DAGSTUHL
REPORTS

Volume 1, Issue 3, March 2011

Reasoning about Interaction: From Game Theory to Logic and Back (Dagstuhl Seminar 11101)
Jürgen Dix, Wojtek Jamroga, and Dov Samet 1

Computational Geometry (Dagstuhl Seminar 11111)
Pankaj Kumar Agarwal, Kurt Mehlhorn, and Monique Teillaud 19

Computational Complexity of Discrete Problems (Dagstuhl Seminar 11121)
Martin Grohe, Michal Koucký, Rüdiger Reischuk, and Dieter van Melkebeek 42

Exploration and Curiosity in Robot Learning and Inference (Dagstuhl Seminar 11131)
Jeremy L. Wyatt, Peter Dayan, Ales Leonardis, and Jan Peters 67

Dagstuhl Reports, Vol. 1, Issue 3 ISSN 2192-5283

ISSN 2192-5283

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany.
Online available at http://www.dagstuhl.de/dagrep

Publication date
August, 2011

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at http://dnb.d-nb.de.

Aims and Scope
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:

- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e.g. summaries from panel discussions or open problem sessions.

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de

Digital Object Identifier: 10.4230/DagRep.1.3.i

Report from Dagstuhl Seminar 11101

Reasoning about Interaction: From Game Theory to Logic and Back

Edited by

Jürgen Dix[1], Wojtek Jamroga[2], and Dov Samet[3]

1 **TU Clausthal, DE,** dix@tu-clausthal.de
2 **University of Luxembourg, LU,** wojtek.jamroga@uni.lu
3 **Tel Aviv University, IL,** dovsamet@gmail.com

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11101 "Reasoning about Interaction: From Game Theory to Logic and Back".

The notion of interaction is crucial in several disciplines, including social science, operational research, and economics. Two frameworks are most prominent in the formal treatment of interaction: game theory and mathematical logic. Quantitative analysis is usually conducted using models and tools of game theory. At the same time, logic provides vocabulary and methods to study interaction in a qualitative way.

The aim of the seminar was to bring together researchers who approach interaction-related phenomena from different perspectives (and with different conceptual tools). We hoped that, by synergy and exchange of expertise, a more integrative view of interaction could be obtained. In particular, we focussed on how interaction between individual entities (be it humans, robots and/or virtual creatures) can lead to emergence of social structures, collective behavior, and teamwork - and, ultimately, help all involved parties benefit from cooperation.

Seminar 06.–11. March, 2011 – www.dagstuhl.de/11101

1998 ACM Subject Classification I.2 Artificial Intelligence, F.4.1 Mathematical Logic, I.2.4 Knowledge Representation Formalisms and Methods

Keywords and phrases Game Theory, Logic, Mechanism Design, Security, Cooperation, Model Checking, Rationality, Knowledge

Digital Object Identifier 10.4230/DagRep.1.3.1

Edited in cooperation with Nils Bulling

1 Executive Summary

Jürgen Dix
Wojtek Jamroga
Dov Samet

The group (48 participants from 13 countries) convened in Dagstuhl in March 2011, for a five day meeting.

The aim of the seminar was to bring together researchers who approach interaction-related phenomena from different perspectives (and with different conceptual tools). We also wanted to identify potentials for coordination, and to discuss general models and methodologies for future research.

Of particular importance was the choice of the participants and the areas they working in, namely: (1) classical game theory, (2) mathematical logic, and (3) economics. While there are some relations between these areas, we felt that more work should be done on the overlapping parts to make tools and methods from one area available in the others (if possible).

In particular, we wanted to find answers to the following questions:

- Are existing models of interaction adequate? Can models used by different disciplines be integrated in a meaningful way?
- How can we use game-theoretical concepts to construct logics that support strategic reasoning? What are the necessary features of such logics?
- How can epistemic-logic reasoning and definitions lead to the definitions of new solution concepts is strategic-form games?
- How can epistemic and strategic logic be adapted to the empirical findings from game theoretic laboratory experiments, manifesting bounded rationality of a variety of types?
- How can issues of computational complexity be addressed vis-a-vis the demand for efficiency/optimality in the design of economic mechanisms under asymmetric information?

The seminar resulted in making the first step towards answering these questions. We did not obtain the ultimate formal answers, especially in the sense of enabling implementation in the form of ready-to-use tools and methodologies. However, researchers with different background shared their views on how games and multi-agent systems can be modeled and reasoned about, which led to several discussions on fundamental questions (like: *what features/concepts are indispensable when analyzing interaction between agents?*). In particular, the issue of whether probabilities (and, more generally: quantities) are necessary to give good account on how agents interact was hotly debated.

The results of the seminar were somewhat constrained by the unbalanced composition of participants. We had invited equally many researchers from computer science (especially computational logic) and economics (game theory). However, while most computer scientists accepted our invitation, the same did only a few economists. This is probably due to the fact that Dagstuhl seminars have an extremely high reputation within computer science, but they are relatively unknown in other disciplines. In consequence, the synergy between different views of interaction occurred only partially. In our opinion, it was especially fruitful on the basic level. That is, economists and computer science logicians learned about the basic models and patterns of analysis used in the other discipline. Even more importantly, they exchanged views on what *research questions* are relevant and viable when analyzing game-like interaction. Most synergy occurred within the subgroup of participants coming from the community of *modal logic in computer science.* Talks on modal logic-related topic triggered intensive discussion and ideas for joint research which are currently being pursued by several participants.

We thank the Dagstuhl staff for a very fruitful and interesting week. We are planning a special issue (in Annals of Math and AI) as a concrete outcome of the seminar. Moreover, it was a general consensus that a follow-up seminar would be highly interesting – this time with more specific topics being the focus. The follow-up is currently in the planning phase.

2 Table of Contents

Executive Summary
Jürgen Dix, Wojtek Jamroga, Dov Samet 1

Overview of Talks

Modal Logic and Strategic Interaction in Multi-agent Systems (overview talk)
Thomas Agotnes 5

Belief revision in dynamic games
Giacomo Bonanno 5

Modal Logic for Reasoning in Game Situations
Jan M. Broersen 5

Programming normative mechanisms
Nils Bulling 6

Strategy synthesis for multi-agent systems
Jan Calta 6

Binary Aggregation
Ulle Endriss 7

On the Dynamics of Information and Abilities of Players in Multi-Player Games: a preliminary report
Valentin Goranko 7

A dynamic logic of normative systems
Andreas Herzig 8

Abstraction for Model Checking Modular Interpreted Systems over ATL
Michael Koester 8

Conditional Equilibrium Outcomes via Ascending Price Processes
Ron Lavi 9

Playing games in large multi-agent simulations
Viliam Lisy 9

From Individualistic to Social Rationality in Strategic Games: a Logical Analysis
Emiliano Lorini 10

Secure interaction and security protocols
Sjouke Mauw 10

Imperfect Information and Intention in Non-Repudiation Protocols
Matthijs Melissen 11

Simultaneous Ad Auctions
Dov Monderer 13

Mission planning: Logic and game theory in multi-robot applications
Peter Novak 13

Reasoning with Plans under Imperfect Information
Eric Pacuit 14

The Power of Knowledge in Games
Rohit Parikh 14

Neighbourhood structures in large games
Ramaswamy Ramanujam 14

Epistemic, Strategic ATL* with Explicit Strategies
Henning Schnoor 15

Efficiency Levels in Sequential Auctions with dynamic Arrivals
Ella Segev 15

Logics for social choice and perspectives on their software implementation
Nicolas Troquard 15

Turning Competition into Cooperation and Cooperation into Competition
Paolo Turrini 16

A Proof System for Message Passing
Jan van Eijck 16

Open Problems 16

Panel Discussions 17

Participants 18

3 Overview of Talks

3.1 Modal Logic and Strategic Interaction in Multi-agent Systems (overview talk)

Thomas Ågotnes (University of Bergen, NO)

A recent trend in logic for multi-agent systems and social mechanisms is logics for coalitional ability, interpreted in game-like structures. In this overview talk I give an overview of some of the most popular basic frameworks and extensions that have been proposed in order to increase their expressiveness.

3.2 Belief revision in dynamic games

Giacomo Bonanno (University of California – Davis, US)

URL http://www.econ.ucdavis.edu/faculty/bonanno/PDF/pbe_2.pdf

We investigate belief revision in dynamic (or extensive-form) games in accordance with the the postulates of the so-called AGM theory of belief revision (developed by Alchourrón, Gärdenfors and Makinson, 1985). We show that consistency with the AGM theory requires that the players' ex ante beliefs and disposition to change those beliefs be rationalizable by a total pre-order on the set of histories, which we call a plausibility order. When an information set is reached, the player's revised beliefs are given by the set of most plausible histories among the ones that constitute the information set; furthermore, the player chooses an action with positive probability if and only if that action maintains the plausibility of the history at which it is taken.

If we add the assumption that the players have a common prior; (that is, common initial beliefs and a common disposition to revise those beliefs) then we obtain a solution concept for dynamic games which is intermediate between subgame-perfect equilibrium and sequential equilibrium and captures the idea of applying Bayes; rule whenever possible; (on and off the equilibrium path(s)).

3.3 Modal Logic for Reasoning in Game Situations

Jan M. Broersen (Utrecht University, NL)

Main reference Modelling Attempt and Action Failure in stit Logic, Jan Broersen, Proceedings of Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011), 2011.

The solution concepts of game theory describe possible equilibria for agent interactions, under specific uniform assumptions about the rationality of these agents. Logics for games often take this same general 'outsider' perspective; they are used to characterize equilibria and aim at reproducing game theoretic results as logic theorems. In my talk I propose to take an 'insider' perspective to logic for games. Our aim could be to design logics that model the reasoning of agents involved in taking decisions in game situations. We should leave open the possibility that different agents reason according to different logics (decision rules). By means of an example I will show that this leads to fundamental questions about the modeling of opponent reasoning, the projection of reasoning into others, and the interaction of projected reasoning with an agent's 'native' reasoning. I will briefly discuss probabilistic stit logic as an example of a logic that takes the inside perspective of reasoning in game situations. As a corollary of the proposed view, I will suggest to look at notions of Nash-action and Pareto-action.

3.4 Programming normative mechanisms

Nils Bulling (TU Clausthal, DE)

Joint work of Dastani, Mehdi
Main reference Nils Bulling and Mehdi Dastani, Verification and Implementation of Normative Behaviours in Multi-Agent Systems, Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI), July, 2011 (to appear).

The environment is an essential component of multi-agent systems, which is often used to coordinate the behaviour of individual agents. Recently many programming languages have been proposed to facilitate the implementation of such environments. This paper is motivated by a programming language that is designed to implement environments in terms of normative concepts such as norms and sanctions. We provide a formal analysis of programmed normative environments based on concepts from mechanism design. By doing this we relate normative environments to implementation theory setting the stage for studying formal properties of normative environments such as whether a set of norms implements specific normative choice function in specific equilibria. This allows, for example, to analyse whether groups of agents are willing to obey the rules specified by a normative system.

3.5 Strategy synthesis for multi-agent systems

Jan Calta (HU Berlin, DE)

Joint work of Calta, Jan; Shkatov Dmitry; Schlingloff Holger
Main reference Jan Calta and Dmitry Shkatov and Holger Schlingloff, Finding Uniform Strategies for Multi-agent Systems, LNCS, pages 135–152, volume 6245
URL http://dx.doi.org/10.1007/978-3-642-14977-1_12

We can specify desired computations of a multi-agent system with temporal properties that these computations satisfy. An executable strategy for an agent is a function assigning an action of the agent to a state of the agent. Thus, a strategy for a coalition of agents in a multi-agent system can be seen as an algorithm that enforces only system computations satisfying given properties.

The executable strategies for a given property and a multi-agent system can be synthesized using an approach based on global model checking. We focus on the synthesis of an executable strategy in two settings. In the general setting, the multi-agent system is composed of agents with different abilities and divided into coalitions of agents and the properties of the system computations are expressed in Alternating-time Temporal Logic. In the restricted setting, the multi-agent system is homogenous and the properties of the system computations are expressed in Linear Temporal Logic. We discuss how the restrictions on the agent interaction and available information influence the synthesis of the strategies.

3.6 Binary Aggregation

Ulle Endriss (University of Amsterdam, NL)

Joint work of Grandi, Umberto; Endriss, Ulle;
Main reference U. Grandi and U. Endriss. Lifting Rationality Assumptions in Binary Aggregation. Proc. 24th AAAI Conference on Artificial Intelligence (AAAI-2010), 2010.
URL http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1592

Binary aggregation deals with situations where several individuals each make a yes/no choice regarding a number of issues and these choices then need to be aggregated into a collective choice. Depending on the application at hand, different combinations of yes/no may be considered rational. We can use an integrity constraint, modelled in terms of a formula of propositional logic, to define the set of those rational choices. Important frameworks of social choice theory, such as Arrovian preference aggregation and judgment aggregation, can be cast as binary aggregation problems, for specific choices of integrity constraints. In this talk, I presented some of our recent work on the interplay of the propositional language used to express integrity constraints and the axiomatic properties of aggregation procedures that can be guaranteed to respect those constraints, and I showed how this general perspective not only is helpful to understand binary aggregation itself, but also has applications in both preference and judgment aggregation.

