Report from Dagstuhl Seminar 14341

Resource-bounded Problem Solving

Edited by

Yll Haxhimusa[1], Iris van Rooij[2], Sashank Varma[3], and H. Todd Wareham[4]

1 **Vienna University of Technology, AT,** yll@prip.tuwien.ac.at
2 **Radboud University Nijmegen, NL,** i.vanrooij@donders.ru.nl
3 **University of Minnesota, US,** sashank@umn.edu
4 **Memorial Univiversity of Newfoundland, CA,** harold@mun.ca

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 14341 'Resource-bounded Problem Solving'. This seminar is a successor to Dagstuhl Seminar 11351: 'Computer Science & Problem Solving: New Foundations', held in August 2011, which was the first Dagstuhl event to bring together computer scientists and psychologists to exchange perspectives on problem solving in general. The current seminar built on the previous seminar by (1) narrowing the focus to issues related to resource-bounded problem solving and (2) broadening the range of perspectives on the specific topic by including a balanced number of attendees with expertise in computer science, psychology, artificial intelligence, and cognitive neuroscience.

Seminar August 17–22, 2014 – http://www.dagstuhl.de/14341

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, F.2.3 Tradeoffs Between Complexity Measures, I.2.0 [Artificial Intelligence] General, I.2.6 Learning, I.2.4 Knowledge Representation Formalisms and Methods, I.2.8 Problem Solving, Control Methods, and Search, J.4 Social and Behavioral Sciences

Keywords and phrases complexity theory, problem solving, cognitive psychology, computational trade-offs

Digital Object Identifier 10.4230/DagRep.4.8.45

1 Executive Summary

Yll Haxhimusa
Iris van Rooij
Sashank Varma
H. Todd Wareham

This Dagstuhl Seminar on 'Resource-bounded Problem Solving' was a successor to Dagstuhl Seminar 11351: 'Computer Science & Problem Solving: New Foundations', held in August 2011, which was the first Dagstuhl event to bring together computer scientists and psychologists to exchange perspectives on problem solving in general. Before summarizing the content of the seminar itself, we describe the theoretical motivations for the topic of 'Resource-bounded Problem Solving', and the choice for the interdisciplinary composition of participants, ranging complexity theory, cognitive psychology, artificial intelligence and cognitive neuroscience.

Resource-bounded Problem Solving, *Dagstuhl Reports*, Vol. 4, Issue 8, pp. 45–72
Editors: Yll Haxhimusa, Iris van Rooij, Sashank Varma, and H. Todd Wareham
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Background and Motivation

Problem solving, whether by humans or machines, is bounded by the resources at hand. For machines, these resources fundamentally include hardware and processing speed. For humans, important resources also include mental representations, memory capacities, inferential capacities, and time to name a few. All these resources can be severely limited, by constraints imposed by both the implementing hardware (current technology in the case of machines, the organization of our brains in the case of humans) and the physical and social operating environments.

The study of resource-bounded problem solving has a long and productive history within computer science, which has resulted in a number of efficient exact-solution algorithms and algorithm design techniques as well as, where such algorithms are not possible, various widely-applicable approximate-solution heuristics. Given that brains have evolved to solve problems under severe resource constraints, resource-bounded problem solving may also provide one of the best windows on the organization of cognitive brain function. Trying to model exactly how humans solve really hard—that is, resource demanding—problems efficiently seems a good strategy from the perspective of deriving predictive and explanatory cognitive models. After all, many cognitive models may be able to match human performance on really easy problems, but only a few can for really hard problems. Hence, if one finds even one cognitive model that can solve a hard problem as well as humans one can be much more confident that it captures fundamental principles of human cognition.

What makes tasks resource demanding? Here computational complexity theory is a key source of information for the study of human and machine problem solving. Computational complexity theory studies the intrinsic resource demands of various computational problems. It allows us to assess when resource demands are low, reasonable, high, or impractically high. Though the importance of computational complexity has been recognized in computer science for decades, it has been underutilized to date by cognitive scientists. This is not for want of opportunities, as cases are known where cognitive scientists have studied principles of resource-bounded problem solving in apparent ignorance of relevant computational complexity results. The following example illustrates such a situation.

- Solving constrained Traveling Salesman problems: MacGregor and Ormerod (1996,Attention, Perception & Psychophysics) hypothesized that the difficulty of solving Euclidean versions of the Traveling Salesperson problem (E-TSP) may be more a function of the number of inner points (i. e., the points not lying on the convex hull of the point set) than of the total number of points. This hypothesis was tested and confirmed for human subjects solving pen-and-paper instances of E-TSP. Independently, it was shown that E-TSP is fixed-parameter tractable when the number of inner points is the parameter (Deineko et al.COCOON 2004). In other words, it is possible to solve E-TSP in time $f(k)n^c$, where fs is a non-polynomial function of the number k of inner points, n is the total number of points, and c is a constant. This result is relevant for explaining the findings of MacGregor and Ormerod (1996) as it gives a computational formalization of their hypothesis.

There may be many other such opportunities waiting to be noticed. There is also evidence that when such opportunities are exploited and cognitive scientists and complexity theorists establish collaborations, these collaborations can yield novel perspectives and approaches. Below are two examples of such ongoing collaborations and their products.

- Analyzing resource demands of cognitive models: Using computational complexity concepts and techniques, psychologists can systematically study how human problem solving proceeds under various resource demands. Also, complexity theory can predict how resource demands scale as a function of a problem's parameters. Psychologists can then in turn use these predictions to test the models being studied. This approach has been successfully implemented by Iris van Rooij, Todd Wareham, and others in a wide variety of domains, including decision-making (2005, Journal of Mathematical Psychology) and analogical problem solving (2011, Journal of Problem Solving). This program has also led to the development of a theory of structure approximability which has produced new results within both computer science (2007, Proceedings of Dagstuhl seminar07281) and psychology (2012, Journal of Mathematical Psychology).

- Pyramid data structures and efficient search: Humans and animals are able to delineate, detect and recognize objects in complex scenes at a blink of an eye. Tsotsos (1990, Behavioral and Brain Sciences) performed a complexity analysis and showed that hierarchical representation of visual information and hierarchical processing of this information is one of the best, if not the best, way for brains to solve visual problems. Pyramid data structures provide an effective model for efficient hierarchical search of the problem space. This perspective has led to fruitful collaborations between Yll Haxhimusa, Zygmunt Pizlo, Walter Kropatsch and others, yielding new algorithmic techniques in computer vision (2005, Pattern Recognition Letters; 2009, Vision & Computing),as well as inspiring new cognitive models of visual problem solving in psychology (2006 and 2011, Journal of Problem Solving).

