empirical studies which uses computational complexity analysis to assess the situations under which the mechanisms proposed by a theory can and cannot operate efficiently. Such assessments, in turn, suggest both predictions that can be verified by experiment as well as viable refinements to the theories. We will illustrate this methodology with an analysis of problem solving by insight as formulated under the Representation Change Theory of Knoblich et al.

3.22 The Way of the Ouroboros: How to Represent Problems by Solving Them

Jelle van Dijk (TU Eindhoven, NL)

This talk is based on observations and design efforts in support of the practice of 'creative problem solving' in groups. I will set the scene discussing the difference between the well-known map and the territory, and I ask whether research on human problem solving should ask how one navigates the former, or rather the latter, or even how one deals with both. Perhaps we have been concerning ourselves too much with the map. I then speculate on three alternative ways by means of which people deal with problems in everyday practice. These 'real-world tactics' may not always be in the central attention of problem solving research. They are: (1) The way of the Oyster: Not solving the problem, but encapsulating it (2) The way of the River: To let things implicitly flow into a solution; and (3) The way of the Ouroboros: Representing the problem by first executing the solution. This strange option I call the Ouroboros, i.e. creating a representation by executing a problem solution, seems to put the cart before the horse and therefore nonsense. I will nonetheless discuss a number of variations of this strategy that I think exist, and actually work, in everyday practical circumstances. In order to ground the idea I will relate it to some theoretical notions from Embodied and Situated Cognition theory, as well as to my empirical observations of creative group sessions (aka brainstorms). My question for the seminar is whether these ideas implies a completely new line of research, or whether it is possible to integrate these ideas into existing models and theories on problem solving (or whether they are really just nonsense).

4 Working Groups

4.1 Key terms and their meanings

During Day 1 it became clear that speakers used terms whose meanings were unclear for some or many participants in the audience. It was decided to keep track of these terms by listing them on the blackboard. All subsequent speakers were asked to define these terms whenever they used them in their talks. Over the course of the seminar, the list grew to include the following terms:

- optimization
- representation
- search
- heuristic
- problem
- insight
- chunk
- cognition
- relevant
- complexity
- embodied cognition
- model
- situated cognition
- knowledge
- distributed cognition
- communication

In a workgroup session on Day 2, small groups of 4 participants (mixed groups of 2 computer scientists and 2 psychologists) were invited to pick one word from the above list and discuss all of its possible meanings. Participants were explicitly instructed not to try to decide on one 'proper' or agreed upon meaning, but rather to generate as many possible different meanings as seemed relevant to the domain of problem solving. Interestingly, of the long list of words there were three words that were a recurrent topic of discussion. These were 'problem', 'representation', and 'model'. It became clear that the meanings of even these three central words differed both between and within computer science and psychology.

The exercise was intended to raise awareness of the different usages of words by different researchers in the area of problem solving. The reasoning of the organizers was as follows: For interdisciplinary collaborations to get off the ground researchers need to be able to speak each other's languages, negotiate meanings, and develop new terminology as the need arises. This exercise helped to foster such an open minded atmosphere.

4.2 Promoting interdisciplinary discussion

On Day 3 small groups of 2 computer scientists and 2 psychologists were formed to discuss the following questions.

1. Why study problem solving in an interdisciplinary manner? Where could one discipline help the other?

2 What do you need to understand each other? What do you need to know to know how to help the other?
3 Do you see insurmountable differences, obstacles, or challenges?
4 What is (human) problem solving? Why study it? What are important research questions?
5 Can we identify different classes of problem solving, and characterize their relative difficulty?

There was general agreement among the participants that an interdisciplinary approach to problem solving would be desirable and possible, and that seminars like this one are useful for building the necessary 'common ground'. Subsequently, the participants added several more questions to the list:

6 When is a problem solved?
7 What type of problems do we want to include in this area of study?
8 What computational methods may be used to investigate these research questions?
9 What experimental paradigm to use?
10 Is there to be one or multiple theories of problem solving?