This is joint work with Umberto Grandi (Amsterdam).

References

1 U. Grandi and U. Endriss.
Lifting Rationality Assumptions in Binary Aggregation.
Proc. 24th AAAI Conference on Artificial Intelligence (AAAI-2010), 2010.
2 U. Grandi and U. Endriss.
Binary Aggregation with Integrity Constraints.
Proc. 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011), To appear in 2011.

3.7 On the Dynamics of Information and Abilities of Players in Multi-Player Games: a preliminary report

Valentin Goranko (Technical University of Denmark, DK)

I will discuss first steps towards a more realistic treatment and logical formalization of the abilities of players to achieve objectives in multi-player games under incomplete, imperfect, or simply wrong information that they may have about the game and about the course of the play. In this talk, after some motivating examples I will introduce a variation of the multi-agent logic ATL as a logical framework for capturing the interplay between the dynamics of information and the dynamics of abilities of players. This framework takes into account both the a priori information of players with respect to the game structure and the empirical information that players develop over the course of an actual play. It associate with them respective information relations and notions of 'a priori' and 'empirical' strategies and strategic abilities. I will briefly discuss the problem of model checking of statements formalized in the new logic under different assumptions about the abilities of the players to observe, remember, and reason.

3.8 A dynamic logic of normative systems

Andreas Herzig (Université Paul Sabatier – Toulouse, FR)

Joint work of Herzig, Andreas;Lorini, Emiliano; Troquard, Nicolas
Main reference Andreas Herzig, Emiliano Lorini, Nicolas Troquard. A dynamic logic of normative systems. Proc. IJCAI 2011
URL http://www.irit.fr/~Andreas.Herzig/P/Ijcai11.html

We propose a logical framework to represent and reason about agent interactions in normative systems. Our starting point is a dynamic logic of propositional assignments whose satisfiability problem is PSPACE-complete. We show that it embeds Coalition Logic of Propositional Control CL-PC and that various notions of ability and capability can be captured in it. We finally show how the logic can be extended in order to represent constitutive rules which are also an essential component of the modelling of social reality.

3.9 Abstraction for Model Checking Modular Interpreted Systems over ATL

Michael Koester (TU Clausthal, DE)

I present an abstraction technique for model checking multi-agent systems given as modular interpreted systems (MIS). MIS allow for succinct representations of compositional systems, they permit agents to be removed, added or replaced and they are modular by facilitating control over the amount of interaction.

Specifications are given as arbitrary ATL formulae: One can therefore reason about strategic abilities of groups of agents. The technique is based on collapsing each agent's local state space with handcrafted equivalence relations, one per strategic modality. Performing model checking on abstractions (which are much smaller in size) rather than on the concrete system which is usually too complex saves space and time.

3.10 Conditional Equilibrium Outcomes via Ascending Price Processes

Ron Lavi (Technion – Haifa, IL)

Joint work of Fu, Hu; Kleinberg, Robert D.; Lavi, Ron
URL http://ie.technion.ac.il/~ronlavi/papers/CE.pdf

A Walrasian equilibrium in an economy with non-identical indivisible items exists only for small classes of players' valuations (mostly "gross substitutes" valuations), and may not generally exist even with decreasing marginal values. This paper studies a relaxed notion, "conditional equilibrium", that requires individual rationality and "outward stability", i.e., a player will not want to *add* items to her allocation, at given prices. While a Walrasian equilibrium outcome is unconditionally stable, a conditional equilibrium outcome is stable if players cannot choose to drop only *some* of their allocated items.

With decreasing marginal valuations, conditional equilibrium outcomes exhibit three appealing properties: (1) An approximate version of the first welfare theorem, namely that the social welfare in any conditional equilibrium is at least half of the maximal welfare; (2) A conditional equilibrium outcome can always be obtained via a natural ascending-price process; and (3) The second welfare theorem holds: any welfare maximizing allocation is supported by a conditional equilibrium. In particular, each of the last two properties independently implies that a conditional equilibrium always exists with decreasing marginal valuations (whereas a Walrasian equilibrium generally does not exist for this common valuation class).

Given these appealing properties we ask what is a maximal valuation class that ensures the existence of a conditional equilibrium and includes unit-demand valuations. Our main technical results provide upper and lower bounds on such a class. The lower bound shows that there exists such a class that is significantly larger than gross-substitutes, and that even allows for some (limited) mixture of substitutes and complements. For three items or less our bounds are tight, implying that we completely identify the unique such class. The existence proofs are constructive, and use a "flexible-ascent" auction that is based on algorithms previously suggested for "fractionally subadditive" valuations. This auction is slightly different from standard ascending auctions, as players may also decrease prices of obtained items in every iteration, as long as their overall price strictly increases.

3.11 Playing games in large multi-agent simulations

Viliam Lisy (Czech Technical University, CZ)

Joint work of Lisy, Viliam; Bosansky, Branislav; Jakob, Michal; Pechoucek, Michal
Main reference Viliam Lisy and Branislav Bosansky and Michal Jakob and Michal Pechoucek: Adversarial Search with Procedural Knowledge Heuristic. In The Eighth International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2009

We investigate creating plans (sequences of actions) for agents acting in a shared environment towards achieving their own, often mutually exclusive goals. We focus on using our techniques in simulations of various kinds of (military) operations in physical world.

I will identify the problems we face in creating rational agents for those scenarios. The problems include limited modeling power of the available approaches and computational demands of the existing algorithms.

In order to overcome these problems, it is necessary to rely on a significant amount of hand-coded (or automatically learned) domain specific knowledge. This fact has been recognized in the field of planning and led to invention of HTN or TLPlan. I will present our method for using domain specific procedural knowledge in game playing.

3.12 From Individualistic to Social Rationality in Strategic Games: a Logical Analysis

Emiliano Lorini (Université Paul Sabatier – Toulouse, FR)

I propose a modal logic that enables to reason about different kinds of rationality in strategic games. This logic integrates the concepts of joint action, belief, individual preference and group preference. The first part of the talk is focused on the notion of individualistic rationality assumed in classical game theory: an agent decides to perform a certain action only if the agent believes that this action is a best response to what he expects the others will do. The second part of the talk explores different kinds of social rationality such as fairness and reciprocity. Differently from individualistically rational agents (alias self-interested agents), social rational agents also consider the benefits of their choice for the group.

Moreover, their decisions can be affected by their beliefs about other agents' willingness to act for the well-being of the group. The analysis also considers team-directed reasoning, i.e. the mode of reasoning that people use when they take themselves to be acting as members of a group or a team

3.13 Secure interaction and security protocols

Sjouke Mauw (University of Luxembourg, LU)

Quoting the late Roger Needham: "Security protocols are three-line programs that people still manage to get wrong." Given the famous history of the Needham-Schroeder protocol, Roger Needham clearly knew what he was talking about.

In this overview talk, I will sketch some of the pitfalls in security protocol design by showing examples from three areas: secrecy protocols, authentication protocols, and non-repudiation protocols.

3.14 Imperfect Information and Intention in Non-Repudiation Protocols

Matthijs Melissen (University of Luxembourg, LU)

Introduction

Repudiation [1] means the *denial* of an entity of having participated in all or part of a communication. When Alice sends a message m to Bob, we can distinguish *non-repudiation of origin* (NRO), which expresses that Alice cannot deny having sent m, and *non-repudiation of receipt* (NRR), which expresses that Bob cannot deny having received m from Alice. Proof of the origin of a message or receipt is provided by digital signatures obtained through a *public key infrastructure*, part of which is a *certificate authority* which links public keys to user identities. An important difference between non-repudiation protocols and most other security protocols is that in non-repudiation protocols, the information that the agents have acquired at the end of the protocol run should not only convince the agent itself, but also serve as a proof towards external agents (such as a judge). Applications of non-repudiation include contract signing and certified e-mail [2]. In this abstract, we show why it is important to take (imperfect) information and intention into account when modelling non-repudiation.

Fair exchange and imperfect information

It is often desirable to have a guarantee of fair exchange [3] of non-repudiation. For example, when Alice sends message m to Bob, it should hold that Alice receives her NRR if and only if Bob receives his NRO. One way to formally verify such constraints is by means of ATL [4], a modal logic of strategic ability. In ATL, the formula $\langle\langle A\rangle\rangle\Diamond\varphi$ stands for "Group of agents A has the ability to make sure that at some point in the future φ holds". Kremer & Raskin [1] use the following formula to express one of the conditions of fairness:

$$\Phi \equiv \neg\langle\langle Bob\rangle\rangle\Diamond(\mathsf{NRO} \wedge \neg\langle\langle Alice\rangle\rangle\Diamond\mathsf{NRR})$$

This formula is supposed to express that Bob should not be able to obtain his NRO at some point while at the same time making Alice unable to obtain her NRR. However, in [1], Φ is interpreted in the basic version of ATL which implicitly assumes perfect information.

That is, agents are assumed to know precisely the current global state of the system, including the local states of the other agents.

Obviously, the assumption is wrong for communication protocols in general (if everybody knows everything, no communication is needed).

Moreover, in the specific case of non-repudiation protocols, it can be the case that – even if Alice has a way of behaving that causes her to obtain her NRR – the behaviour cannot be represented as an executable strategy because it requires executing different actions in situations that look the same to her.

And even if she has an executable strategy, she may be unaware of having it, and unable to identify it [5]. For example, one can construct a protocol in which Alice needs to send a different message depending on whether Bob did or did not receive some other message. Alice is not aware of messages being received by Bob, so although she has a perfect information strategy, she is not able to follow it.

Conversely, it is also possible that Bob has a perfect information strategy to put Alice at a disadvantage, but the strategy cannot be executed under imperfect information. Thus, it can both happen that a protocol satisfying Φ is clearly unfair, and that a protocol is intuitively fair while satisfying $\neg\Phi$. This problem needs to be addressed by interpreting specifications in the appropriate version of ATL with imperfect information [6].

Intention

In many applications, the notion of non-repudiation is not sufficient. Recall that NRO means that the sender cannot deny having sent message m. However, often we require that the sender cannot deny having *intended to* sent m.

When Alice is presented with NRR signed by her, it is clear that she is the origin of the message. However, she can claim that she never executed the protocol in question, and that in fact NRR resulted from another protocol, in which m was used as a random string without meaning. For example, NRR could be interpreted in one protocol as “I am the originator of the message ‘I owe you money’”, while in a second protocol it is merely used as a confirmation that the participant is still on-line. In this situation, we know that one of the parties cheated: either Bob forced Alice to run the second protocol, or Alice lies about never having executed the first protocol. However, it is not possible to find out which of both parties is the culprit. We call an attack like this a *virtual multi-protocol attack*, as it is not necessary for Alice to actually execute any part of the second protocol; merely the existence of a protocol that could result in the same evidence is sufficient.

In general, for each NRR resulting from a non-repudiation protocol, it is possible to construct another protocol where the message NRR has a different interpretation. Note that adding flags to each message that indicate the purpose of the message does not work, as it is not possible to guarantee that flags are unique across protocols. Note furthermore that the same problem occurs for NRO.

The above problem can be solved by making sure that the certificate authority does not only link public keys and user identities, but additionally stores the exact protocol (and version) for which the key will be used with. Note also that this registration needs to be happening in the physical world, as an electronic registration leads to a bootstrapping problem: the registration authority needs to have non-repudiation of origin, because otherwise the agent whose key is registered could later deny having registered her key.

Conclusion

We have studied the security requirement of non-repudiation. Often, fair exchange of non-repudiation needs to be guaranteed. When modelling this property formally, it should be taken into account that agents in a protocol have imperfect information, and therefore a model that takes this into account should be used. Furthermore, we indicate that in many applications the notion of non-repudiation is not sufficient, as the parties are required to send their messages *intentionally*. This can be solved by letting the certificate authority store the exact protocol for which a key will be used.

Acknowledgements

This work was supported by the FNR (National Research Fund) Luxembourg under projects S-GAMES, C08/IS/03 and GMASec - PHD/09/082.