With this seminar we aimed to actively stimulate the exchange of ideas and results between computational complexity theorists, psychologists, cognitive neuroscientists, and AI-researchers studying problem solving. In particular, we wanted to ensure that this exchange would be of use to each (and not just one) of these communities. We believe that such n-way productivity is crucial to fruitful long-term interdisciplinary collaboration, in that it encourages the continued interest of members of all communities in collaborating.

Organization of seminar

On Day 1 of the seminar, Iris van Rooij opened the seminar by explaining its history, motivation and aims. This was following by a round of introductions, in which each participant introduced themselves, who they are, what their home disciplines were, what their relevant research interests were, and what they hoped to both bring to the seminar and get out off it.

Given the interdisciplinary nature of the questions of interest and the wide range of expertise of the seminar participants, it was crucial that a common understanding of the different goals and assumptions of the disciplines involved at the meeting be established early in the meeting. To this end, the first full day of the seminar was devoted to primers on basic terminology, assumptions, and goals of four major sub disciplines represented at the seminar (namely, computational complexity theory, artificial intelligence, psychology, and cognitive neuroscience).

The four introductory keynote speakers, Zygmunt Pizlo (Cognitive Psychology), Todd Wareham (Computational Complexity Theory), Rineke Verbrugge (Artificial Intelligence), and David Noelle (Cognitive Neuroscience), all addressed each of the following questions for their own respective fields.

- What are the goals of that discipline?
- What are the techniques used in that discipline?

- What do the terms “problem solving” and “resource bounds” mean in that discipline?
- What does that discipline have to offer to other disciplines in the context of this seminar?
- What does that discipline want from the other disciplines in the context of this seminar?
- What are some tentative research questions and collaboration opportunities?

Following these introductory key notes there was a panel discussion. Day 1 closed with a Town Hall meeting, in which the set-up and organization for the next days was discussed with all participants and a preliminary schedule was established (later on, as needed, this schedule was updated). At the Town Hall meeting the concept of Birds-of-a-Feather (B.O.F.) sessions was also explained (see Section 5), which turned out to be a very successful format for self-organized working groups.

Days 2, 3 and 4 involved a mix of talks, posters, B.O.F. sessions, and Town Hall sessions. Pairs or triplets of talks were followed by panel discussions to stimulate cross talk connections. Poster sessions allowed for more in-depth discussion in an informal setting, and B.O.F. sessions allowed people to gather and discuss more specific topics of common interest. At Town Hall sessions, a plenary report on the insights gained from each B.O.F. session was given, so that all participants were kept up to date of the outcome of such events. All in all, this set-up worked very well, stimulating active exchange and discussion between participants that crossed disciplinary boundaries.

On the morning of Day 5 we closed with reflections on the overall organization and content of the seminar. It was a shared perspective among participants that the seminar had been exceptionally successful in bringing together the different fields involved in the seminar, initiating many first-time theoretical exchanges and conceptualizations of possible common research questions. Many participants indicated that following this seminar, they would be interested in more focused seminars specializing in subtopics within the domain of problem solving or specializing in specific modeling techniques.

2 Table of Contents

3 Overview of Talks

3.1 Where is the chocolate?: Modeling Development of Reasoning about False Beliefs of Others

Burcu Arslan (University of Groningen, NL)

Reasoning about false beliefs of others develops with age. We constructed a computational cognitive model in order to show the developmental transitions. These start from a child's reasoning from his/her own point of view (zero-order) to taking into consideration another agent's beliefs (first-order), and later to taking into consideration another agent's beliefs about again other agents' beliefs (second-order). The model is based on a combination of rule-based and simulation approaches. We modeled the gradual development of reasoning about false beliefs of others by using activation of declarative knowledge instead of utility learning. Initially, in addition to the story facts, there is only one strategy chunk, namely a zero-order reasoning chunk, in declarative memory. The model retrieves this chunk each time it has to solve a problem. Based on the feedback, the model will strengthen a successful strategy chunk, or it will add or strengthen an alternative strategy if the current one failed.

3.2 When Thinking Never Comes to a Halt: Tractability and Approximability in AI

Tarek R. Besold (Universität Osnabrück, DE)

Joint work of Besold, Tarek R.; Robere, Robert
Main reference T. R. Besold, R. Robere, "When Thinking Never Comes to a Halt: Using Formal Methods in Making Sure Your AI Gets the Job Done Good Enough," in V. C. Mueller, editor, Fundamental Issues of Artificial Intelligence (Synthese Library), Springer, to appear.

A growing number of researchers in Cognitive Science advocate the thesis that human cognitive capacities are constrained by computational tractability. If right, this thesis also can be expected to have far-reaching consequences for work in Artificial General Intelligence: Models and systems considered as basis for the development of general cognitive architectures with human-like performance would also have to comply with tractability constraints, making in- depth complexity theoretic analysis a necessary and important part of the standard research and development cycle already from a rather early stage. We present an application case study for such an analysis based on results from a parameterized complexity and approximation theoretic analysis of the Heuristic Driven Theory Projection (HDTP) analogy-making framework.

3.3 The Challenge of Optimization

Sarah Carruthers (University of Victoria, CA)

How do humans cope with complexity? One approach that can assist our understanding in this area is to study human performance on hard computational problems. Studies to date have focused on the optimization version of a handful of problems, including the Traveling Salesperson Problem, Minimum Vertex Cover, and n-Puzzle. In this work, we show that for these hard optimization problems, participants' internal representations cannot be consistent with the given problem. Identifying how internal representations are generated under these conditions presents an opportunity for new interpretation of results from previous studies. With this in mind, we propose alternative ways of presenting hard computational problems such that participants' internal representations can be consistent with the intended task. Finally we present study results, which support our hypothesis that the internal representation generation differs between optimization and other versions of two hard computational problems.

3.4 Brain-Wiring Optimization: Lessons for Cognitive Modeling

Christopher Cherniak (University of Maryland – College Park, US)

A (a) bounded-resource rationality framework links together independent research programs in (b) cognitive psychology and (c) computer science. Neural wiring minimization sometimes appears either virtually perfect, or as good as can be currently detected (vs only moderately well-engineered). – A predictive success story. Quick & dirty heuristics for such neuroanatomy optimization may be germane to problem solving procedures:

1. Neural network minimization 'for free, directly from physics'. Via exploiting anomalies of the mathematical / computational order / the optimization landscape.
   a. Evading exhaustive search of alternative topologies, when 'topology does not matter'.
   b. Avoiding local minima traps in energy-minimization layout processes.
   c. Finetuning a fast heuristic to match only best cases of wirecost minimization.
2. Detecting large-scale system optimization indirectly, via a 'Size Law': The larger the proportion of a complete optimal system that an evaluated sub-system is, the better the optimization of the subset.