Question 6 was motivated by the observation that in computer science a problem $f : X \rightarrow Y$ is said to be solved when an input $x \in X$ is translated to an output $f(x) \in Y$. Yet, solving the problems presented by the psychologist Batchelder sometimes meant something like 'understand the reason or motive for the behavior or situation in the scenario'. Can the latter notion also be mapped to the input-computation-output paradigm? Question 7 was raised because many different cognitive abilities could count as examples of problem solving, e.g., visual problem solving and common sense reasoning. Should we focus our research on some of these, or consider all of them? Question 8 was raised because a variety of computational methods could be adopted in problem solving research, such as models, architectures and algorithms. Question 9 was raised because fruitful research in psychology often depends on a stable experimental paradigm with interpretable dependent measures (such as accuracy and speed). Last, Question 10 was raised because some participants felt that problem solving theories may be as diverse as the different types of problems out there, whereas other participants were committed to building unified or integrative accounts of problem solving in general.

4.3 New Ideas and Collaborations

On Day 4 of the seminar a working group was scheduled in which participants were invited to think about and try to come up with new cross-disciplinary collaborative projects and/or identify important open problems in the field of problem solving. The ideas generated in this workgroup were presented either the same day or, when ideas needed to be first further developed, on the morning of Day 5. A large group of participants presented ideas inspired by the seminar and/or new collaborations. All participants have furthermore been invited to be submit their work presented at or inspired by the seminar for consideration for publication in *The Journal of Problem Solving* (see Section 6). In the Section 5, we present a selection of the ideas that were presented at the seminar that we judge to be particularly original or important, and of general interest.

4.4 Wrap-up session: Evaluation and outlook

Iris van Rooij chaired the wrap-up session, and considered the following questions.

- Did we get (closer to) a shared notion of 'problem solving' (and 'the study of problem solving')?
- Did we get (closer to) new foundations?
- Did we get (closer to) new ideas for formalizing notions such as 'representational complexity' (e.g., 're-representation', 'insight', 'ill-defined', etc.).

She argued that we could answer all these questions in the affirmative. The questions 'what is problem solving?' and 'what distinguishes problem solving from other cognitive domains?' was also a recurring topic of discussion at the two preceding Purdue workshops on Problem Solving in 2005 and 2008. It appears that these questions plague the domain of problem solving more than other cognitive domains (for reasons unknown, though speculations range from the idea that problem solving is not a unified category of cognitive processes, to the idea that it is but that we have too few good experimental paradigms for studying problem solving in the lab). Even though this seminar has not produced definite answers to these questions, there does seem to be a convergence of ideas on what defines the different types of key mental processes involved in problem solving. This consideration of subprocesses of problem solving has even motivated a new classification of areas of problem solving research that may inform and guide future research and theorizing in the field (see section 5.3). As for the hoped-for progress in the formal foundations of problem solving research, and the notion of 'representational complexity' in particular, novel ideas have also been put forth, for instance in the presentations by Wareham, Kwisthout, and Ragni, and the open problem proposed by Haxhimusa and van Rooij.

In the wrap-up session participants were also asked for their feedback about the seminar and for recommendations for a follow-up seminar. One idea that was raised was that a follow-up seminar could aim to prepare for a Handbook on Problem Solving, and that the first day of the program could include tutorials, e.g., about computer science for psychologists and about psychology for computer scientists. A brainstorm on topics for a follow-up seminar resulted in the following list: Problem solving in the real-world versus the lab; Cognitive architectures and problem solving; Problem solving in the large; Problem solving of dynamic problems; Social problem solving; Human-inspired machine problem solving; Spatial problem solving; Problem solving with bounded resources (Bounded Rationality).

5 New Ideas and Open Problems

In this section we present a selection of open problems and new ideas that were presented at the morning session on Day 5 of the seminar.