References

1 Kremer, S., Raskin, J.F.: A game-based verification of non-repudiation and fair exchange protocols. Journal of Computer Security **11** (2003)

2 Dashti, M.T.: Keeping Fairness Alive. PhD thesis, Vrije Universiteit (2008)
3 Ben-Or, M., Goldreich, O., Micali, S., Rivest, R.: A fair protocol for signing contracts. IEEE Transactions on Information Theory **IT-36** (1990) 40–46
4 Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. Journal of the ACM **49** (2002) 672–713
5 Jamroga, W., van der Hoek, W.: Agents that know how to play. Fundamenta Informaticae **63** (2004) 185–219
6 Schobbens, P.Y.: Alternating-time logic with imperfect recall. Electronic Notes in Theoretical Computer Science **85** (2004) 82–93

3.15 Simultaneous Ad Auctions

Dov Monderer (Technion – Haifa, IL)

Joint work of Ashlagi, Itai;Monderer, Dov; Tennenholtz Moshe
Main reference Ashlagi, Itai;Monderer, Dov; Tennenholtz Moshe, Simultaneous Ad Auctions, Mathematics of Operations Research, Vol. 36, No. 1, February 2011, pp. 1-13.
URL http://dx.doi.org/10.1287/moor.1100.0475

Two models for a pair of simultaneous ad auctions, A, B, are discussed:

(i) single-campaign advertisers, which participate in a single ad auction, and (ii) multi-campaign advertisers, which participate in both auctions. We prove the existence and uniqueness of a symmetric equilibrium in the first model. Moreover, when the click rates in A are point-wise higher than those in B, we prove that the expected revenue in A is greater than the expected revenue in B in this equilibrium. In contrast, we show that higher click-rates do not necessarily imply higher revenues in the second model.

3.16 Mission planning: Logic and game theory in multi-robot applications

Peter Novak (Czech Technical University, CZ)

I presented a research project towards study and implementation of a temporally extended multi-agent planner. The core idea revolves around using tools and techniques developed for reasoning about temporal modal logics for computing multi-agent plans. The main motivation for the mission planner is the actual need in many, especially military and defense domains, such as e.g., planning for tactical missions in urban operations, rescue operations, undersea anti-submarine warfare missions, or port guarding. I discussed use of position logic and algorithmic game theory could assume in such a planning problem.

3.17 Reasoning with Plans under Imperfect Information

Eric Pacuit (Tilburg University, NL)

Joint work of Pacuit, Eric; Simon, Sunil

Various combinations of temporal logics, epistemic and doxastic logics, and action logics have been used to reason about (groups of) agents in social situations. A key issue that has emerged is how best to represent and reason about the underlying protocol that governs the agents' interactions in a particular social situation. In this paper, we propose a PDL-style logic for reasoning about protocols, or plans, under imperfect information.

Our paper touches on a number of issues surrounding the relationship between an agent's abilities, available choices and information in an interactive situation. The main question we address is under what circumstances can the agent commit to a protocol or plan, and what can she achieve by doing so?

3.18 The Power of Knowledge in Games

Rohit Parikh (City University of New York, US)

We develop a theory of the interaction between knowledge and games.

Epistemic game theory is of course a well developed subject But there is also a need for theory of how some agents can *affect* the outcome of a game by affecting the knowledge which other agents have and thereby affecting their actions.

We concentrate on games of incomplete or imperfect information, and study how cautious, median seeking, or aggressive players might play such games. We provide models for the behavior of a knowledge manipulator who seeks to manipulate the knowledge states of active players in order to affect their moves and to maximize her own payoff even while she herself remains inactive.

3.19 Neighbourhood structures in large games

Ramaswamy Ramanujam (The Institute of Mathematical Sciences – Chennai, IN)

Joint work of Paul, Soumya; Ramanujam, R.

We study repeated normal form games where the number of players is large and suggest that it is useful to consider a neighbourhood structure on the players.

The game proceeds in rounds where in each round the players of every clique play a strategic form game among each other. Based on what a player observes, i.e., the strategies and the outcomes in the previous round of the players visible to her, the player may switch strategies in the same neighbourhood, or migrate to another neighbourhood. We show that given the initial neighbourhood graph and the types of the players in a simple modal logic, we can effectively decide if the game eventually stabilises. We also prove a characterisation result for games for arbitrary types using potentials. We then offer some applications to the special case of weighted co-ordination games where we can compute bounds on how long it takes to stabilise.

3.20 Epistemic, Strategic ATL* with Explicit Strategies

Henning Schnoor (Universität Kiel, DE)

Strategic logics have been used successfully to reason about game-like situations such as multi-agent systems. An important aspect of a game is the amount of information available to each player. We discuss a variation of the standard strategic logic ATL with knowledge-related aspects. In particular, we require that players are able to identify strategies for a given goal, and we allow formulas to explicitly reason about strategies.

3.21 Efficiency Levels in Sequential Auctions with dynamic Arrivals

Ella Segev (Ben Gurion University – Beer Sheva, IL)

Joint work of Lavi, Ron; Segev, Ella

In an environment with dynamic arrivals of players who wish to purchase only one of multiple identical objects for which they have a private value, we analyze a sequential auction mechanism with an activity rule. If the players play undominated strategies then we are able to bound the efficiency loss compared to an optimal mechanism that maximizes the total welfare. We have no assumptions on the underlying distribution from which the players' arrival times and valuations for the object are drawn.

Moreover we have no assumption of a common prior on this distribution.

3.22 Logics for social choice and perspectives on their software implementation

Nicolas Troquard (University of Essex, GB)

Joint work of Troquard, Nicolas; van der Hoek, Wiebe; Wooldridge, Michael
Main reference to appear in Journal of Philosophical Logic
URL http://arxiv.org/abs/1102.3341

In this talk, I will present a logic to reason about voting procedures, proposed in a common work with W. van der Hoek and M. Wooldridge. I will discuss some perspectives on its software implementation and try to assess its practicality. I will also make a short demonstration of a prototype.

3.23 Turning Competition into Cooperation and Cooperation into Competition

Paolo Turrini (Utrecht University, NL)

Joint work of Goranko, Valentin; Grossi, Davide; Jamroga, Wojtek; Turrini, Paolo
Main reference Goranko, V., Jamroga, W., Turrini, P.: Strategic Games and Truly Playable Effectivity Functions Proceedings of the 10th international conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2011) Taipei, Taiwan (to appear)
URL http://www2.imm.dtu.dk/ vfgo/papers/GJTcamera-ready-final.pdf

The work presented relates strategic games, object of study of non-cooperative game theory, and effectivity functions, object of study of cooperative game theory. It consists of two parts:

- The first part, joint work with Valentin Goranko and Wojtek Jamroga, shows that a believed correspondence between strategic games and a class of effectivity functions, known as Pauly's Representation Theorem, is not correct as it stands. A proof of this fact is presented and an alternative correspondence established.

- The second part, joint work with Davide Grossi, presents a weakening of the classical relation between strategic games and effectivity functions and study the reciprocity among players as precondition of coalition formation.

3.24 A Proof System for Message Passing

Jan van Eijck (CWI – Amsterdam, NL)

Joint work of van Eijck, Jan; Sietsma, Floor
Main reference work under submission

We propose a framework for message passing that combines the best properties of dynamic epistemic semantics and history-based approaches. We assume that all communication is truthful and reliable.

Our framework consists of Kripke models in which we keep a history of messages that have been sent. We introduce an update operation for message sending, with a corresponding action modality. With this update we can study the exact epistemic consequences of the sending of a message.

We define a class of models that is generated from initial Kripke models by means of message updates, we axiomatize our update modality, and we give examples of how the framework can be used to prove properties of message passing systems.

4 Open Problems

The discussion between different groups of researchers remained mostly on fundamental level. In consequence, besides open technical problems, we also identified some entry points to more fundamental research. Here is the list of main problems that have been posed and/or emerged during the seminar:

- Are quantitative probabilities necessary to analyze interaction? How can they be used in logic-based approaches?

- How can logical descriptions be helpful in quantitative analysis of games?
- What is the relation between different views of interaction (static vs. dynamic, a priori vs. runtime)
- What is the exact relationship between models and methodologies from noncooperative game theory, coalitional game theory, and some theories from social science (especially dependency theory)?
- What is the formal relationship between knowledge and power? Can it be characterized using a suitable logical language?
- Are game logics useful in analysis of security protocols? What features are necessary to extend the existing game logics like ATL to obtain a sufficiently expressive language?
- How can very large games be modeled and solved effectively?

5 Panel Discussions

We had three panel discussions. The first one was on *Strategic Analysis*. We discussed the following questions:

1. Why is *logic* important for *game theory*? Or is it?
2. Why is *probability* so rare in logic of *game theory*?
3. The *zoo of logics*: Is there a common framework?
4. *Logics of games* vs. *games of logic* (game semantics, game-based verification).

The second discussion was on *Algorithms and Complexity*:

1. *Complexity*: Is it just theory and of *no practical use*?
2. Are there any *specific features* of complexity analysis for *games*?
3. Can we really use *game theory* and *logics of games* (despite complexity obstacles)?

The final discussion treated *Applications*:

1. Can we really use game theory and logics of games *out of academia*?
2. Are existing *models of interaction* adequate?
3. Can models used by different disciplines be *integrated* in a meaningful way?

The discussions identified a number of meeting points between the main disciplines (logic, game theory, applications). The need for common framework is strong, not only between logic and game theory, but also one that unifies different logic-based approaches. Modal logic seems to be *lingua franca* in the latter case, but there are many variations of modal logic being used, and the exact relationship between them is unclear.

In order to *use* the frameworks being developed in logic-based as well as economics approaches to interaction, efficient algorithms are essential. It has been also pointed out that classical complexity theory – while useful – often gives misleading estimations of the “real” complexity that one faces in applications. Studying *average complexity* (instead of worst-case complexity) can be a good way out, though average complexity is usually much harder to establish.

Participants

- Thomas Agotnes
University of Bergen, NO
- Liad Blumrosen
The Hebrew University of Jerusalem, IL
- Giacomo Bonanno
Univ. of California – Davis, US
- Jan M. Broersen
Utrecht University, NL
- Manuela-Luminita Bujorianu
Manchester University, GB
- Nils Bulling
TU Clausthal, DE
- Jan Calta
HU Berlin, DE
- Mehdi Dastani
Utrecht University, NL
- Alfredo Di Tillio
Universitá Bocconi-Milan, IT)
- Catalin Dima
LACL, FR
- Jürgen Dix
TU Clausthal, DE
- Naipeng Dong
University of Luxembourg, LU
- Ulle Endriss
University of Amsterdam, NL
- Sujata Ghosh
University of Groningen, NL
- Valentin Goranko
Technical Univ. of Denmark, DK
- Paul Harrenstein
LMU München, DE
- Andreas Herzig
Université Paul Sabatier – Toulouse, FR
- Wojtek Jamroga
University of Luxembourg, LU
- Piotr Kazmierczak
University of Bergen, NO
- Max Knobbout
Utrecht University, NL
- Michael Koester
TU Clausthal, DE
- Ron Lavi
Technion – Haifa, IL
- Viliam Lisy
Czech Technical University, CZ
- Emiliano Lorini
Université Paul Sabatier – Toulouse, FR
- Sjouke Mauw
University of Luxembourg, LU
- Matthijs Melissen
University of Luxembourg, LU
- Dov Monderer
Technion – Haifa, IL
- Peter Novak
Czech Technical University, CZ
- Eric Pacuit
Tilburg University, NL
- Rohit Parikh
City University of New York, US
- Alan Perotti
University of Torino, IT
- Ramaswamy Ramanujam
The Institute of Mathematical Sciences – Chennai, IN
- Olivier Roy
LMU München, DE
- Dov Samet
Tel Aviv University, IL
- Federico Schlesinger
TU Clausthal, DE
- Henning Schnoor
Universität Kiel, DE
- Pierre-Yves Schobbens
University of Namur, BE
- Ella Segev
Ben Gurion University – Beer Sheva, IL
- Inanc Seylan
Free Univ. Bozen-Bolzano, IT
- Floor Sietsma
CWI – Amsterdam, NL
- Sunil Simon
CWI – Amsterdam, NL
- Marija Slavkovik
University of Luxembourg, LU
- Sonja Smets
University of Groningen, NL
- Nicolas Troquard
University of Essex, GB
- Paolo Turrini
Utrecht University, NL
- Leon van der Torre
University of Luxembourg, LU
- Jan van Eijck
CWI – Amsterdam, NL

Report from Dagstuhl Seminar 11111

Computational Geometry

Edited by

Pankaj Kumar Agarwal[1], Kurt Mehlhorn[2], and Monique Teillaud[3]

1 **Duke University, US,** pankaj@cs.duke.edu
2 **MPI für Informatik - Saarbrücken, DE,** mehlhorn@mpi-inf.mpg.de
3 **INRIA Sophia Antipolis - Méditerranée, FR,** Monique.Teillaud@inria.fr

Abstract

This report documents the outcomes of Dagstuhl Seminar 11111 "Computational Geometry". The Seminar gathered fifty-three senior and younger researchers from various countries in the unique atmosphere offered by Schloss Dagstuhl. Abstracts of talks are collected in this report as well as a list of open problems.