3.5 Ifs and Ands and Ors

Nicole Cruz de Echeverria Loebell(University of London, GB)

Joint work of Loebell, Nicole Cruz de Echeverria ; Baratgin, Jean; Feeney, Aidan; Oaksford, Mike; Over, David

The probabilistic approach in the study of reasoning generalizes binary logic to cover the correctness of deductive inferences from uncertain premises. In this new framework, the concept of binary validity is extended to that of probabilistic validity, p-validity, and the concept of binary consistency is extended to that of coherence. Many results in judgment and decision making show that people are not always coherent in their probability judgments. Our question was whether they are more coherent, and conform more to p-validity, when they make explicit deductive inferences. This question was investigated in an experiment using conditionals and disjunctions. Participants gave confidence judgments about a list of inferences (inference group) or about the statements these inferences were composed of (statements group) based on information about a character from a short context story. Overall participants' probability judgments conformed to p-validity and coherence at above chance levels. They conformed less often to p-validity for binary valid, but p-invalid inferences, and they conformed more often to p-validity and coherence in the inference group than in the statements group. The results provide support for the psychological relevance of the normative standards of p-validity and coherence. And they suggest that conformance to these standards is not automatic, but increases in the context of an explicit inference task. The possibility of reducing the complexity of computing coherence when engaged in an inference task through the manipulation of graphic representations of logical partitions is discussed.

3.6 Emergence of Complex, Problem-Solving Cognitive Structure

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

How does a problem solving cognitive architecture come about? Humans are unique in extent to which we take constraints and resources at hand into account by adapting existing problem solving methods to new situations to meet own goals, preferences. This requires more than heuristics. It requires more than assimilation (reframe information in terms of what you already know) and accommodation (revise what you know to account for new information). It requires that representations be iteratively restructured by re-viewing them from different, relevant perspectives until they achieve a form that solves problems. This in turn requires integrated, self-modifying conceptual structure. In the talk I discuss how such a problem-solving structure comes about.

3.7 An Analytic Tableaux Model for Deductive Mastermind Empirically Tested with a Massively Used Online Learning System

Nina Gierasimczuk (University of Amsterdam, NL)

The paper is concerned with the psychological relevance of a logical model for deductive reasoning. We propose a new way to analyze logical reasoning in a deductive version of the Mastermind game implemented within a popular Dutch online educational learning system (Math Garden). Our main goal is to derive predictions about the difficulty of Deductive Mastermind tasks. By means of a logical analysis we derive the number of steps needed for solving these tasks (a proxy for working memory load). Our model is based on the analytic tableaux method, known from proof theory. We associate the difficulty of Deductive

Mastermind game-items with the size of the corresponding logical trees obtained by the tableaux method. We derive empirical hypotheses from this model. A large group of students (over 37 thousand children, 5–12 years of age) played the Deductive Mastermind game, which gave empirical difficulty ratings of all 321 game-items. The results show that our logical approach predicts these item ratings well, which supports the psychological relevance of our model.

3.8 Real-World Problem-Solving: Back to the Future

Vinod Goel (York University – Toronto, CA)

The machinery of Newell and Simon's information processing theory has been the basis cognitive theory for the past 50 years. It has been nurtured on toy or well structured problems and plagued by doubt about its ability to scale up to explain real-world problems. I introduce some data from the neuropsychology literature that reinforces the distinction between toy problems and real world problems and further sheds doubt on the ability of information processing theory to explain significant aspects of human cognition.

3.9 Coordinated Reasoning

Justine Jacot (Lund University, SE)

Logical reasoning tasks have focused on testing the way lay people use some logical skills – be it an innate ability or an acquired competence – often without questioning the nature of the reasoning itself, taking for granted that the 'logic' used is classical logic and that reasoning equates with being rational. A failure to complete reasoning tasks is thus either interpreted as a failure of logical skills or a failure of rationality. However, if one agrees with Stenning and Van Lambalgen (2008) that people must reason 'to' an interpretation before reasoning 'from' an interpretation when performing a reasoning task, coordination on the meaning of instructions between experimenters and subjects is not only a precondition to drawing conclusions about reasoning (logical or otherwise), but also a requirement all along the performing of the tasks. Moreover, when those tasks are formulated in natural language, pragmatic constraints apply, identified as 'intuitions' of relevance by Relevance Theory (Sperber & Wilson 1995), which takes them as an argument that experimental tasks do not say much about the theory behind the reasoning competence of subjects. Although I agree with Relevance Theory on the fact that (un)successful performance in reasoning tasks is compatible with different theories of reasoning, the explanation they provide in terms of intuitions of relevance does not help much in understanding how people reason, nor does it allow to understand how experimenters assess the (lack of) success of subjects. I propose a simple model of experimental reasoning tasks in terms of signaling games, which highlights the need for coordination between experimenter and subject throughout the task, as well as the difficulty for experimenter to evaluate properly the performance of subject.

3.10 Abductive Reasoning Using Random Examples

Brendan Juba (Harvard University, US)

In 'abductive reasoning,' a known or desired condition is given, and one seeks to find plausible (but not necessarily entailed) premises that imply the given condition. Several key cognitive processes, notably including (sub)goal formation, can be viewed as instances of abductive reasoning problems. I propose a new model of abductive reasoning as searching for a conditional distribution using random examples, in the same spirit as Valiant's PAC-learning model. In contrast to previous probabilistic models of abductive reasoning, this model does not rely on a prior distribution over the possible conditions to indicate which conditions are more plausible.

Much like PAC-learning, this simple model is well suited to the analysis of algorithms. As a consequence, the model suggests an interesting picture of which representations can be found efficiently, and which are computationally intractable. In particular, k-DNF representations (for small k) can be found efficiently in this model, but recent results suggest that an efficient algorithm for finding conjunctive representations (or any stronger representation) may not exist.

3.11 Prefrontal Function, Reasoning and Adaptive Behavior

Etienne Koechlin (ENP – Paris, FR)

The prefrontal cortex subserves executive control and decision-making, i. e. the coordination and selection of thoughts and actions in the service of adaptive behavior. We present a computational theory describing the evolution of the prefrontal cortex in humans as gradually adding new inferential Bayesian capabilities for dealing with a computationally intractable decision problem: exploring and learning new behavioral strategies vs. exploiting and adjusting previously learned ones through reinforcement learning. The theory clarifies the integration of model-free and model-based reinforcement learning through the notion of strategy creation. The theory also shows that counterfactual inferences in humans yield to the notion of hypothesis testing, a critical reasoning ability for approximating optimal adaptive processes and presumably endowing humans with a qualitative evolutionary advantage in adaptive behavior. We also present recent empirical data from behavioral and neuroimaging experiments conducted in our lab supporting the proposed theory.