5.1 A 'Turing Test' for Problem Solving

Ken Forbus presented an analog of the Turing test for intelligence for the domain of problem solving. He coined this the *Dagstuhl-Batchelder test*, as it was specifically inspired by the computational challenges posed by the 19 problems presented by Bill Batchelder at this seminar. The idea behind this test is that a system could be said to engage in genuine problem solving if it could solve at least these 19 problems. Importantly, the test should be

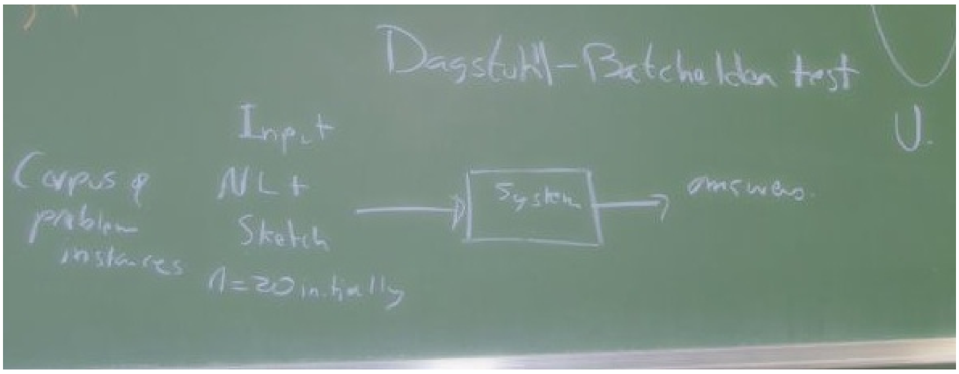

Figure 1 The Dagstuhl-Batchelder test for Problem Solving

fair and representative of how humans can solve the problems. Therefore Forbus imposed the constraint that the inputs to the system should be the raw text and images as presented in Appendix 7.1. In addition the system is allowed to have a knowledge data base, which could for instance consist in sketches of situations etc. Figure 1 illustrates the idea.

5.2 A Complexity Hierarchy of Insight Problems

Haxhimusa and van Rooij posed an open question: Is it possible to formulate a hierarchy of complexity classes for insight problems analogous to the computational complexity classes for search problems? They proposed that such a hierarchy may define classes in terms of the number of changes c to the input representation required to turn an insight problem (conceived as an ill-defined search problem) into a well-defined (potentially trivial) search problem. Here c may be thought of as the number of basic 'insights', 'pieces of information', 'hints' or 're-representation steps' needed to turn an insight problem into a well-defined problem (cf. Todd Wareham's proposal for a similar framework). In the proposed hierarchy $C_0 \subset C_1 \subset C_2 \subset ... \subset C_{n-1} \subset C_n$, the class C_0 would denote the class of well-defined problems, i.e., problems requiring no change in the input in order to become well-defined. Further, each class C_k in the hierarchy consists of those problems that are at most $c = k$ changes away from some problem in C_0. Interesting open questions are the following: Can this idea for a complexity hierarchy for insight problems be worked out to a formal framework? And if so, would it possible to use the framework to characterize, explain and/or predict the difficulty of different classes or types of insight problems for humans?

5.3 A Classification of Problem Solving Research(ers)

Todd Wareham proposed a classification of problems in terms of the nature of the subprocesses that are (assumed to be) invoked during the problem solving process (see Table 1). He observed a scale of problem classes ranging from search only, to restructuring combined with search, to problems that require access to and processing of world knowledge over and above the restructuring and search processes involved. Wareham observed that each class of problem seems to have at least one associated representative researcher, each of which was present at the seminar. This way of conceptualizing different classes of problems seems very intuitive and may prove useful for the field to understand how different problems and models of problem solving differ and relate to each other.

Table 1 Classification of Problems.

Processes assumed to be involved in problem solving	Representative researcher in psychology	Representative researcher in computer science
search	Pizlo	Stege
restructuring + search	Chu	Wareham
restructuring + search + world knowledge	Batchelder	Forbus