Seminar 13.–18. March, 2011 – www.dagstuhl.de/11111
1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity (Geometrical problems and computations), G.4 Mathematical Software (Algorithm design and Analysis)
Keywords and phrases Algorithms, geometry, combinatorics, topology, theory, applications, implementation
Digital Object Identifier 10.4230/DagRep.1.3.19
Edited in cooperation with Saurabh Ray

1 Executive Summary

Pankaj Kumar Agarwal
Kurt Mehlhorn
Monique Teillaud

Computational Geometry and its Evolution

The field of computational geometry is concerned with the design, analysis, and implementation of algorithms for geometric problems, which arise in a wide range of areas, including computer graphics, CAD, robotics computer vision, image processing, spatial databases, GIS, molecular biology, and sensor networks. Since the mid 1980s, computational geometry has arisen as an independent field, with its own international conferences and journals.

In the early years mostly theoretical foundations of geometric algorithms were laid and fundamental research remains an important issue in the field. Meanwhile, as the field matured, researchers have started paying close attention to applications and implementations of geometric algorithms. Several software libraries for geometric computation (e.g. LEDA, CGAL, CORE) have been developed. Remarkably, this emphasis on applications and implementations has emerged from the originally theoretically oriented computational geometry community itself, so many researchers are concerned now with theoretical foundations as well as implementations.

Computational Geometry, *Dagstuhl Reports*, Vol. 1, Issue 3, pp. 19–41
Editors: Pankaj Kumar Agarwal, Kurt Mehlhorn, and Monique Teillaud
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Seminar Topics

The emphasis of the seminar was on presenting the recent developments in the field as well as identifying new challenges. We have identified a few broad topics, listed below, that cover both theoretical and practical issues in computational geometry and that we believe are some of the most interesting subareas in the field.

- *Theoretical foundations* of computational geometry lie in combinatorial geometry and its algorithmic aspects. They are of an enduring relevance for the field, particularly the design and the analysis of efficient algorithms require deep theoretical insights.
- Various *applications* such as robotics, GIS, or CAD lead to interesting variants of the *classical topics* originally investigated, including convex hulls, Voronoi diagrams and Delaunay triangulations, and geometric data structures. For example, Voronoi diagrams and nearest-neighbor data structures under various metrics have turned out to be useful for many applications and are being investigated intensively.
- Because of applications in molecular biology, computer vision, geometric databases, *shape analysis* has become an important topic. Not only it It raises many interesting geometric questions ranging from modeling and reconstruction of surfaces to shape similarity and classification, but it has also led to the emergence of the so-called field *computational topology*.
- In many applications the data lies in very high dimensional space and typical geometric algorithms suffer from the curse of dimensionality. This has led to extensive work on dimension-reduction and embedding techniques.
- Massive geometric data sets are being generated by networks of sensors at unprecedented spatial and temporal scale. How to store, analyze, query, and visualize them has raised several algorithmic challenges. New computational models have been proposed to meet these challenges, e.g., streaming model, communication-efficient algorithms, and maintaining geometric summaries.
- *Implementation issues* have become an integral part of the research in computational geometry. Besides general software design questions especially *robustness* of geometric algorithms is important. Several methods have been suggested and investigated to make geometric algorithms numerically robust while keeping them efficient, which lead to interaction with the field of computer algebra, numerical analysis, and topology.

Participants

53 researchers from various countries and continents attended the meeting. This high number shows the strong interest of the community for this event. The feedback from participants was very positive.

Dagstuhl seminars on computational geometry have been organized since 1990, lately in a two year rhythm. They have been extremely successful both in disseminating the knowledge and identifying new research thrusts. Many major results in computational geometry were first presented in Dagstuhl seminars, and interactions among the participants at these seminars have led to numerous new results in the field. These seminars have also played an important role in bringing researchers together and fostering collaboration. They have arguably been the most influential meetings in the field of computational geometry.

A session of this Seminar was dedicated to our dear friend Hazel Everett, deceased on July 20th, 2010.

The place itself is a great strength of the Seminar. Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions.

Therefore, we warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl!

2 Table of Contents

Executive Summary
Pankaj Kumar Agarwal, Kurt Mehlhorn, and Monique Teillaud 19

Overview of Talks

Computing the depth of an arrangement of axes-parallel rectangles in parallel
Helmut Alt 24

Memory-constrained algorithms
Tetsuo Asano 24

Rips complexes for Shape Reconstruction
Dominique Attali 25

The Height of a Homotopy
Erin Moriarty Wolf Chambers 25

A Generalization of Kakeya's Problem
Otfried Cheong 25

Star Trek Replicators via Staged Assembly
Erik D. Demaine 26

The Effect of Noise on the Number of Extreme Points
Olivier Devillers 26

Improved Bound for the Union of Fat Triangles
Esther Ezra 27

Exact Solutions and Bounds for General Art Gallery Problems
Sandor Fekete 27

Range Queries in Distributed Networks
Jie Gao 27

Intersection patterns of convex sets
Xavier Goaoc 28

On Path Quality in Sampling-Based Motion Planning
Dan Halperin 28

Algorithms for Persistent Homology
Michael Kerber 29

Polygonal paths of bounded curvature
David G. Kirkpatrick 29

VC Dimension of Visibility
Rolf Klein 30

The complexity of Ham-Sandwich cuts in high dimension
Christian Knauer 30

On the topology of plane algebraic curves
Sylvain Lazard 30

The Phsyarum Computer
Kurt Mehlhorn 31

Recent Progress on Some Covering Tour Problems
Joseph S. Mitchell 31

Natural Neighbor Interpolation Based Grid DEM Construction Using a GPU
Thomas Molhave 32

Tight bounds for epsilon-nets
Janos Pach 32

Double Permutation Sequences and Arrangements of Planar Families of Convex Sets
Richard Pollack 33

Shortest Paths in Time-Dependent FIFO Networks
Joerg-Ruediger Sack 33

Can Nearest Neighbor Search be Simple and always Fast?
Raimund Seidel 34

From joints to distinct distances and beyond: The dawn of an algebraic era in combinatorial geometry
Micha Sharir 34

Fixed Points of the Restricted Delaunay Triangulation Operator
Jonathan Shewchuk 35

The potential to improve the choice: List coloring for geometric hypergraphs
Shakhar Smorodinsky 35

Shapely Measures, and Measures on Shapes: The story of the Kernel distance
Suresh Venkatasubramanian 36

Kasteleyn and the Number of Crossing-Free Matchings and Cycles on a Planar Point Set
Emo Welzl 36

Computing the Frechet Distance Between Folded Polygons
Carola Wenk 37

Triangular Meshes: Registration and GPU-based Distance Computation
Nicola Wolpert 37

Cxyz: Isotopic Subdivision Algorithm for Non-Singular Surfaces
Chee K. Yap 38

Median trajectories
Marc van Kreveld 38

Open Problems 39

Participants 41

3 Overview of Talks

3.1 Computing the depth of an arrangement of axes-parallel rectangles in parallel

Helmut Alt (FU Berlin, DE)

Joint work of Alt, Helmut; Hagerup, Torben; Scharf, Ludmila

We consider the problem of determining the depth of an arrangement of axes parallel rectangles in the plane, i.e., the highest number of rectangles intersecting in one point. Sequentially, standard procedures have quadratic runtime since the size of the arrangement is quadratic in the number of rectangles. However, $O(n \log n)$ algorithms are known for this problem.

We design *parallel* algorithms for this problem. We consider a structure related to interval trees which is traversed top-down level by level propagating and adjusting the information belonging to the interval associated to a node from parent to child. We obtain a parallel runtime of $O(log^2 n)$ with a total of O(n log n) operations or a parallel runtime of $O(logn)$ with $O(n^(1 + \varepsilon))$ operations for any constant $\varepsilon > 0$.

3.2 Memory-constrained algorithms

Tetsuo Asano (JAIST – Nomi, JP)

Main reference T. Asano, W. Mulzer, and Y. Wang, "Constant-Work-Space Algorithm for a Shortest Path in a Simple Polygon," Invited talk, Proc. 4th International Workshop on Algorithms and Computation, WALCOM, Dhaka, Bangladesh, 2010, LNCS 5942, pp.9–20.s
URL http://dx.doi.org/10.1007/978-3-642-11440-3_2

In this talk I will introduce algorithms with limited work space, which are desired for applications to highly functional hardware such as scanners, digital cameras, and Android cellular phones. One extreme set of memory- constrained algorithms have been studied under the name of log-space algorithms which use only $O(\log n)$ bits for their work space. This talk starts with a simple example of a memory-constrained algorithms based on a general paradigm for designing such algorithms. Then, it is extended to a problem on image processing, where work space of square root of an input size is more realistic. A practically efficient algorithm with work space of $O(\sqrt{n} \log n)$ bits is given for a problem of extracting an interesting region in a given color image.

3.3 Rips complexes for Shape Reconstruction

Dominique Attali (GRIPSA Lab – Saint Martin d'Hères, FR)

Joint work of Attali, Dominique; Lieutier, André; Salinas, David
Main reference D. Attali, A. Lieutier and D. Salinas, "Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes," Proc. 27th Ann. Sympos. Comput. Geom., Paris, France, June 13–15 2011
URL http://hal.archives-ouvertes.fr/hal-00579864/en/

We associate with each compact set X of $\mathbb{R}^n$ two real-valued functions c_X and h_X defined on $\mathbb{R}^+$ which provide two measures of how much1 the set X fails to be convex at a given scale. First, we show that, when P is a finite point set, an upper bound on $c_P(t)$ entails that the Rips complex of P at scale r collapses to the Čech complex of P at scale r for some suitable values of the parameters t and r. Second, we prove that, when P samples a compact set X, an upper bound on h_X over some interval guarantees a topologically correct reconstruction of the shape X either with a Čech complex of P or with a Rips complex of P. Regarding the reconstruction with Čech complexes, our work compares well with previous approaches when X is a smooth set and surprisingly enough, even improve constants when X has a positive μ-reach. Most importantly, our work shows that Rips complexes can also be used to provide topologically correct reconstruction of shapes. This may be of some computational interest in high dimension.

3.4 The Height of a Homotopy

Erin Moriarty Wolf Chambers (St. Louis University, US)

In our paper in CCCG 2009, we examined the problem of computing the height of a homotopy. This can be phrased as a very combinatorial problem, and indeed has been studied in at least one very different context, submodular percolation. We proved several properties of homotopies that obtain the minimum height value, but were unable to completely characterize or compute the minimum height homotopy. More recently, new work has been completed to compute an $O(\log n)$ approximation, but again no exact algorithms or hardness results are known. We will survey known results and techniques for this problem.

3.5 A Generalization of Kakeya's Problem

Otfried Cheong (KAIST – Daejeon, KR)

Given a not necessarily finite family F of line segments, we show that the smallest-area convex figure P such that every segment in F can be translated to lie in P can always be chosen to be a triangle.

This generalizes the result by Pal from 1921 for the case where F contains a unit-length segment of every possible orientation.

Our result can be rephrased as follows: Given a convex figure P, there is always a triangle T of area at most the area of P such that for every possible direction, the width of T in that direction is not less than the width of P in that direction.

We also given an algorithm that computes the smallest-area triangle as above when the input F is a set of n line segments in time $O(n \log n)$.

3.6 Star Trek Replicators via Staged Assembly

Erik D. Demaine (MIT – Cambridge, US)

Tile self-assembly is an intriguing approach to manufacturing desired shapes with nano-scale feature size. A recent direction in this theory allows the use of multiple stages—operations performed by the experimenter, such as mixing two self-assembling systems together. This flexibility transforms the experimenter from a passive entity into a parallel algorithm, and vastly reduces the number of distinct parts required to construct a desired shape, possibly making the systems practical to build. The staged-assembly perspective also enables the possibility of additional operations, such as adding an enzyme that destroys all tiles with a special label. By enabling destruction in addition to the usual construction, we can perform tasks impossible in a traditional self-assembly system, such as replicating many copies of a given object's shape, without knowing anything about that shape, and building an efficient nano computer.

3.7 The Effect of Noise on the Number of Extreme Points

Olivier Devillers (INRIA Sophia Antipolis, FR)

Joint work of Attali, Dominique; Devillers, Olivier; Goaoc, Xavier
Main reference D. Attali, O. Devillers, X. Goaoc, "The Effect of Noise on the Number of Extreme Points", Research Report 7134, INRIA, 2009.
URL http://hal.inria.fr/inria-00438409/

Assume that Y is a noisy version of a point set X in convex position. How many vertices does the convex hull of Y have, that is, what is the number of extreme points of Y?