3.12 Computing Probabilities and Probabilistic Computations – A Primer

Johan H. P. Kwisthout (Radboud University Nijmegen, NL)

In this presentation, I elaborate on the difference (complexity-wise) between 'computing probabilities' and 'probabilistic computations' and show how and why they are distinct. I show that NP-hard problems cannot be approximated both efficiently (in polynomial time) and effectively (with small probability of error) using such a probabilistic computation unless $NP \subseteq BPP$.

3.13 Complexity of Computations in Spiking Neural Networks

Wolfgang Maass (TU Graz, AT))

Spike-based computations promise to be substantially more power-efficient than traditional clocked processing schemes. However it turned out to be surprisingly difficult to design networks of spiking neurons that are able to carry out demanding computations. We present a new method for organizing such networks out of simple stereotypical network motifs in such a way that they can solve hard constraint satisfaction problems from the domains of planning / optimization and verification / logical inference. We use here noise as a computational resource, in spite of the fact that the timing of spikes (rather than just spike rates) are essential for the resulting computations. This new organization scheme is supported by a new theoretical understanding of spike-based stochastic computations. Surprisingly, one can identify in this context also a concrete computational advantage of spiking networks compared with traditional non-spiking stochastic neural networks (Boltzmann machines).

3.14 Cognitive Neuroscience: Foundations and Insights for Problem Solving

David Noelle (University of California – Merced, US)

This presentation provides a general introduction to the field of cognitive neuroscience, including its goals, its primary methods, and its major findings. The domain of 'problem solving' is discussed from a cognitive neuroscience perspective, focusing on cognitive control and the prefrontal cortex. A variety of 'resource bounds' on neural computation are introduced. Finally, opportunities for cross-fertilization between the fields of computational complexity, artificial intelligence, cognitive psychology, and cognitive neuroscience are suggested, using as an example some recent work involving how anatomical connection patterns in the prefrontal cortex naturally give rise to a mechanism that resembles 'pointers' in computer science, supporting novel compositional knowledge structures in the brain.

3.15 From Channel Capacity to Semantic Relevance: Accurate Retrieval as the Ultimate Boundary on Complex Cognition

Stellan Ohlsson (University of Illinois at Chicago, US)

Inspired by work on channel capacity in statistical information theory, the first wave of cognitive psychology emphasized limits on transient storage (short-term memory, working memory). Digital computers focused computer scientists on boundaries on long-term memory storage and cpu processing speed. Processing limits of these sorts put boundaries on the amount of heuristic search an intelligent agent could engage in, and hence which problems it can solve. To overcome this boundary, intelligent agents need large amounts of knowledge to focus search on a few promising search paths. This focus in turn implies that learning rate is the limiting factor: It takes 10 years or more to acquire the required knowledge base. This, in turn, focus attention on the retrieval of the knowledge that is needed at each moment in time. J. R. Anderson has argued that because the world is unpredictable, retrieval cannot, even, in principle, be 100% accurate. The boundary on cognition is thus the unpredictable nature of reality. In this talk, I'll take the focus on retrieval one step further, and argue that the ultimate boundary is the need for a massively parallel relevance calculation that is independent of experience. This limit might be ultimate in the sense that there is no way to overcome it.

3.16 Psychology of Problem Solving

Zygmunt Pizlo (Purdue University – West Lafayette, US)

The talk begins with a brief discussion of problem solving as a goal vs. a task to achieve some goal. Visual integration of a contour is an example of the latter. In the presence of noise in the retinal image, this problem is computationally intractable. But, if natural constraints are used, such as smoothness, proximity, convexity and closure, this problem becomes tractable. Specifically, it can be solved as a shortest path problem in $n \log n$ time using the cortical representation known as a log-polar mapping. The second part of the talk discusses a recent model of how humans produce near-optimal TSP tours using working memory that can hold only a few pieces of information at a time. This is an illustration of problem solving as a goal in itself. The new model is an elaboration of a multiresolution-multiscale pyramid algorithm whose computational architecture is based on the known anatomical and physiological characteristics of the human visual system. The computational complexity of the model is linear and its working memory stores at any given time information about 2-5 clusters. The model has visual attention that is moved the way humans move their eyes. Performance of this model matches closely performance of human subjects.

3.17 Problem Solving: Representing the (right) Computational Problem

Ulrike Stege (University of Victoria, CA)

When problem solving computational optimization problems with the goal to understand how good human performance is when solving instances of the problem or understanding their problem solving strategy in general, it can be crucial to describe the task in a way that is easy to understand and that does not hinder human performance.We give examples of different computation problems that are – in a sense – equivalent. We further discuss different ways of parameterized problems using equivalent problem definitions. This can be particularly informative for NP-hard problems as different parameterization can result in different parameterized complexity (FPT versus W[1]). We then discuss kernelization, a powerful technique for fixed-parameter tractable problems. In the second part of the talk (joint work with Balasubramanian, Muller and Srinivasan) we discuss the problem of reverse engineering definition of computational problems when the problem solving strategy is known. Using the example of the greedy technique, we discuss how tweaking the constraints or objective function of an optimization problem can yield quality or optimum solution by introducing structure to the problem. We further explain the analogy of algorithm w.r.t. to the quality of their solution with policies employed, e. g., in autonomic systems.

3.18 Towards a Theory of Learning Problem-solving Skills

Niels A. Taatgen (University of Groningen, NL)

A general theory of human problem solving that everyone starts out with a general set of problem solving algorithms that are specialized for specific tasks when necessary. This idea originated from Newell and Simon, and has also been instantiated by Anderson in ACT* and later ACT-R. However, there is a problem with this theory: it cannot deal with the fact that people are generally able to deal with problems that are computationally intractable. Therefore, the general problem solving algorithms cannot 'solve' them, and therefore cannot lead to appropriate specialized skills. The solution is to assume that humans generally build up general problem solving skills while solving specific problems. These general skills are adapted to the specific needs of an individual, and can therefore, after longer periods of time, lead to the ability to solve increasingly complex instances of a particular intractable problem, with possible transfer to other intractable problems. The new PRIMs cognitive architecture (Taatgen, 2013) has been developed to learn new general skills, including problem-solving skills. Although it has not been applied explicitly to problem solving yet, it has been used for mathematics, working memory, control and theory of mind.