5.4 Verbal Reports Revisited

Frank Jäkel proposed that, lacking a firm theoretical foundation to date, research on problem solving by *insight* may do well to revisit the methodology that lay the foundations for the theory of problem solving by *search* developed by Newell and Simon, viz., verbal reports made by humans about what they are thinking while they are engaged in problem solving. Jäkel pointed out that this methodology has fallen out of favor in psychology because it is a form of introspection and therefore considered unreliable for understanding the nature of cognitive processes. As a consequence, the methodology has been replaced by methodologies using simpler behavioral measures such as speed and accuracy of problem solving. Jäkel makes an important observation however. Even though verbal protocols are based on a form of introspection and may be to some extent unreliable, they are very rich in information that can be useful for hypothesis generation and theory formation. For instance, revisiting the 1972 book by Newell and Simon on problem solving reveals that many of their hypotheses about 'means-end analysis' and 'heuristic search' were a direct consequence of the meticulous analysis of verbal reports of people solving search problems. It is also noteworthy that the methodology for verbal reports has been refined considerably since the heyday of introspection in early psychology without the mainstream of cognitive psychology really taking notice of these developments [1]. In addition, verbalizing and inner speech are an important part of problem solving anyway. Hence, even if verbal reports do not constitute a rigorous *test* of theories of problem solving, they may prove useful for *coming up* with new theories of insight problem solving, which later can be tested using other measures.

References

1 K.A. Ericsson and H.A. Simon. Verbal Reports as Data. *Psychological Review*, 87(3):215–251, 1980.

6 Dissemination of Results

All participants have been invited to submit their research presented at this seminar or inspired by this seminar for consideration for publication in *The Journal of Problem Solving* (JPS). JPS is an open access journal with an interdisciplinary readership. We plan to have two special issues: one in the Spring and the other in the Fall of 2012. Considering the fact that papers in JPS can be accessed by everyone (no subscription is required), the proceedings from this workshop are expected to be read widely and have large impact. Once the special issues are published, Purdue University Press (the publisher) will print a book with the published papers.

JPS (ISSN 1932-6246) is a multidisciplinary journal that publishes empirical and theoretical papers on mental mechanisms involved in problem solving. The journal welcomes original and rigorous research in all areas of human problem solving, with special interest in solving difficult problems (e.g., problems in which human beings outperform artificial systems). Examples of topics include (but are not limited to) optimization and combinatorial problems, mathematics and physics problems, theorem proving, games and puzzles, knowledge discovery problems, insight problems and problems arising in applied settings. JPS encourages submissions from psychology, computer science, mathematics, operations research and neuroscience. More information on journal web site: http://docs.lib.purdue.edu/jps/

Editor-in-Chief: Zygmunt Pizlo, Department of Psychological Sciences, Purdue University.

7 Appendices

7.1 19 Classic Insight Problems, by Bill Batchelder

Classic insight problems presented by Bill Batchelder are shown in Figures 2, 3, 4. In these figures Batchelder has adapted several famous problems from the folklore of brain teasers. Batchelder selected these problems because they are not move problems in the sense of Newell and Simon, and the main barrier to solution is finding a productive representation. Batchelder acknowledges the original author of some of these problems, even though they are not shown in the figures.