We consider the case where X is an (ϵ, κ)-sample of a sphere in $\mathbb{R}^d$ and the noise is random and uniform: Y is obtained by replacing each point $x \in X$ by a point chosen uniformly at random in some region $S(x)$ of size δ around x. We give upper and lower bounds on the expected number of extreme points in Y when $S(x)$ is a ball (in arbitrary dimension) or an axis-parallel square (in the plane). Our bounds depend on the size n of X and δ, and are tight up to a polylogarithmic factor. These results naturally extend in various directions (more general point sets, other regions $S(x)\ldots$).

We also present experimental results, showing that our bounds for random noise provide good estimators of the behavior of snap-rounding, that is when Y is obtained by rounding each point of X to the nearest point on a grid of step δ.

3.8 Improved Bound for the Union of Fat Triangles

Esther Ezra (New York University, US)

Joint work of Ezra, Esther; Aronov, Boris; Sharir, Micha
Main reference E. Ezra, B. Aronov, M. Sharir, "Improved Bound for the Union of Fat Triangles," Proc. 22nd Annual ACM-SIAM Symp. on Discrete Algorithms (SODA'11), to appear.
URL http://www.cims.nyu.edu/ esther/Publications/fatri.pdf

We show that, for any fixed $\delta > 0$, the combinatorial complexity of the union of n triangles in the plane, each of whose angles is at least δ, is $O(n2^{\alpha(n)} \log^* n)$, with the constant of proportionality depending on δ. This considerably improves the twenty-year-old bound $O(n \log \log n)$, due to Matousek etal.

3.9 Exact Solutions and Bounds for General Art Gallery Problems

Sandor Fekete (TU Braunschweig, DE)

Joint work of Fekete, Sándor; Kröller, Alexander; Schmidt, Christiane; Kamphans, Tom; Baumgartner, Tobias
Main reference Proceedings of the SIAM-ACM Workshop on Algorithm Engineering and Experiments (ALENEX'10), pp. 11–22.
URL http://www.siam.org/proceedings/alenex/2010/alx10_002_baumgartnert.pdf

The classical Art Gallery Problem asks for the minimum number of guards that achieve visibility coverage of a given polygon. This problem is known to be NP-hard, even for very restricted and discrete special cases. For the general problem (in which both the set of possible guard positions and the point set to be guarded are uncountable), neither constant-factor approximation algorithms nor exact solution methods are known. We present a primal-dual algorithm based on linear programming that provides lower bounds on the necessary number of guards in every step and (in case of convergence and integrality) ends with an optimal solution. We describe our implementation and give results for an assortment of polygons, including non-orthogonal polygons with holes.

3.10 Range Queries in Distributed Networks

Jie Gao (SUNY – Stony Brook, US)

Joint work of Gao, Jie; Sarkar, Rik

Consider mobile targets moving in a plane and their movements being monitored by a network such as a field of sensors. We develop distributed algorithms for in-network tracking and range queries for aggregated data (for example returning the number of targets within any user given region). Our scheme stores the target detection information locally in the network, and answers a query by examining the perimeter of the given range. The cost of updating data about mobile targets is proportional to the target displacement.

The key insight is to maintain in the sensor network a function with respect to the target detection data on the graph edges that is a differential one-form such that the integral of this one-form along any closed curve C gives the integral within the region bounded by C.

3.11 Intersection patterns of convex sets

Xavier Goaoc (INRIA Lorraine, FR)

Joint work of Colin de Verdière, Éric; Ginot, Grégory; Goaoc, Xavier
Main reference Eric Colin de Verdiere, Gregory Ginot, Xavier Goaoc, "Helly numbers of acyclic families," arXiv:1101.6006v2
URL http://arxiv.org/abs/1101.6006

The Helly number of a family of sets (with empty intersection) is the size of its largest inclusion-wise minimal sub-family with empty intersection. We show how techniques from homology theory lead to fairly general conditions under which the Helly number of a family can be bounded.

Our typical result is along the following lines: "if F is a family of sets in $\mathbb{R}^d$ such that the intersection of any subfamily has at most r connected components, each of which is a homology cell then the Helly number of F is at most $r(d+1)$". This result can be generalized so as to imply, in a unified way, bounds on Helly numbers in geometric transversal theory that had previously been studied by ad hoc techniques.

3.12 On Path Quality in Sampling-Based Motion Planning

Dan Halperin (Tel Aviv University, IL)

Joint work of Angela Enosh, Angela; Halperin, Dan; Nechsushtan, Oren; Raveh, Barak
Main reference (1) Raveh-Enosh-Halperin, A Little More, a Lot Better: Improving Path Quality by a Path-Merging Algorithm, IEEE Trans. on Robotics, 27/2, 2011, pp 365–371.
(2) Nechushtan-Raveh-Halperin, Sampling-Diagram Automata: A Tool for Analyzing Path Quality in Tree Planners, WAFR 2010.
URL http://acg.cs.tau.ac.il/projects

Sampling-based motion planners are a central tool for solving motion-planning problems in a variety of domains, but the theoretical understanding of their behavior remains limited, in particular with respect to the quality of the paths they generate (in terms of path length, clearance, etc.). We prove, for a simple family of obstacle settings, that the popular dual-tree planner Bi-RRT may produce low-quality paths that are arbitrarily worse than optimal with modest but significant probability, and overlook higher-quality paths even when such paths are easy to produce. At the core of our analysis are probabilistic automata designed to reach an accepting state when a path of significantly low quality has been generated.

Complementary experiments suggest that our theoretical bounds are conservative and could be further improved. We also present a method to improve path quality by merging an arbitrary number of input motion paths into a hybrid output path of superior quality, for a broad and general formulation of path quality. Our approach is based on the observation that the quality of certain sub-paths within each solution may be higher than the quality of the entire path.

3.13 Algorithms for Persistent Homology

Michael Kerber (IST Austria – Klosterneuburg, AT)

Joint work of Kerber, Michael; Chen, Chao

Persistent homology is a quickly-growing area of research in the analysis of topological spaces. One reason for its success is the existence of an efficient algorithm to solve the problem using Gaussian elimination. I will present two recent results in this context. First, I show a simple optimization technique of the default algorithm that avoids column operations on roughly half of the columns.

This yields both significant practical improvements, and provides new insights on the complexity of persistence for certain special cases.

Second, I will present a divide-and-conquer approach to compute persistence based on rank computations of submatrices instead of Gaussian elimination.

The algorithm only outputs homology classes with persistence larger than a given threshold, and permits an output-sensitive complexity analysis. In particular, using a Monte-Carlo algorithm for rank computation and assuming that the number of returned classes is logarithmic in the input size, an quadratic algorithm for persistence computation is achieved, modulo logarithmic factors.

3.14 Polygonal paths of bounded curvature

David G. Kirkpatrick (University of British Columbia – Vancouver, CA)

Joint work of Kirkpatrick, David G.; Polishchuk, Valentin

In one of the seminal papers [Dubins57] in non-holonomic motion planning, L.E. Dubins developed a strong characterization of minimum length curves of bounded curvature joining fixed initial and final configurations (specified by position and direction): in the absence of obstacles, such paths consist of two (possibly degenerate) circular arcs joined by either a straight line segment or a third circular arc. Dubins' original proof uses advanced calculus; subsequently the same result was reproved using control-theoretic techniques.

We revisit bounded-curvature in the context of polygonal paths (paths consisting of a sequence of straight line segments) and formulate a natural notion of bounded-curvature for such paths. While polygonal paths clearly violate curvature bounds in sufficiently small neighbourhoods, our notion still manages to capture the less restrictive constraint that they "do not turn too sharply". We present an elementary and purely geometric proof of a characterization result, analogous to that of Dubins, for minimum length polygonal paths satisfying our bounded curvature property. This not only provides a discrete analogue of continuous motion of bounded curvature, but it also gives a fundamentally new proof of Dubins' original result as a limiting case of our constructions.

3.15 VC Dimension of Visibility

Rolf Klein (Universität Bonn, DE)

Main reference A. Gilbers and R. Klein, "A New Upper Bound For the VC-Dimension of Visibility Regions," Proc. 27th ACM Symp. on Computational Geometry (SoCG'11), pp.380–386.
URL http://doi.acm.org/10.1145/1998196.1998259

We lower the upper bound for the VC dimension of visibility in simple polygons from 23 to 14.

3.16 The complexity of Ham-Sandwich cuts in high dimension

Christian Knauer (Universität Bayreuth, DE)

We study a canonical decision problem arising from the Ham-Sandwich cut theorem. We show it to be W[1]-hard (and NP-hard) if the dimension is part of the input. This is done by an fpt-reduction (which is actually a ptime-reduction) from the d-SUM problem. Our reduction also implies that the problem cannot be solved in time $n^{o(d)}$ unless the Exponential-Time Hypothesis (ETH) is false.

3.17 On the topology of plane algebraic curves

Sylvain Lazard (INRIA - Nancy, FR)

Joint work of Lazard, S.; Bouzidi, Y.; Cheng, J.; Penaranda, L.; Pouget, M.; Rouillier, F.; Tsigaridas, E.

I will present some recent results on the computation of the topology of planar algebraic curves. Our approach is based on a rectangular decomposition of the plane, instead of the standard CAD, which has the advantage of being oblivious to degenerate configurations. It is also based on a new algorithm for computing the critical points of the input curve. We first decompose the corresponding system into subsystems according to the number of roots (counted with multiplicities) in vertical lines, as presented by Gonzalez-Vega and Necula in 2002. We then show how these systems can be efficiently solved by computing their lexicographic Grobner bases and Rational Univariate Representations. We also show how this approach can be performed using modular arithmetic, while remaining deterministic. We finally demonstrate that our approach yields a substantial gain of a factor between 1 to 200 on curves of degree up 40 compared to state-of-the-art implementation.

3.18 The Phsyarum Computer

Kurt Mehlhorn (MPI für Informatik - Saarbrücken, DE)

Physarum is a slime mold. It was observed over the past 10 years that the mold is able to solve shortest path problems and to construct good Steiner networks (Nakagaki-Yamada-Toth, Tero-Takagi-et al). In a nutshell, the shortest path experiment is as follows: A maze is built and the mold is made to cover the entire maze. Food is then provided at two positions s and t and the evolution of the slime is observed. Over time, the slime retracts to the shortest s-t-path.

A mathematical model of the slime's dynamic behaviour was proposed by Tero-Kobayashi-Nakagaki.

Extensive computer simulations confirm the experimental findings; the slime retracts to the shortest path. We (joint work with Vincenzo Bonifaci, Girish Varma) have recently proved convergence.

I will start with a video showing the mold in action. Then I review the mathematical model, explain the computer simulation, and run some computer experiments. Next, I will discuss how we formulated the right conjecture based on computer experiments. Finally, I will briefly discuss the convergence proof.

3.19 Recent Progress on Some Covering Tour Problems

Joseph S. Mitchell (SUNY – Stony Brook, US)

Joint work of Mitchell, Joseph S.; Arkin, E.; Polishchuk, V.; Yang, S.

We give some recent results on some covering tour problems:

(1) We answer the question initially posed by Arik Tamir at the Fourth NYU Computational Geometry Day (March, 1987):

"Given a collection of compact sets, can one decide in polynomial time whether there exists a convex body whose boundary intersects every set in the collection?"

We prove that when the sets are segments in the plane, deciding existence of the convex stabber is NP-hard. We also show that in 3D the stabbing problem is hard when the sets are balls. On the positive side, we give a polynomial-time algorithm to find a convex polygonal transversal (on a given discrete set of vertices) of a maximum number of segments in 2D if the segments are pairwise-disjoint. Our algorithm also finds a convex stabber of the maximum number of a set of polygonal pseudodisks in the plane.

The stabbing problem is related to "convexity" of point sets and polygons, measured as the minimum distance by which the points/polygons must be shifted in order that there is a convex stabber.

(2) Watchman routes in polygons with holes: We prove a lower bound of $\Omega(\log n)$ on approximation, from SET-COVER. We also give an $O(\log n)$-approximation algorithm for the capae that a certain "bounded perimeter assumption" holds.

3.20 Natural Neighbor Interpolation Based Grid DEM Construction Using a GPU

Thomas Molhave (Duke University, US)

Joint work of Beutel, Alex; Mølhave, Thomas; Agarwal, Pankaj K.
Main reference Natural Neighbor Interpolation Based Grid DEM Construction Using a GPU. Alex Beutel, Thomas Malhave, Pankaj K. Agarwal. GIS '10: Proceedings of the 18th ACM SIGSPATIAL International Symposium on Advances in Geographic Information Systems, 2010.
URL http://dx.doi.org/10.1145/1869790.1869817

With modern LiDAR technology the amount of topographic data, in the form of massive point clouds, has increased dramatically.