3.19 Complexity, Attention, Learning and Goldilocks

John K. Tsotsos (York University – Toronto, CA)

This presentation begins by overviewing results aimed at illuminating the complexity complexity of visual processing. These results are then used to steer development of a theory of visual attention, Selective Tuning, to ensure a tractable formulation. The theory has made many predictions of how the human brain attends to visual stimuli, which many years after, are now strongly supported by experiment. One remaining problem, of many, was highlighted, namely, how to learn the parameters of the model. A hypothesis was presented on this, specifically that we consider how each mechanism within the theory might be improved via training, thus leading to overall model improvements. A pilot experiment was shown to confirm that human visual attention may be trained in such a manner. The tie to the complexity reduction of each mechanism via training was presented. Finally, the overall optimality of the process was discussed, with the conclusion that the brain is not optimal. Rather, it exhibits Goldilocks intelligence – it behaves in 'just the right manner' for organisms to be successful.

3.20 Higher-order theory of mind in mixed-motive settings: An agent-based simulation study

Harmen de Weerd (University of Groningen, NL)

In daily life, people often encounter situations in which they have little to no data on how others will react. To solve this problem, people use theory of mind by reasoning about the beliefs, goals and intentions of others. This allows them to predict the behaviour of others in new situations, as well as interpret observed behaviour. People can also use this ability recursively: they use higher-order theory of mind to reason about the theory of mind abilities of others, as in 'he thinks that I don't know that he sent me an anonymous letter'. However, humans appear to be the only species to do so. In this poster, we investigate whether exposure to mixed-motive situations, which involve both cooperative and competitive elements, can explain the emergence of human-like theory of mind abilities. We focus on repeated one-shot negotiations within the setting of the Colored Trails game. Simulation results show that in this setting, both first-order and second-order theory of mind agents outperform agents that are more limited in their theory of mind abilities. Moreover, the presence of such first-order and second-order theory of mind agents also benefits other agents. Agents did not obtain any additional benefit from the ability to make use of third-order theory of mind. However, in contrast to findings in strictly competitive settings, we find that in negotiations, agents do benefit from fourth-order theory of mind.

4 Overview of Posters

In order to provide an alternative forum to talks for presentation of research results and interests, posters were solicited prior to the workshop. These posters were on display from the second through fourth full days of the meeting, and several periods of time were allocated in the schedule for presenters to stand by their posters and answer questions.

In addition, a cash prize of $200 was donated by Purdue University Press, publisher of the *Journal of Problem Solving*, for the best poster with a student as first author. A team of four judges, one from each of the four disciplines represented at the workshop (John Tsotsos (CS), Stellan Ohlsson (PY), Reineke Verbrugge (AI), Vinod Goel (CN)), made their assessments over the course of the poster display. The prize was awarded to Harmen de Weerd on the final evening of the workshop by Zygmunt Pizlo.

Several of the posters described above were also presented as talks. To avoid duplication of material, though the author and title are given below, a reference is made to the appropriate subsection of Section 3 of this report for the associated abstract.

4.1 Where is the chocolate? Modeling Development of Reasoning about False Beliefs of Others

Burcu Arslan (University of Groningen, NL)

(See abstract in Section 3.1.)

4.2 Formally Making Sure Your AI Gets the Job Done (Good Enough)

Tarek R. Besold (Universität Osnabrück, DE)

The recognition that human minds/brains are finite systems with limited resources for computation has led researchers in cognitive science to advance the Tractable Cognition thesis: Human cognitive capacities are constrained by computational tractability. As also human-level AI in its attempt to recreate intelligence and capacities inspired by the human mind is dealing with finite systems, transferring this thesis and adapting it accordingly may give rise to insights that can help in progressing towards meeting the classical goal of AI in creating machines equipped with capacities rivaling human intelligence. Therefore, we develop the "Tractable Artificial and General Intelligence Thesis" and corresponding formal models usable for guiding the development of cognitive systems and models by applying notions from parameterized complexity theory and hardness of approximation to a general AI framework. This poster an overview of work putting special emphasis on connections and correspondences to the heuristics framework as recent development within cognitive science and cognitive psychology.

4.3 The Challenge of Optimization

Sarah Carruthers (University of Victoria, CA)

(See abstract in Section 3.3.)

4.4 Ifs and Ands and Ors

Nicole Cruz de Echeverria Loebell (University of London, GB)

(See abstract in Section 3.5.)

4.5 Computational Evidence that Self-regulation of Creativity is Good for Society

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

Excess individual creativity can be detrimental to society because creators invest in unproven ideas at the expense of propagating proven ones. Moreover, a proportion of individuals can benefit from creativity without being creative themselves by copying creators. We hypothesized that (1) societies increase their rate of cultural evolution by tempering the novelty-generating effects of creativity with the novelty-preserving effects of imitation, and (2) this is carried out by selectively rewarding and punishing creativity according to the value of the individuals' creative outputs. We tested this using an agent-based model of cultural evolution in which each agent self-regulated its invention-to-imitation ratio as a function of the fitness of its cultural outputs. In self-regulating societies, agents segregated into creators and imitators. The mean fitness of cultural outputs was higher than in non-self-regulating societies, and changes in diversity were rapider and more pronounced. We discuss limitations and possible social implications of our findings.

4.6 Fast and Loose: A Bounded-rational Analysis of Wason's Selection Task

Emmanuel Genot (Lund University, SE)

This paper proposes a bounded-rational version of the Bayesian rational analysis of Wason's Selection Task (ST), which was called for by its proponents—who acknowledged that it presupposes computations with unrealistic computational cost—but never realized. We model ST as a sequential decision problem of information-seeking by questioning, based on a non-classical understanding of the conditional. Our assumptions generalize those of the rational analysis, and are supported by independent empirical arguments, pertaining to both ST and other reasoning tasks. Our analysis generalizes the Bayesian rational analysis: it

converges to the same selection under its special assumptions, but vindicates the selection under weaker assumptions, and without the need for Bayesian rationality. In conclusion, we discuss the extension of our analysis to the deontic case.

4.7 Pearl Diving – Data for TSP Research from an Online Game

Alexandra Kirsch (Universität Tübingen, DE)

For a better understanding of how humans solve instances of the Traveling Salesperson Problem (TSP) we have implemented an online game, in which player shave to solve different variants of TSPs. The game contains features of social games that are now common on social network platforms, with an appealing colorful design, an introduction level to introduce the goal of the game, and purchasable game advantages. The levels are structured in blocks of eight levels, each containing a slightly different variant of the TSP. In previous laboratory studies of human TSP solving, our players get feedback and can repeat levels until they find the optimal solution. Up to August 2014, more than 50people have produced data of more than 2000 tours. The regularly updated dataset is available online for all interested researchers:http://hci.uni-tuebingen.de/DatasetsThe game itself can be played at www.perlentaucher.medieninformatik.uni-tuebingen.de

4.8 Resource-Rational Decision-Making

Falk Lieder (University of California – Berkeley, US)

Extreme events come to mind very easily and people overestimate their probability and overweight them in decision-making.In this paper we show that rational use of limited cognitive resources can generate these ‘availability biases’. We hypothesize that availability helps people to quickly make good decisions in very risky situations. Our analysis shows that agents who decide by simulating a finite number of possible outcomes (sampling) should over-sample outcomes with extreme utility. We derive a cognitive strategy with connections to fast-and-frugal heuristics, and we experimentally confirm its prediction that an event’s extremity increases the factor by which people overestimate its frequency. Our model also explains three context effects in decision-making under risk: framing effects, the Allai’s paradox, and preference reversals.