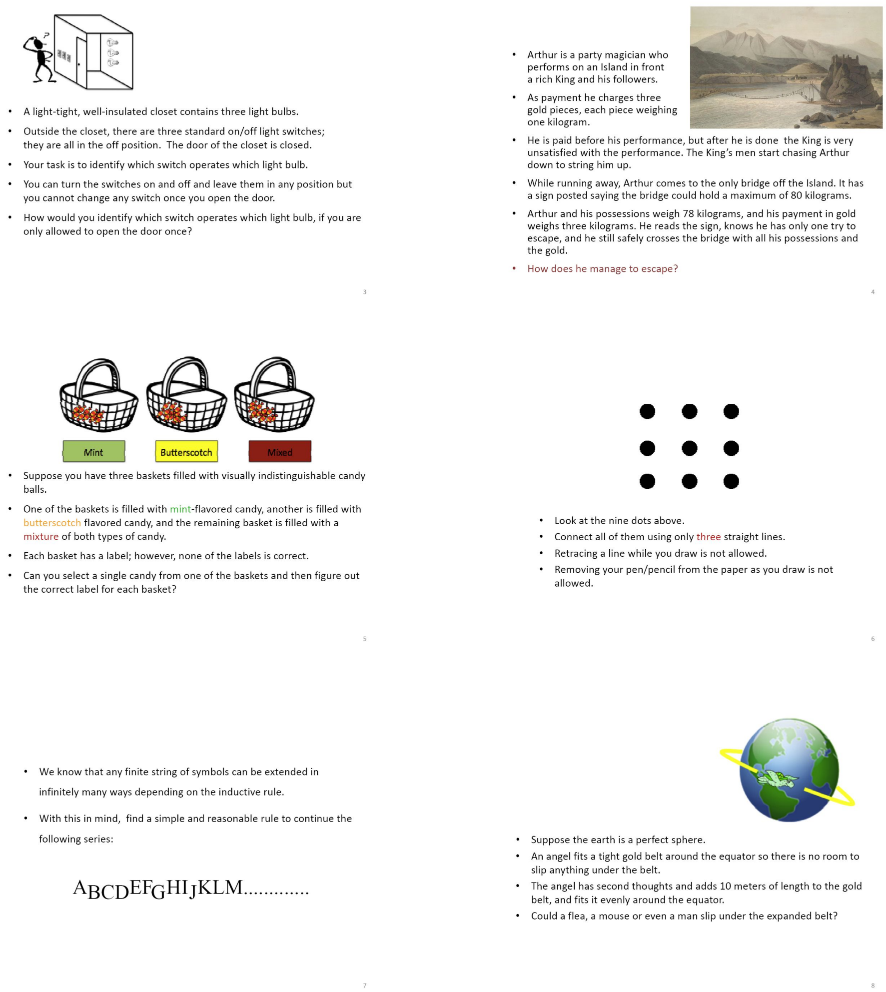

Figure 2 19 insight problems.