One of the most fundamental GIS tasks is to construct a grid digital elevation model (DEM) from these 3D point clouds. In this paper we present a simple yet very fast algorithm for constructing a grid DEM from massive point clouds using natural neighbor interpolation (NNI). We use a graphics processing unit (GPU) to significantly speed up the computation. To handle the large data sets and to deal with graphics hardware limitations clever blocking schemes are used to partition the point cloud. For example, using standard desktop computers and graphics hardware, we construct a high-resolution grid with 150 million cells from two billion points in less than thirty-seven minutes. This is about one-tenth of the time required for the same computer to perform a standard linear interpolation, which produces a much less smooth surface

3.21 Tight bounds for epsilon-nets

Janos Pach (EPFL – Lausanne, CH)

According to a well known theorem of Haussler and Welzl (1987), any range space of bounded VC-dimension admits an ϵ-net of size $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$. Using probabilistic techniques, Pach and Woeginger (1990) showed that there exist range spaces of VC-dimension 2, for which the above bound is sharp. The only known range spaces of small VC-dimension, in which the ranges are geometric objects in some Euclidean space and the size of the smallest ϵ-nets is superlinear in $\frac{1}{\epsilon}$, were found by Alon (2010). In his examples, the size of the smallest ϵ-nets is $\Omega\left(\frac{1}{\epsilon}g(\frac{1}{\epsilon})\right)$, where g is an extremely slowly growing function, closely related to the inverse Ackermann function.

We show that there exist geometrically defined range spaces, already of VC-dimension 2, in which the size of the smallest ϵ-nets is $\Omega\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$. We also construct range spaces induced by axis-parallel rectangles in the plane, in which the size of the smallest ϵ-nets is $\Omega\left(\frac{1}{\epsilon}\log\log\frac{1}{\epsilon}\right)$. By a theorem of Aronov, Ezra, and Sharir (2010), this bound is tight.

3.22 Double Permutation Sequences and Arrangements of Planar Families of Convex Sets

Richard Pollack (New York University, US)

We (re)introduce Double Permutation Sequences, which provide a combinatorial encoding of arrangements of convex sets in the plane. We also recall the notion of a topological affine plane and several (some new) of its properties.

In particular, that there is a universal topological affine plane P (i.e. any finite arrangement of pseudolines is isomorphic to some arrangement of finitely many lines of P).

All of this work is joint with Jacob E. Goodman and some involves numerous other people, among whom are Raghavan Dhandapani, Andreas Holmsen, Shakhar Smorodinsky, Rephael Wenger, and Tudor Zamfirescu.

3.23 Shortest Paths in Time-Dependent FIFO Networks

Joerg-Ruediger Sack (Carleton University – Ottawa, CA)

Main reference F.Dehne, M.T.Omran, and J.-R.Sack, "Shortest paths in time-dependent FIFO networks," Algorithmica, to appear

In this talk, we study the time-dependent shortest paths problem for two types of time-dependent FIFO networks. First, we consider networks where the availability of links, given by a set of disjoint time intervals for each link, changes over time. Here, each interval is assigned a non-negative real value which represents the travel time on the link during the corresponding interval.

The resulting shortest path problem is the time-dependent shortest path problem for availability intervals, which asks to compute all shortest paths to any (or all) destination node(s) d for all possible start times at a given source node s. Second, we study time-dependent networks where the cost of using a link is given by a non-decreasing piece-wise linear function of a real-valued argument.

Here, each piece-wise linear function represents the travel time on the link based on the time when the link is used. The resulting shortest paths problem is the time-dependent shortest path problem for piece-wise linear functions which asks to compute, for a given source node s and destination d, the shortest paths from s to d, for all possible starting times.

We present an algorithm for both problems which improve significantly on the previously known algorithms.

3.24 Can Nearest Neighbor Search be Simple and always Fast?

Raimund Seidel (Universität des Saarlandes, DE)

We show that any method for nearest neighbor search that is simple, in the sense that the only operations involving query points are distance to site comparisons, cannot be fast on all inputs. In particular we show that any such methods can be forced to make $n-1$ such comparisons for some input, where n ist the number of sites.

3.25 From joints to distinct distances and beyond: The dawn of an algebraic era in combinatorial geometry

Micha Sharir (Tel Aviv University, IL)

In November 2010 the earth has shaken, when Larry Guth and Nets Hawk Katz posted a nearly complete solution to the distinct distances problem of Erdős, open since 1946. The excitement was twofold:

(a) The problem was one of the most famous problems, as well as one of the hardest nuts in the area, resisting solution in spite of many attempts (which only produced partial improvements).

(b) The proof techniques were algebraic in nature, drastically different from anything tried before.

The distinct distances problem is to show that any set of n points in the plane determine $Omega(n/\sqrt{\log n})$ distinct distances.

(Erdős showed that the grid attains this bound.) Guth and Katz obtained the lower bound $Omega(n/\log n)$.

Algebraic techniques of this nature were introduced into combinatorial geometry in 2008, by the same pair Guth and Katz. At that time they gave a complete solution to another (less major) problem, the so-called joints problem, posed by myself and others back in 1992.

Since then these techniques have led to several other developments, including an attempt, by Elekes and myself, to reduce the distinct distances problem to an incidence problem between points and lines in 3-space. Guth and Katz used this reduction and gave a complete solution to the reduced problem.

One of the old-new tools that Guth and Katz bring to bear is the Polynomial Ham Sandwich Cut, due to Stone and Tukey (1942). I will discuss this tool, including a “1-line” proof thereof, and its potential applications in geometry. One such application, just noted by Matoušek, is an algebraic proof of the classical Szemerédi-Trotter incidence bound for points and lines in the plane.

In the talk I will review all these developments, as time will permit.

Only very elementary background in algebra and geometry will be assumed.

3.26 Fixed Points of the Restricted Delaunay Triangulation Operator

Jonathan Shewchuk (Univ. of California – Berkeley, US)

Consider the problem of reconstructing an unknown smooth 2-manifold Σ embedded in three dimensions from a set of points S sampled from the surface. It is well known that if the point sample is sufficiently dense, then the restriction of the Voronoi diagram of S to the surface Σ dualizes to a triangulation that is homeomorphic to Σ. This triangulation is called the restricted Delaunay triangulation of S with respect to Σ. What if we take the restricted Delaunay triangulation of S with respect to this triangulation? And then repeat with the new triangulation? I show that the iteration always "converges" to a "fixed point," that is, a fixed set of triangles; and that for sufficiently dense samples, this set of triangles is likely to be a particularly nice reconstruction of Σ. Best of all, the idea works equally well for 2-manifolds embedded in much higher-dimensional spaces.

3.27 The potential to improve the choice: List coloring for geometric hypergraphs

Shakhar Smorodinsky (Ben Gurion University – Beer Sheva, IL)

Joint work of Smordinsky, Shakhar; Cheilaris, Panagiotis; Sulovsky, Marek

Given a hypergraph $H = (V, \mathcal{E})$, a coloring of its vertices is said to be conflict-free if for every hyperedge $S \in \mathcal{E}$ there is at least one vertex whose color is distinct from the colors of all other vertices in S.

This notion has been studied for several geometric hypergraphs and in various generalizations.

We study the list version of this notion:

In this version we are interested in the minimum number $k = k(H)$ such that if each vertex v is associated with a set of colors L_v of size k then one can pick a color for each vertex v from its "list" L_v such that the resulting coloring is conflict-free. Denote this number by $ch(H)$ (the conflict free "choice" number of H) It is easy to see that the minimum number of colors needed for a conflict-free coloring of H is bounded by $ch(H)$.

Let C be some absolute constant. We prove that for hypergraphs H with n vertices which are hereditarily C colorable (in the non-monochromatic sense) we have $ch(H) = O(\log n)$ and this bound is asymptotically tight. More over, we show that one can color the vertices of such a hypergraph from lists of size $O(\log n)$ such that the maximum color in any hyperedge is unique. The proof is constructive and uses a suitable potential function for constructing such coloring and analyzing the size of lists needed.

3.28 Shapely Measures, and Measures on Shapes: The story of the Kernel distance

Suresh Venkatasubramanian (University of Utah, US)

Joint work of Joshi, Sarang; Kommaraju, Raj Varma; Phillips, Jeff M.; Venkatasubramanian, Suresh
Main reference S.C. Joshi, R.V. Kommaraju, J.M. Phillips, S. Venkatasubramanian, "Comparing Distributions and Shapes using the Kernel Distance," Proc. 27th ACM Symp. on Computational Geometry, 2011, pp. 47–56.
URL http://doi.acm.org/10.1145/1998196.1998204

There are many ways to compute a distance between probability measures. However, if the underlying domain of the measures itself carries a metric, then the standard way to compare probability measures in a way that respects this is via the earthmover distance (or the transportation distance). This is a popular distance, especially in computer vision, but is expensive to compute.

There are also many ways to compare shapes, using methods from computational geometry, topology, and differential geometry. But these distances are all expensive to compute, are limited in important ways, and are difficult to use for applications that involve analyzing collections of shapes, rather than just pairs.

It turns out that there is a single mechanism that makes comparing distributions over metrics easy, and makes comparing (noisy) shapes (even surfaces) easy. It uses a kernel-based method to embed distributions (or shapes) in an (infinite-dimensional) Hilbert space, so that distance computation is a matter of computing the induced Hilbertian metric.

In this talk, I will describe algorithms for (i) estimating this "kernel distance" approximately and efficiently, (ii) computing sparse representations of shape s that are close to the original under the kernel distance, (iii) comparing shapes under rigid transformations using the kernel distance. I'll also briefly mention applications in the realm of clustering and metaclustering.

3.29 Kasteleyn and the Number of Crossing-Free Matchings and Cycles on a Planar Point Set

Emo Welzl (ETH Zürich, CH)

In 1967 Piet Kastelyn showed how to count the number of perfect matchings in a planar graph, simply by looking at the determinant of an appropriate skew symmetric variant of the adjacency matrix. Raimund Seidel observed an extremal combinatorics implication, namely that a planar graph on n vertices has at most $\sqrt[4]{6}^n$ perfect matchings (via the so-called Hadamard bound). We show how, for a planar set P of n points, that can be used to relate the number, $sc(P)$, of crossing-free straight-line spanning cycles (simple polygonizations) of P to the number of triangulations, $tr(P)$, of P: $sc(P) < 8\sqrt[4]{12}^n tr(P)$. We conjecture, though, that this can be significantly improved, actually to $sc(P) = O(c^n)tr(P)$ for some constant $c < 1$.

3.30 Computing the Frechet Distance Between Folded Polygons

Carola Wenk (Univ. of Texas at San Antonio, US)

Joint work of Cook, Atlas F. IV; Driemel, Anne; Har-Peled, Sariel; Sherette, Jessica; Wenk, Carola;

We present the first results showing that the Frechet distance between non-flat surfaces can be approximated within a constant factor in polynomial time. Computing the Frechet distance for surfaces is a surprisingly hard problem. It is not known whether it is computable, it has been shown to be NP-hard, and the only known algorithm computes the Frechet distance for flat surfaces (Buchin et al.). We adapt this algorithm to create one for computing the Frechet distance for a class of surfaces which we call folded polygons. Unfortunately, if extended directly the original algorithm no longer guarantees that a homeomorphism exists between the surfaces. We present three different methods to address this problem. The first of which is a fixed-parameter tractable algorithm.

The second is a polynomial-time approximation algorithm which approximates the optimum mapping. Finally, we present a restricted class of folded polygons for which we can compute the Frechet distance in polynomial time.

3.31 Triangular Meshes: Registration and GPU-based Distance Computation

Nicola Wolpert (Hochschule für Technik – Stuttgart, DE)

We present two projects with industrial partners. The first one is on the registration of triangular meshes. The motivation comes from sheet metal forming where after the stamping process the formed metal has to be compared with the original CAD-data. The main step here, based on the ICP-algorithm, is to compute the translation and rotation which best possibly aligns the two meshes.

The second project is called RASAND: Robust Algorithms for Distance Computation of large moving Triangular Meshes. The motivation comes from the digital mock up process where early in the construction phase of a mechanical part constructions and motions have to be evaluated. Mathematically we are given two large triangular meshes, one moving in discrete steps over time.

The question is to find all time steps and all pairs of triangles that become closer than a given epsilon. The challenge is the mass data and we are looking for data structures and algorithms that can answer this question using massively the computing-capability of modern graphics cards.

3.32 Cxyz: Isotopic Subdivision Algorithm for Non-Singular Surfaces

Chee K. Yap (New York University, US)

Joint work of Yap, Chee K.; Long Lin
URL http://cs.nyu.edu/exact/papers/

We consider exact algorithms for isotopic approximating of implicit surfaces given as the zero set of a C1 real function $f(x, y, z)$.