4.9 Indirection and Symbol-like Processing in the Prefrontal Cortex and Basal Ganglia

David Noelle (University of California – Merced, US)

The ability to flexibly, rapidly, and accurately perform novel tasks is a hallmark of human behavior. In our everyday lives we are often faced with arbitrary instructions that we must understand and follow, and we are able to do so with remarkable ease. It has frequently been argued that this ability relies on symbol processing, which depends critically on the ability to represent variables and bind them to arbitrary values. Whereas symbol processing is a fundamental feature of all computer systems, it remains a mystery whether and how this ability is carried out by the brain. Here, we provide an example of how the structure and functioning of the pre-frontal cortex/basal ganglia working memory system can support variable binding, through a form of indirection(akin to a pointer in computer science). We show how indirection enables the system to flexibly generalize its behavior substantially beyond its direct experience (i. e., systematicity). We argue that this provides a biologically plausible mechanism that approximates a key component of symbol processing,exhibiting both the flexibility, but also some of the limitations, that area associated with this ability in humans.

4.10 Predictive Processing: A Standard Deviation from the Real World

Maria Otworowska (Radboud University Nijmegen, NL)

The question how a brain may give rise to cognition, such as the ability to form beliefs about the world and act upon them, has been a subject of decades of interdisciplinary research. Currently, a leading theory about how a brain may just do this is Predictive Processing theory. However in its current version, Predictive Processing is based on a Laplace assumption – it means that causes of sensory input are assumed to be represented probabilistically (so-called: recognition density) as an unimodal probability density function. We have shown that Laplace assumption violates expectations about typical human search behavior. Predictive Processing theory could be adjusted to explain such behavior, but this requires multimodal distributions, violating the Laplace assumption. This raises important questions about the possible (neural) representation and approximation of such multimodal distributions within a Predictive Processing framework.

4.11 Assessing the Computational Adequacy of the General Problem Solver Model

Zahra Sajedinia (Memorial University of Newfoundland, CA)

Problem solving is a core cognitive ability. Human performance varies widely in solving different types of problems. Ideally, cognitive models of problem solving should explain these variations in two ways: (1) the model should reproduce the sequences of actions applied by humans during problem solving (empirical adequacy), and (2) the time required by the model should match that required by humans, i. e., the model should be fast (slow) when humans are fast (slow) (computational adequacy). The former can be assessed by traditional psychological experiments; however, the latter requires the application of techniques from computational complexity theory. In this poster, we describe the first formal assessment of the computational adequacy of Newell and Simon's General Problem Solver model. We also discuss how our results can be used in both designing new psychological experiments and refining models of human problem solving.

4.12 Higher-order Theory of Mind in Mixed-motive Settings: An Agent-based Simulation Study

Harmen de Weerd (University of Groningen, NL)

(See abstract in Section 3.20.)

5 Working Groups and Discussions

Early in the planning for this workshop, the co-organizers realized that it would be difficult to specify topics for working groups given the diversity of disciplines represented. To mitigate this, we adopted the Bird-of-a-Feather (B.O.F.) mechanism employed at large conferences to allow people with similar interests to find each other in an unsupervised manner. A bulletin board was designated for the duration of the meeting as the B.O.F. 'clearing house'. Anyone could post a suggested B.O.F. topic along with a meeting time and venue on the board, and anyone with an interest in that topic was free to attend and contribute. Our only requirement was that the organizer of each B.O.F. provide a brief oral report summarizing the discussion and prepare a written report for inclusion in the seminar report. This mechanism was very successful, yielding 9 B.O.F. sessions. A subset of the written reports from these sessions is included below.

5.1 New Brain-Resource Inspired Computational Formalisms

Organizers: I. van Rooij, J. Kwisthout.

Participants: S. Carruthers, Ch. Cherniak, M. Fellows, N. Gierasimczuk, F. Jaekel, B. Juba, A. Kolokolova, F. Lieder, W. Maass, D. Noelle, F. Rosamond, C. Rothkopf, Z. Sajedinia, U. Stege, J. Tsotsos, S. Varma, T. Wareham, S. Watson.

One of the expected outcomes of this workshop was 'suggestions for new theories of computational complexity that better address the needs of researchers in psychology and cognitive science.' A sub-group of participants discussed this expected outcome. It was agreed that in order to be useful for the above-mentioned community of researchers, the abstraction level of such new theories should be 'just right.' It should be sufficiently abstract to generalize over concrete experiments or implementations, yet specific enough to be able to capture the key brain resources and constraints on them without generalizing them away. From the complexity-theoretic point of view, it was emphasized that the formalisms must discriminative, i. e., that important resources should not be 'lost in the abstraction.'

An important outcome of the discussion was agreement that 'networks of spiking neurons', as proposed by W. Maass and colleagues, are a promising starting point for future investigations. One of the foundational problems to be solved is how to translate the defining aspects of traditional computation devices into this neurally plausible model. For instance, in a Turing Machine the defining aspects are the input (the tape with symbols), the state transitions, the resources (e. g., time, space, non-determinism, randomness), the acceptance criteria, and the notion of reductions. These aspects must be mapped to their relevant counterparts in networks of spiking neurons, where the resources are defined as the convergence speed to the stationary distribution of the network, the number of samples from the stationary distribution that are needed to give the target output with a certain accuracy, the number or maximum frequency of spikes, the number of neurons, the number of connections, and so on.

We believe that researchers from various disciplines and with various backgrounds will be interested in this research program. Complexity theorists in computer science and artificial neural network researchers in cognitive and neuroscience can join forces to investigate the foundational problem identified above, with neuroscientists involved to make sure the formal models capture the relevant aspects of neuronal computations. The complexity theorists might then continue, for example, to prove class inclusion or separation, establish hierarchies, identify relations with more traditional complexity classes, and so on. Computational modelers in cognitive science and neuroscience might use the formalism and proven results to show the tractability and intractability of competing theories, compare empirical data with model predictions, and generate testable hypotheses about particular brain computations.

5.2 Cognitive Architectures and Problem Solving

Organizer: T. Buwalda.