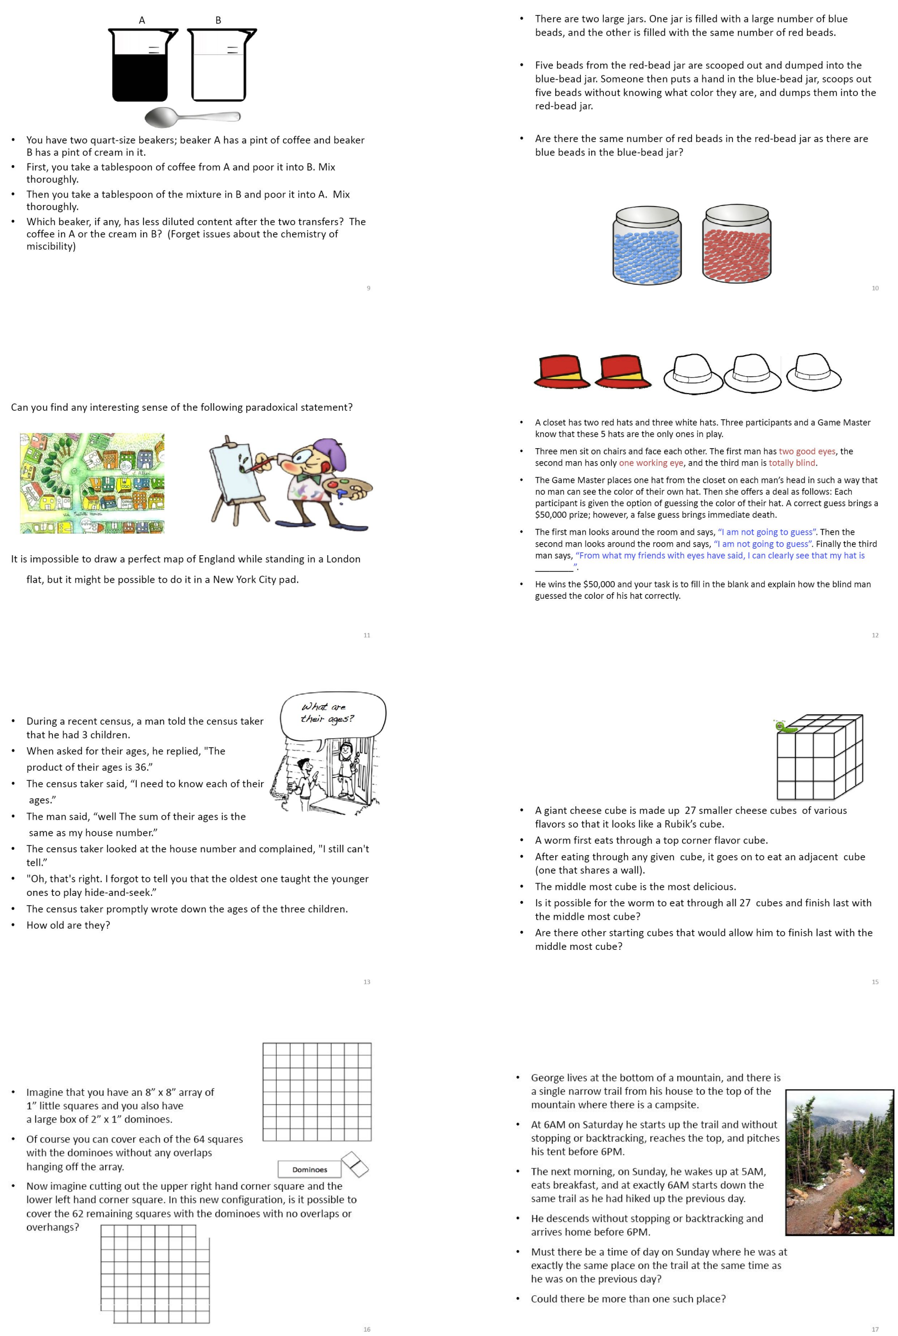

Figure 3 19 insight problems, cont.

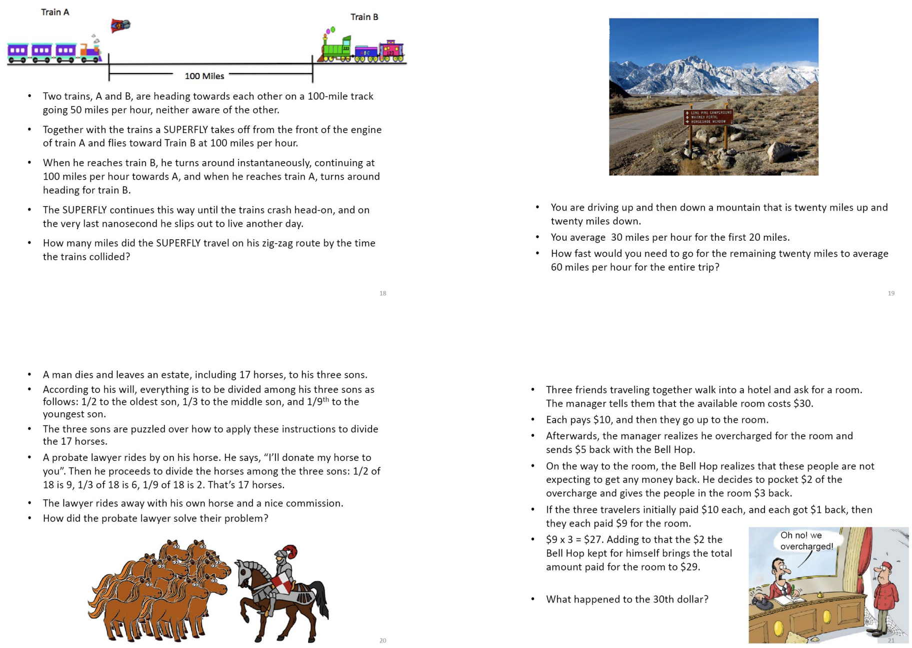

Figure 4 19 insight problems, cont.

7.2 Definition of Basic Terms in Insight Problem Solving, by Yun Chu

Some possible definitions for terms often used in insight problem solving.

▶ **Definition 1** (Problem)**.** A problem occurs when there is an obstacle between a present state and a goal state and it is not immediately obvious how to get around the obstacle.

▶ **Definition 2** (Well-defined problem)**.** It is a clear problem representation with the initial state, goal state, obstacles to the goal state, and the solution path stated.

▶ **Definition 3** (Ill-defined problem)**.** It lacks a clear path to the solution or the operators are not specified.