We focus on (domain) subdivision methods based on numerical predicates because they tend to have adaptive complexity, locality, and are easy to implement.

Numerical predicates, unlike algebraic or geometric predicates, admits a trade-off between efficacy (how tight is the bound) and efficiency (how quickly it can be computed); this can be exploited. The challenge of numerical predicates is how to ensure global correctness.

We describe an algorithm called Cxyz Algorithm (the name derives from the key predicate Cxyz in the algorithm).

It is a generalization our earlier Cxy algorithm for plane curves.

The algorithm exploits two properties of two previous approaches: non-local isotopy [Plantinga and Vegter, 2004] and parameterizability [Snyder, 1994].

Our preliminary implementations in Core Library suggest that our algorithm is very efficient compared to previous methods.

The proof of correctness is quite non-trivial. Termination is guaranteed for non-singular surface. Following Plantinga-Vegter, we first prove the correctness of a simpler algorithm called Regular Cxyz, and extend it to the Cxyz Algorithm which involves balancing of boxes.

As in the 2-D case, we need to deal with ambiguity.

A new phenomenon is that the arc connection rules require some global consistency properties.

We also briefly extend the method to discuss boundary processing, admitting a general Region-of-Interest (ROI), and anisotropic subdivision (boxes can be half-split, quarter-split or full-split producing arbitrary aspect ratio boxes).

3.33 Median trajectories

Marc van Kreveld (Utrecht University, NL)

Given a collection of more or less similar trajectories, what would constitute a "middle" trajectory? We address this problem by constructing a median trajectory from pieces of the input trajectories, while always staying in the middle. We consider two possible methods. One is simple, often fails when the trajectories have self-intersections. The other one works well and uses the concept of homotopy. We motivate our choices, give algorithms, and show some experimental results.

4 Open Problems

▶ Problem 1 (Jack Snoeyink). *Solvent Inaccessible Volume, or Nature Abhors a Vacuum:* Given a set $S = \{B_1, \ldots, B_n\}$ of n balls in 3D, give a very fast method to estimate the volume of the portion of 3-space that is inaccessible to a ball W that is allowed to be placed anywhere that it is not intersecting a ball of S. (i.e., give a very fast method to estimate the volume of an α-hull determined by n balls) The motivating application (solvent-accessibility for molecules) has balls B_i of radii in some range (1–2Å), W of radius 1.4Å, and density constraints that no pair of balls is within .8Å. One could also fix 2/3 of the atom positions, and choose from a fixed set of positions for the rest (using what is known as a rotamer library). (Danny Halperin mentions the related work he and Mark Overmars had in SoCG'94, "Spheres, Molecules, and Hidden Surface Removal".)

▶ Problem 2 (Shakhar Smorodinsky). *Coloring vertices of a simplicial complex:* Given a set P of n points in 3D and a set S of triangles spanned by P, with no two triangles having their interiors intersecting. How many colors suffice to color P so that no triangle of S is monochromatic? (This is a natural generalization of the problem of coloring the vertices of a planar (straight-line) graph, for which the answer is 4 colors always suffice.) Shakhar knows an upper bound of $O(\sqrt{n})$. The conjecture is that a constant or polylog number of colors suffice. Günter Rote asked if the special case of a simplicial complex of triangles given by a tetrahedralization of the convex hull of a set of points may be of interest (even if the points P are in convex position).

▶ Problem 3 (Janos Pach). We define the *obstacle number*, $f(G)$, of a graph $G = (V, E)$ as the minimum number of obstacles needed in a "representation" of G as a visibility graph of a set of $|V|$ points in the plane. Let $F(n)$ be the maximum obstacle number of graphs having n vertices. Clearly, $F(n) \leq \binom{n}{2}$. Janos knows that $F(n) \geq \Omega(n/\log n)$. He asks if there exists an n-vertex graph G having $f(G) = \Omega(n)$. He conjectures that "most" graphs (e.g., random graphs, expanders) have nearly quadratic obstacle number.

▶ Problem 4 (Janos Pach, and Gabor Tardos). Is it true that for $S = \{p_1, \ldots, p_n\} \subset \Re^2$ one can always find $n/2$ points not necessarily in S such that every axis-parallel rectangle that avoids T contains at most 1000 points of S? (i.e., are $n/2$ points enough to stab every rectangle containing at least 1000 elements of S?)

▶ Problem 5 (Suresh Venkatasubramanian). (Shout out to cstheory.stackexchange.com) Given a set S of n points in 2D. Consider the complete graph whose edge weights are the squared Euclidean distances between pairs of points of S. Is it always possible to find a cut whose weight is at least 2/3 of the total edge weight? (Note that it is not possible in 3D.)

▶ Problem 6 (Ferran Hurtado (6.1)). Let $S = \{p_1, \ldots, p_n\}$ be a set of n "blue" (not white!) points in 2D. Let W be a set of "red" points in 2D; these are "witness" points. Define the "witness Delaunay graph", $WDG(S, W)$ of S with respect to W that joins blue points $p_i \in S$ and $p_j \in S$ with an edge if and only if there exists a circle through p_i and p_j whose interior has no red points of W. Note that $WDG(S, \emptyset) = K_n$ and $WDG(S, S) = Del(S)$, the usual Delaunay graph of S. Let $f(n)$ be the number of witness points W that always suffice, and are sometimes necessary, to make $WDG(S, W) = \emptyset$, for an n-element set S (equivalently, no pair of blue points have adjacent regions in $Vor(S \cup W)$, as they have been "surrounded" by the red witnesses). Note that it always suffice to place a red point interior to each Delaunay edge of $Del(S)$ in order to "kill" all edges and make $WDG(S, W)$ empty; thus, $f(n) \leq 3n$. Better bounds are known: $n \leq f(n) \leq (3/2)n$, where the lower bound is from (6.1) and the upper bound from (6.2). Close the gap! The authors of (6.2) conjecture that $f(n) = n$. Also, what can be said algorithmically about computing a smallest set W to "kill" all Delaunay edges of points S?

(6.1) Witness (Delaunay) Graphs. B. Aronov, M. Dulieu, F. Hurtado. Computational Geometry: Theory and Applications, Volume 44, Issues 6-7, August 2011, Pages 329–344.
(6.2) O. Aichholzer, R. Fabila-Monroy, T. Hackl, M. van Kreveld, A. Pilz, P. Ramos, B. Vogtenhuber, Blocking Delaunay triangulations, in: Proc. 22nd Annual Canadian Conference on Computational Geometry CCCG 2010, Winnipeg, Manitoba, Canada, 2010, pp. 21–24.

▶ Problem 7 (Boris Aronov). (Posed previously, but still open.) Let K be a convex body in 3D. Consider n translates of K, $K_1, \ldots, K_n$. What is the maximum possible number, $c(n)$, of connected components of the set, $(K_1 \cup \cdots \cup K_n)^c$, the complement of the union of the translates? He knows that $\Omega(n^2) \leq c(n) \leq \binom{n}{3}$ and conjectures that $c(n) = \Theta(n^2)$. (Some special classes of convex bodies have a linear number of components in the complement of the union.)

▶ Problem 8 (Boris Aronov). Let S be a set of n blue points in 2D, and let W be a set of n red points in 2D. Let D_{pq} be the diametrical disk determined by $p, q \in S$. We want to compute all m pairs (p, q) for which $D_{pq} \cap W \neq \emptyset$. How efficiently can it be done? ($O(n^2 \log n)$, and possibly $O(n^2)$, is easy) Give an output-sensitive algorithm, e.g., taking time $O(m + n \log n)$.

▶ Problem 9 (Joe Mitchell). The following problem arises in the "Last Byte" column of Peter Winkler in CACM in 2010. Given n points in 2D, can they always be covered by n disjoint unit disks ("pennies")? The answer is "yes" for $n = 10$ (by a simple probabilistic method argument), for $n = 11$, and for $n = 12$. It is not known for $n = 13$. Janos mentions that there is a paper (by physicists) showing that for $n \geq 50$ there are configurations of points for which coverage by n disjoint disks is not possible. There remains a gap between 12 and 50.

▶ Problem 10 (Pankaj Agarwal). Let P be a set of n points in $\Re^2$, and let D be a set of m disks in $\Re^2$. We are given a positive integer k. Our goal is to pick a subset $R \subseteq P$ of k points in order to maximize the number of disks that are "double-stabbed" (i.e., that contain at least 2 points of R). THe problem is known to be NP-hard. No nontrivial approximation is known. (The related problem of minimizing the size of R in order that *every* disk is double-stabbed has an $O(\log OPT)$-approximation, by Chekuri, Clarkson, and Har Peled. The abstract set version is known not to have any nontrivial approximation algorithm.)

Participants

- Pankaj Kumar Agarwal
Duke University, US
- Oswin Aichholzer
TU Graz, AT
- Helmut Alt
FU Berlin, DE
- Victor Alvarez
Universität des Saarlandes, DE
- Lars Arge
Aarhus University, DK
- Boris Aronov
Polytech. Univ. – New York, US
- Tetsuo Asano
JAIST – Nomi, JP
- Dominique Attali
GRIPSA Lab – Saint Martin d'Héres, FR
- Erin Moriarty Wolf Chambers
St. Louis University, US
- Bernard Chazelle
Princeton University, US
- Otfried Cheong
KAIST – Daejeon, KR
- Kenneth L. Clarkson
IBM Almaden Center – San José, US
- Erik D. Demaine
MIT – Cambridge, US
- Olivier Devillers
INRIA Sophia Antipolis, FR
- Herbert Edelsbrunner
IST Austria – Klosterneuburg, AT
- Alon Efrat
Univ. of Arizona – Tucson, US
- Esther Ezra
New York University, US
- Sándor Fekete
TU Braunschweig, DE
- Jie Gao
SUNY – Stony Brook, US
- Joachim Giesen
Universität Jena, DE
- Xavier Goaoc
INRIA Lorraine, FR
- Dan Halperin
Tel Aviv University, IL
- Ferran Hurtado
UPC – Barcelona Tech, ES
- Michael Kerber
IST Austria – Klosterneuburg, AT
- David G. Kirkpatrick
University of British Columbia – Vancouver, CA
- Rolf Klein
Universität Bonn, DE
- Christian Knauer
Universität Bayreuth, DE
- Sylvain Lazard
INRIA – Nancy, FR
- Kurt Mehlhorn
MPI für Informatik – Saarbrücken, DE
- Joseph S. Mitchell
SUNY - Stony Brook, US
- Thomas Molhave
Duke University, US
- Nabil Hassan Mustafa
LUMS – Lahore, PK
- Janos Pach
EPFL – Lausanne, CH
- Richard Pollack
New York University, US
- Saurabh Ray
EPFL – Lausanne, CH
- Günter Rote
FU Berlin, DE
- Jörg-Rüdiger Sack
Carleton Univ. – Ottawa, CA
- Michael Sagraloff
MPI für Informatik - Saarbrücken, DE
- Raimund Seidel
Universität des Saarlandes, DE
- Micha Sharir
Tel Aviv University, IL
- Jonathan Shewchuk
University of California – Berkeley, US
- Shakhar Smorodinsky
Ben Gurion University – Beer Sheva, IL
- Jack Snoeyink
University of North Carolina – Chapel Hill, US
- Bettina Speckmann
TU Eindhoven, NL
- Ileana Streinu
Smith Coll. – Northampton, US
- Monique Teillaud
INRIA Sophia Antipolis, FR
- Marc van Kreveld
Utrecht University, NL
- Suresh Venkatasubramanian
University of Utah – Salt Lake City, US
- Emo Welzl
ETH Zürich, CH
- Carola Wenk
University of Texas at San Antonio, US
- Nicola Wolpert
Hochschule für Technik – Stuttgart, DE
- Chee K. Yap
New York University, US
- Ke Yi
The Hong Kong University of Science & Technology, HK

Report from Dagstuhl Seminar 11121

Computational Complexity of Discrete Problems

Edited by

Martin Grohe[1], Michal Koucký[2], Rüdiger Reischuk[3], and Dieter van Melkebeek[4]

1 **HU Berlin, DE,** grohe@informatik.hu-berlin.de
2 **Academy of Sciences - Prague, CZ,** koucky@math.cas.cz
3 **Universität zu Lübeck, DE,** reischuk@tcs.uni-luebeck.de
4 **University of Wisconsin – Madison, US,** dieter@cs.wisc.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11121 "Computational Complexity of Discrete Problems". The first section gives an overview of the topics covered and the organization of the meeting. Section 2 lists the talks given in chronological order. The last section contains the abstracts of the talks.