Participants: B. Arslan, T. Besold, M. Blokpoel, V. Goel, F. Jaekel, N. Loebell, D. Noelle, S. Ohlsson, N. Taatgen, J. Tsotsos, S. Varma, R. Verbrugge.

This discussion focused on the relation between problem solving and cognitive architectures. We collectively articulated the history of problem solving and cognitive architectures, inventoried the role problem solving is currently playing in cognitive architectures, and speculated about the role problem solving will play in the future development of cognitive architectures. The discussion of the historical relation focused on George Miller, who coined the definition of cognitive architectures (but not the term itself) and on Newell and Simon, who created the General Problem Solver and imported production rules from the theory of computation into cognitive science. The discussion of the contemporary relation focused on production system architectures such as Soar and ACT-R, which put problem solving first, and on Bayesian approaches, which emphasize statistical mechanisms. These architecture provide not just different views of the fundamental computational mechanisms of cognition, but also – when are applied to different domains – yield different theories and models of cognition. One existing architecture that is strongly tied to problem solving and that continues to generate new insights is Soar. A promising new architecture that was created with problem solving as a focus area is Icarus.

A number of questions for future research were identified. One was whether it is more scientifically productive to have an architecture that is relatively specialized for one domain, or to have an architecture that is more general, but that underspecifies (i. e., leaves more degrees of freedom to the builder of) models of specific domains. Although specialized architectures often support 'better' models of the domains for which they are specialized (e. g., problem solving), they make less sense if the value of architecture is in explaining all of cognition. On the other hand, general architectures and the underspecified models they support are often criticized as unfalsifiable, and thus outside of the science. We agreed that the underspecificity and unfalsifiability critiques do not hold when comparing models of the same domain (i. e., problem solving) that are implemented in the same architecture. In that case, model evaluation is possible, and the best model can be retained and the other models rejected.

We discussed the more general question of how to choose the 'best' architecture among competing alternatives based on empirical support. One viewpoint, inspired by Lakatos' philosophy of science and endorsed by scientists such as Newell, Anderson, and Cooper is that architectures can be assessed by their productivity – how many models of how many domains they support, and whether these models 'progress' by generating new empirical predictions and growing to account for new data (However, we noted that major updates or changes to architectures can make their past productivity irrelevant). Another viewpoint is that it might not be necessary to pick the 'best' architecture. Instead, all architectures contribute to progress in cognitive science, with newer architectures 'standing on the shoulders' of older architectures.

With respect to the future, it was observed that there is currently little appetite in psychology for constructing computational models; experimental research reigns supreme. By contrast, in neuroscience, modeling is more appreciated, perhaps because neuroscientists deal with more complex phenomena that generate orders of magnitude more data. This leads to a better understanding of the utility of computational models for making scientific progress. As evidence of this fact, cognitive architecture is flourishing in the interstices between artificial intelligence, cognitive science, and cognitive neuroscience, as evidenced by emerging fields and conferences such as Biologically Inspired Cognitive Architectures (BICA), Artificial General Intelligence (AGI), and Neural-Symbolic Learning and Reasoning (NeSy).

The discussion ended with the mention of two more challenges for the future: (1) building models of problem solving that are constrained by, and can explain, neuroscientific data, and (2) having a reference set of problem solving data against which to evaluate architecture-based models of problem solving.

5.3 Should 'as if' Explanations of Inferential Behavior be Resource-Bounded?

Organizers: E. Genot, J. Jacot.

Participants: N. Fleischhut, N. Loebell, M. Otworowska, S. Varma.

Theoretical frameworks such as Bayesian learning theory, Bayesian decision theory, and game theory can be fruitfully applied to describe and understand inferential behavior. It is common practice in philosophy and economics – and sometimes in cognitive science and psychology – to do so for functional-level theories (e. g., with Bayesian rational analysis models). However, people – the agents whose behavior is to be modeled – are not rational in the strong normative sense that these frameworks assume. Philosophers and economists generally ignore this problem. Cognitive scientists and psychologists do so as well by justifying their models as 'as if' explanations, shifting the burden of accommodating bounds on cognition to lower algorithmic-level theories (in the sense of Marr). This shifting strategy is usually accompanied by hand waving and the invocation of 'approximations' and 'heuristics' as cure-alls for intractability. We agreed that this strategy does not stand close scrutiny, and that an alternative must be found.

The main question we discussed was whether functional-level theories should incorporate from the start constraints on the cognitive resources that are assumed to be available, and not defer such constraints to algorithmic-level theories. We discussed as an example the rational analysis approach, which assumes the 'principle of rationality': that organisms respond optimally to their environment (and thus behave as optimal statistical learners and expected-utility maximizers). If instead organisms can be assumed to respond merely 'well enough,' and if a suitable interpretation of 'well enough' can be given that assumes only bounded resources, then 'as if' functional-level explanations can share the load with algorithmic-level theories. However, this imposes a trade-off and raises an issue of its own. The trade-off is that what constitutes a 'well enough' response to a task is highly context-dependent, and for this reason broad-spectrum theories cannot serve as unifying frameworks without local adjustments. This issue is that what constitutes an acceptable adjustment must be made on a context-by-context basis, and this raises the risk of over-fitting.

We agreed both on the usefulness of 'as if' explanations, and on the necessity to adjust them. Although this topic is broad by nature (as is the spectrum of the theoretical frameworks we discussed), we also discussed concrete cases and applications, such as applications of rational analysis to data on reasoning tasks (e. g., variants of the Wason Selection Task). We agree that 'as if' explanations should incorporate linguistic pragmatics for understanding subjects' responses. In particular, Gricean pragmatics and Relevance Theory incorporate constraints on cognitive resources (such as attention and memory) in communication contexts while taking into account linguistic competence in the interpretation of instructions. Provided that such models can be captured formally, they can yield empirical predictions, by suggesting which parameters should be manipulated experimentally. The discussion concluded on which tools would be appropriate for this formalization.

5.4 Formal Representations of Real World Problem Solving / Complexity of Representing Real World Problems

Organizers: J. Kwisthout, F. Lieder.

Participants: [not recorded]

This meeting addressed two closely related questions: (1) what are suitable formal representations for real world problems, and (2) how can we analyze and characterize the complexity of representing real world problems by human problem solvers? A number of interesting ideas from cognitive science, artificial intelligence, and complexity analysis were brought to bear to understand these questions.

To take one example, interesting links were established between the observation that human problem solvers faced with unfamiliar problems tend to 'probe' the environment (i. e., act to get information), and with the complexity-theoretic aspect of interactive proofs. In such proofs, the traditional notion of 'verification' in NP is extended to be an interactive protocol where the verifier and an oracle exchange information.