▶ **Definition 4** (Insight problem)**.** It is one type of ill-defined problem. 'obvious' solutions do not work, usually low solution rates, sudden realization of the solution.

▶ **Definition 5** (Metacognition)**.** It is 'thinking about thinking.'

▶ **Definition 6** (Verbalization)**.** It is talking about what you are doing/thinking while in the problem solving process.

▶ **Definition 7** (Feelings-of-warmth)**.** It is asking the problem solver to provide this rating in answer to the question, 'how close do you feel to the solution?'

Below a list of insight problems is shown.
Some classic insight problems are listed below. For more see Appendix 7.1.

Verbal: Marsha and Marjorie were born on the same day of the same month of the same year to the same mother and father yet they are not twins. How is this possible?

Math: There are 10 bags, each containing 10 gold coins, all of which look identical. In 9 of the bags, each coin is 16 ounces, but in one of the bags, the coins are 17 ounces each. How is it possible (in a single weighing) to determine which bag contains the 17-ounce coins?

Spatial: 9-dot problem. Connect 3 rows of 3 dots each with 4 straight lines without lifting your pencil or retracing any lines.

Some recent insight problems are shown below.

Matchstick arithmetic: Move 1 matchstick to make the following statement true: `IV = III = I`

Compound remote associates: Find the solution word associated with all words of the triad forming 3 compound words: age mile sand

Rebus: What is the common saying (fill the empty part)? iii ______ ooo

Cheap necklace problem: Make a closed necklace with 4 chains of 3 links each. You have 15 cents total. It costs 2 cents to open a link and 3 cents to close it.

8-Ball problem: There are 8 balls in front of you. One of them is slightly heavier than the other 7. Using a balance scale only two times, how can you find the heavy ball?

7.3 Computational Search Problems, by Ulrike Stege

▶ **Definition 1** (Euclidean TSP)**.** Input: A set of points in the Euclidean Plane. Output: A shortest tour connecting all the points.

▶ **Definition 2** (Euclidean MST)**.** Input: A set of points in the Euclidean Plane. Output: A shortest network/graph connecting all the points.

▶ **Definition 3** (Vertex Cover)**.** Input: A(n undirected) graph $G = (V, E)$ Output: A smallest vertex cover V' for G. That is, a subset V' of V where for each edge xy in E, x or y is in V' and V' is as small as possible.

▶ **Definition 4** (Independent Set)**.** Input: A(n undirected) graph $G = (V, E)$ Output: A largest independent set V' for G. That is, a subset V' of V such that for each pair x, y of vertices in V', xy is not an edge in E and V' is maximized.

▶ **Definition 5** (Dominating Set)**.** Input: A(n undirected) graph $G = (V, E)$ Output: A smallest dominating set V' for G. That is, a subset V' of V such that for each vertex $x \in V$, x is in V' or x is adjacent to a vertex y that is in V' and V' is as small as possible.

8 Seminar Program

Monday, 29th of August 2011

Chair: Iris van Rooij

09:00 – 09.45	Introduction of the Seminar. Short presentation of the participants.
09:45 – 10:45	Todd Wareham. *What Does (and Does not) Make Problem Solving by Insight Easy? A Complexity-Theoretic Investigation.*
11:15 – 11:45	Georg Gottlob. *Living with Computational Complexity.*
11:45 – 12:15	Sarah Carruthers. *Vertex Cover and Human Problem Solving.*
14:00 – 14:30	Discussions
14:30 – 15:30	William Batchelder. *Some Issues in Developing a Theory of Human Problem Representation.*
16:00 – 16:30	Sashank Varma. *Spatial Problem Solving: The Optimal Deployment of Cortical Resources.*
16:30 – 16:40	Jakub Szymanik *Generalizing Muddy Children Puzzle.*
16:40 – 18:00	Working group and Discussions.