To take another example, an interesting observation was made that problem solving is rarely done in complete isolation: we interact both with the environment and with the problem itself. Thus, there is an inherent multi-agent aspect in problem solving. This raises the possibility that communication complexity is relevant for understanding problem solving.

6 Panel Discussions and New Research Directions

The importance of choosing a proper representation in problem solving was discussed extensively. Some new insight were (1) when humans solve problem(s) the problem(s) they are solving may not be what they think they are solving and (2) representation format affects human problem solving. It was realized that research on meta-mathematics (the formal study of problem solving strategies) can benefit psychology and neuroscience knowledge of human problem solving in various ways, e. g.,

- knowledge compilation is the cognitive science equivalent of proposing a lemma in Extended Frege, and
- change detection in evolving data streams in AI can be informed by cognitive neuroscience on learning and switching strategies in changing environment.

New links between memory retrieval and Markov chains on semantic networks were also discussed.

In the B.O.F. sessions (Section 5.2) some notable common issues were (1) there is no agreement upon criterion for success (and how to measure it); (2) making falsifiable claims/models of cognitive / brain computation is important; and (3) computation is the appropriate language for modeling natural systems. Moreover, when comparing 'different' cognitive architectures, one has to know what are the key assumptions and commitments of the different architectures.

Several insights emerged during the panel discussion on alternative resource-bounded computation in spiking neural networks – namely, computational level models need to be at the right level of abstraction: low enough for neuroscience (to relate to concrete data of brain and behavior; should have features that neuroscience consider critical); but also general enough to be suitable for formal analytical methods.

It was also frequently noted that off-the-shelf complexity theory is limited in how much it can directly inform cognitive (neuro)science. Many of the sub-disciplines represented at the seminar could benefit from such a more finely-tuned computational complexity. To address this, a proposal was made for a study on alternative resource based computation. Such a study should incorporate the observation that the structure of the environment is exploited during human problem solving: this came out both in psychological talks as well as in computer science talks on parameterized complexity.

Future possible direction for further study has been considered, just to mention some:

- Machine Learning and Problem Solving,
- Problem solving by nature,
- Problem solving in cognitive architectures,
- Problem representation / interpretation,
- Learning representations and algorithms for problem solving (machine learning; how do we learn the representations and algorithms).

7 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. 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. 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 the journal web site:

- http://docs.lib.purdue.edu/jps/
- Editor-in-Chief: Zygmunt Pizlo, Department of Psychological Sciences, Purdue University.

8 Seminar Program

See Section 3 for more information about the talks and the posters.

Monday, August the 18th

9:00	General Introduction	I. van Rooij
10:15	Intro: Psychology	Z. Pizlo
11:15	Intro: Computational Complexity	T. Wareham
13:30	Intro: AI	R. Verbrugge
14:30	Intro: Cognitive Neuroscience	D. Noelle
16:00	Panel and General Discussion	
19:15	Town Hall: Plans for the Week	

Tuesday, August the 19th

9:00	Talk	J. Tsotsos
9:30	Talk	W. Maass
10:15	Discussion:	Computational Complexity and Brain Computation
11:15	Poster session	
13:30	Talk	S. Ohlsson
14:00	Talk	M. Fellows
14:30	Talk	J. Kwisthout
15:00	Talk	T. Besold
16:00	Discussions	
19:15	Reporting + Panel I	

Wednesday, August the 20th

9:00	Talk	U. Stege
9:30	Talk	Ch. Cherniak
10:15	Talk	N. Taatgen
10:35	Talk	N. Loebell
10:55	Open problem session	
19:15	Panel II	

Thursday, August the 21st

9:00	Talk	N. Gierasimczuk
9:30	Talk	C. Rothkopf
10:30	Talk	B. Juba
10:55	Talk	E. Koechlin
11:20	Talk	V. Goel
11:45	Panel Discussion	
13:30	Talk	L. Gabora
13:55	Poster Session II	
14:50	Break out Session Talk:	B. Arslan / H. de Weerd / S. Carruthers
16:00	Discussions	
19:15	Poster Prize	
19:20	B.O.F Reports	
20:00	Town Hall Discussion	

Friday, August the 22nd

9:00	Seminar Summary

Participants

- Burcu Arslan
University of Groningen, NL
- Tarek R. Besold
Universität Osnabrück, DE
- Mark Blokpoel
Radboud Univ. Nijmegen, NL
- Sarah Carruthers
University of Victoria, CA
- Christopher Cherniak
University of Maryland – College Park, US
- Nicole Cruz de Echeverria Loebell
Birkbeck, Univ. of London, GB
- Harmen de Weerd
University of Groningen, NL
- Michael R. Fellows
Charles Darwin University – Darwin, AU
- Nadine Fleischhut
MPI für Bildungsforschung, DE
- Liane Gabora
University of British Columbia – Vancouver, CA
- Emmanuel Genot
Lund University, SE
- Nina Gierasimczuk
University of Amsterdam, NL
- Vinod Goel
York University – Toronto, CA
- Yll Haxhimusa
TU Wien, AT
- Justine Jacot
Lund University, SE
- Frank Jäkel
Universität Osnabrück, DE
- Brendan Juba
Harvard University, US
- Alexandra Kirsch
Universität Tübingen, DE
- Etienne Koechlin
ENP – Paris, FR
- Antonina Kolokolova
Memorial University of Newfoundland, CA
- Johan H. P. Kwisthout
Radboud Univ. Nijmegen, NL
- Falk Lieder
University of California – Berkeley, US
- Wolfgang Maass
TU Graz, AT
- Matthias Mnich
Universität Bonn, DE
- Martin Möhrmann
Universität Osnabrück, DE
- David Noelle
Univ. of California – Merced, US
- Stellan Ohlsson
Univ. of Illinois – Chicago, US
- Maria Otworowska
Radboud Univ. Nijmegen, NL
- Zygmunt Pizlo
Purdue University – West Lafayette, US
- Frances A. Rosamond
Charles Darwin University – Darwin, AU
- Constantin Rothkopf
TU Darmstadt, DE
- Zahra Sajedinia
Memorial University of Newfoundland, CA
- Ulrike Stege
University of Victoria, CA
- Marieke Sweers
Radboud Univ. Nijmegen, NL
- Niels A. Taatgen
University of Groningen, NL
- John K. Tsotsos
York University – Toronto, CA
- Iris van Rooij
Radboud Univ. Nijmegen, NL
- Sashank Varma
University of Minnesota – Minneapolis, US
- Rineke Verbrugge
University of Groningen, NL
- Todd Wareham
Memorial University of Newfoundland, CA
- Scott Watson
Memorial University of Newfoundland, CA