Tuesday, 30th of August 2011

Chair: Yll Haxhimusa

09:00 – 10.00	Niels Taatgen. *Human problem solving: the search for the right toolkit.*
10:00 – 10:30	Johan Kwisthout. *Relevant Representations.*
11:00 – 12:00	Rina Dechter. *Advanced Reasoning in Graphical models.*
14.00 – 15.00	Ken Forbus. *Analogy as a computational foundation for problem-solving and learning.*
15.00 – 15:30	Jelle van Dijk. *The way of the Ouroboros: How to represent a problem by solving it.*
16:00 – 16:30	Marco Ragni. *In Search of a Cognitive Complexity Measure for Matrix Reasoning Problems.*
16:30 – 18.00	Discussions and Working groups.

Wednesday, 31st of August 2011

Chair: Iris van Rooij

09:00 – 10.00	Dedre Gentner. *The Analogical Mind.*
10:00 – 10:30	Liane Gabora. *Problem Solving as the Recognition and Actualization of Potentiality.*
11:00 – 11:30	Daniel Reichman. *Speed-Accuracy Tradeoffs: A computational Perspective.*
14.00 –	Hiking.

Thursday, 1st of September

Chair: Yll Haxhimusa

09:00 – 09:30	Yun Chu. *Human Performance on Insight Problem Solving: A Review.*
09:30 – 10:00	Ute Schmid. *Learning Productive Rule Sets from Problem Solving Experience.*
10:00 – 10:30	Ulrike Stege. *Using (even more) Foundations from Computer Science to study Aspects of Human Problem Solving.*
11:00 – 12:00	Open Discussion.
14:00 – 14:30	Brendan Juba. *PAC Semantics: a Framework for Heuristic Rules.*
14:30 – 15:00	Nysret Musliu. *Algorithms for Computing (Hyper)tree Decompositions.*
15:00 – 15:15	Jered Vroon. *Problem Solving as Producing a Solution.*
15:15 – 15:30	Zygmunt Pizlo. *Multiresolution-multiscale pyramids and the traveling salesman problem.*
16:00 – 17:00	Preparation: Open Problem Session.
17:00 – 18:00	Open Problem Session.

Friday, 2nd of September 2011

Chair: Iris van Rooij

09:00 – 10.30	Working group, Discussions and Short Talks.
11:00 – 12:00	Wrap-up Session.

Participants

- William H. Batchelder
University of California, US
- Mark Blokpoel
Radboud Univ. Nijmegen, NL
- Sarah Carruthers
University of Victoria, CA
- Yun Chu
La Crescenta, US
- Rina Dechter
Univ. California – Irvine, US
- Kenneth D. Forbus
Northwestern University – Evanston, US
- Liane Gabora
University of British Columbia – Vancouver, CA
- Dedre Gentner
Northwestern University – Evanston, US
- Noah D. Goodman
Stanford University, US
- Georg Gottlob
University of Oxford, GB
- Marcin Grzegorzek
Universität Siegen, DE
- Yll Haxhimusa
TU Wien, AT
- Jens Hedrich
Universität Koblenz-Landau, DE
- Adrian Ion
TU Wien, AT
- Frank Jäkel
Universität Osnabrük, DE
- Brendan Juba
MIT – Cambridge, US
- Markus Knauff
Universität Giessen, DE
- Walter Kropatsch
TU Wien, AT
- Johan Kwisthout
Radboud Univ. Nijmegen, NL
- David Landy
University of Richmond, US
- Zoltan Miklos
EPFL – Lausanne, CH
- Nyrset Musliu
TU Wien, AT
- Zygmunt Pizlo
Purdue University – West Lafayette, US
- Marco Ragni
Universität Freiburg, DE
- Daniel Reichman
Weizmann Inst. – Rehovot, IL
- Ute Schmid
Universität Bamberg, DE
- Ulrike Stege
University of Victoria, CA
- Jakub Szymanik
University of Groningen, NL
- Niels A. Taatgen
University of Groningen, NL
- Susanne Tak
University of Canterbury – Christchurch, NZ
- Jelle van Dijk
TU Eindhoven, NL
- Iris van Rooij
Radboud Univ. Nijmegen, NL
- Sashank Varma
University of Minnesota, US
- Jered Vroon
Radboud Univ. Nijmegen, NL
- H. Todd Wareham
Memorial Univ. of Newfoundland, CA
- Gerhard Woeginger
TU Eindhoven, NL