Seminar 20.–25. March, 2011 – www.dagstuhl.de/11121
1998 ACM Subject Classification F.1.3 Complexity Measures and Classes, F.2 Analysis of Algorithms and Problem Complexity
Keywords and phrases discrete problems, computational complexity, Turing machines, Boolean circuits, quantum computing, communication and query complexity, extractors, pseudorandomness, derandomization, approximation, coding cryptography, algorithmic learning
Digital Object Identifier 10.4230/DagRep.1.3.42
Edited in cooperation with Michael Elberfeld

1 Executive Summary

Martin Grohe
Michal Koucký
Rüdiger Reischuk
Dieter van Melkebeek

Introduction and Goals

Computational models like Turing machines and Boolean circuits work on discrete input data. Even quantum computation and communication studied in the recent past are mainly applied to solve discrete problems. Analysing the computational complexity of such problems with respect to these models is one of the central topics in the theory of computation. Researchers try to classify algorithmic problems according to complexity measures like time and space – both in the uniform and in the nonuniform setting. A variety of specialized computational models have been developed in order to better measure the complexity of certain classes of discrete problems.

Randomness has turned out to be another fundamental measure and added a lot of new intricate questions. Performing probabilistic choices within an algorithm one can design solution strategies for a given computational problem for which there are no obvious deterministic ones. Recently, large effort has been taken to remove randomness from probabilistic algorithms, so called derandomization. Here one tries to develop general techniques that can be applied to a wide range of discrete problems.

Information transfer is investigated according to the amount of communication necessary in different scenarios like 1-way channels or a bounded number of communication rounds. This is a basis for the design of efficient communication protocols. Furthermore, it has been observed that often ordinary computational problems given to a specific computational device can formally be analysed elegantly by concentrating on information flow aspects.

In addition, other computational processes arising in diverse areas of computer science have been studied, each with their own relevant complexity measures. Several of those investigations have evolved into substantial research areas, including:

- approximability (motivated by optimization),
- computational learning theory (motivated by artificial intelligence),
- query complexity (motivated by databases).

The analysis and relative power of basic models of computation remains a major challenge. New lower bound techniques for explicitly defined functions have brought the field a major step forward. For example, close connections have been discovered between circuit lower bounds for certain uniform complexity classes and the existence of pseudorandom generators and the possibility of efficient derandomization.

The seminar "Computational Complexity of Discrete Problems" has evolved out of the series of seminars entitled "Complexity of Boolean Functions," a topic that has been covered at Dagstuhl on a regular basis since the foundation of this research center. Over the years, the focus on nonuniform models has broadened to include uniform ones as well.

A salient feature of the current research in computational complexity is the interpenetration of ideas from different subareas of computational complexity and from other fields in computer science and mathematics. By organizing a generic seminar on computational complexity we have aimed to attract researchers from those various subareas and foster further fruitful interactions.

Organization of the Meeting

47 researchers from all over the world including a good number of young scientists met in Dagstuhl for this seminar. Every day we started with a longer survey talk on recent advances in specific topics that had been selected in advance by the organizers. We thank our colleagues who agreed to prepare and give these presentations:

— Prahladh Harsha took the duty for *threshold functions,*
— Paul Beame for AC^0 *circuits,*
— Troy Lee for *communication complexity,*
— Eli Ben-Sasson for *extractors,* and
— Robin Moser for *constraint satisfaction problems.*

The surveys were followed by shorter talks on new results obtained by the participants. We could schedule 30 of such plenary talks such that enough additional time was left for discussions in smaller groups on a spontaneous basis. In addition, Tuesday evening was devoted to a rump session where everybody could present his favourite open problem.

The first evening the participants could also extend their knowledge and taste in a completely different area, namely arts. We took part in the vernisage of the art exhibit by Irene Zaharoff who presented her colorful paintings in the corridors of the new building.

Everybody got so excited that the spontaneous idea to select one of her paintings and support a donation to Dagstuhl was implemented on the spot.

Topics and Achievements

We shortly review the main topics that have been discussed during the meeting. Further details as well as additional material can be found in the abstracts following.

Randomized Computations, Derandomization, and Testing

The complexity of randomized computations and its theoretical foundation was a major issue of the seminar. It showed up in about half of all contributions. Eric Allender discussed the computational power of Kolmogorov-random strings, while Andrej Bogdanov showed how to construct pseudorandom generators for read-once formulas and Pavel Pudlák for read-once permutation branching programs. Improved pseudorandom generators for a special class of discrete functions called combinatorial checkerboards were given by Thomas Watson.

For randomness extractors improved constructions were presented by Eli Ben-Sasson for the case of two sources and by Xin Li for three sources. For restricted sources that are generated by circuits of constant depth extractors were designed by Emanuele Viola.

Markus Bläser discussed the randomized complexity of identity tests for sparse polynomials, while Beate Bollig did this for integer multiplication in the OBDD model generalizing her deterministic lower bound shown at the previous meeting.

Derandomization techniques were presented by Matthew Anderson for zero tests of multilinear arithmetic formulas and by Robin Moser for Schöning's satisfiability test of k-CNF formulas. Eldar Fischer considered property testing of monotone formulas.

Communication Complexity

How many bits two parties have to exchange in order to compute a given function if the input is distributed among them? This is a fundamental question for the design of communication protocols. To determine the Hamming distance of two n-bit strings it has been well known that the trivial solution of one party sending his string to the other party is optimal. It was open for quite a while if this still holds if the two parties get the additional information that the distance of their inputs is either small or large (let's say not in the region between $[n/2 - \sqrt{n}, n/2 + \sqrt{n}]$). Oded Regev could resolve this question for general probabilistic protocols with unbounded rounds of communication by proving a linear lower bound.

Allowing the two parties also to exchange quantum bits leads to quantum protocols which – in contrast to quantum computers – are already used in practice. There are examples known that quantum bits can lead to an exponential decrease of the amount of communication. Oded Regev in his second contribution showed that such a separation can even be obtained when comparing 1-way quantum protocols with 2-way classical protocols.

Troy Lee investigated the query complexity of quantum states and a generalization of this problem called state conversion. He defined a new norm for the distance of quantum states and proved that this gives an appropriate measure.

Andrew Drucker showed how fault-tolerant protocols can be used to improve upon more complex objects called probabilistically checkable debate systems. This has implications for the approximability of problems in PSPACE.

Complexity Classes

For classical complexity classes some progress was reported, too. Michael Elberfeld showed how graph problems restricted to instances of bounded tree width can be solved in logarithmic space, while Fabian Wagner was able to solve the isomorphism problem for such instances in LogCFL. Meena Mahajan presented a detailed investigation of arithmetic circuits of logarithmic depth and their relation to logspace.

For Boolean circuits the tradeoff between size and depth has been investigated from the very beginning. Anna Gál showed with the help of the pebble game how the bounds can be substantially improved for layered circuits. Or Meir took a closer lock at the breakthrough result IP=PSPACE and showed an alternate proof that uses fairly general error-correcting codes instead of polynomials.

Another famous result that unbounded circuits of constant depth cannot compute the parity function was taken up by Paul Beame and Johan Håstad trying to get a finer quantitative statement of this property. Paul described a simulation of such circuits by decision trees which has implications how well AC^0 can approximate the parity function, while Johan presented a direct proof for an upper bound on the correlation between parity and functions computable in constant depth. Nicole Schweikardt gave a precise characterization of the locality of order-invariant first-order queries with arbitrary predicates, which are closely related to the complexity class AC^0.

Srikanth Srinivasan showed that computing the determinant over simple noncommutative rings is as hard as computing the permanent in the commutative case, thus establishing a huge complexity gap between commutative and noncommutative domains.

Further Topics

Improving the construction of good error correcting codes by combining classical codes with outer codes was adressed by Amnon Ta-Shma. Matthias Krause presented a new technique to prove security properties of cryptographic hash functions. Philipp Woelfel considered random walks on a line towards a target where the searcher can fix an arbitrary distribution on his probabilistically chosen next movement. A lower bound on the first hitting time was shown matching previously known upper bounds for this problem. Jakob Nordström considered linear invariance properties in the realm of property testing. He investigated the problem to decide on the semantic difference for syntactically different descriptions of linear invariances.

Efficiently learning unknown concepts by queries or taking random samples is another topic of general applicability. An important measure in this respect is the VC-dimension which implies bounds on the minimal achievable additive error for any learning algorithm. Ilan Newman asked the same question for the multiplicative error and showed that the triangular rank of a set system gives a corresponding measure. Kristoffer Hansen considered the class of Boolean functions with a constant degree representation where each variable is used only a bounded number of times. He presented a deterministic polynomial time algorithm for this problem. For Boolean functions the notion of sensivity has turned to be an important measure although it seems to be hard to estimate in many cases. Prahladh Harsha considered polynomial threshold functions and showed the first nontrivial upper bound, which has also implications for the learning complexity of these functions.

Conclusion

Investigating the complexity of discrete problems is one of the fundamental tasks in the theory of computation. On the one hand, new algorithmic techniques and new ways to look at a problem have led to better algorithms and protocols. On the other hand, typically more demanding is the task to prove lower bounds on the computational complexity of a concrete problem. Progress is still continuing, as seen for example in testing, derandomization and explicit constructions of combinatorial objects like extractors, that improves our knowledge considerably. Despite these significant steps forward that have been achieved in several subareas since our previous meeting three years ago, the general feeling among the participants was that we still have to work hard for many more years to get a good understanding what are the limits of efficient computation.

We like to thank the staff at Dagstuhl who – as usual – provided a marvellous surrounding to make this a successful meeting with ample space for undisturbed interactions between the participants.

2 Table of Contents

Executive Summary
Martin Grohe, Michal Koucký, Rüdiger Reischuk, and Dieter van Melkebeek 42

Overview of Talks

A Survey of Recent Advances in Threshold Functions
Prahladh Harsha 49

Bounding the Average and Noise Sensitivity of Polynomial Threshold Functions
Prahladh Harsha 49

Isomorphism and Canonization of Bounded Treewidth Graphs
Fabian Wagner 50

Counting Classes and the Fine Structure between NC^1 and DLOG
Meena Mahajan 50

The Hardness of the Noncommutative Determinant
Srikanth Srinivasan 51

IP = PSPACE using Error Correcting Codes
Or Meir 51

Efficient Probabilistically Checkable Debates
Andrew Drucker 51

Limits on the Computational Power of Random Strings
Eric Allender 52

How well do AC^0 Circuits Approximate Parity? Approximating AC^0 by "Small" Height Decision Trees
Paul Beame 53

Randomized OBDDs for the Most Significant Bit of Multiplication Need Exponential Size
Beate Bollig 53

The Size and Depth of Layered Boolean Circuits
Anna Gál 54

Preimage Resistance Beyond the Birthday Barrier – The Case of Blockcipher Based Hashing
Matthias Krause 54

Pseudorandom Generators for Group Products
Pavel Pudlák 55

Pseudorandom Generators for Combinatorial Checkerboards
Thomas W. Watson 55

Pseudorandomness for Read-Once Formulas
Andrej Bogdanov 55

Testing Assignments for Satisfying a Monotone Formula
Eldar Fischer 56

On the Semantics of Local Characterizations for Linear-Invariant Properties
Jakob Nordström 56

Quantum Query Complexity of State Generation
Troy Lee 57

Quantum One-Way Communication can be Exponentially Stronger Than Classical Communication
Oded Regev 58

An Optimal Lower Bound on the Communication Complexity of Gap-Hamming-Distance
Oded Regev 58

From Affine to Two-Source Extractors via Approximate Duality
Eli Ben-Sasson 59

Improved Constructions of Three Source Extractors
Xin Li 59

Extractors for Circuit Sources
Emanuele Viola 59

What Binary Codes can be Obtained by Concatenating AG Codes with Hadamard?
Amnon Ta-Shma 60

Derandomizing Polynomial Identity Testing for Multilinear Constant-Read Formulae
Matthew Anderson 60

Randomness Efficient Testing of Sparse Blackbox Identities of Unbounded Degree over the Reals
Markus Bläser 61

Learning Read-Constant Polynomials of Constant Degree Modulo Composites
Kristoffer Arnsfelt Hansen 61

Locality of AC^0-Computable Graph Queries
Nicole Schweikardt 62

A Survey of Exponential Algorithms for Constraint Satisfaction Problems
Robin Moser 62

A Full Derandomization of Schöning's k-SAT Algorithm
Robin Moser 63

Algorithmic Meta Theorems Inside Logspace and Their Applications
Michael Elberfeld 63

Triangular Rank and Sampling With Multiplicative Errors
Ilan Newman 63

Tight Lower Bounds for Greedy Routing in Uniform Small World Rings
Philipp Woelfel 64

On the Correlation of Parity with Small Depth Circuits
Johan Håstad 65

Participants 66

3 Overview of Talks

3.1 A Survey of Recent Advances in Threshold Functions

Prahladh Harsha (TIFR Mumbai, IN)

Threshold functions (halfspaces, intersection of halfspaces, polynomial threshold functions) commonly occur in various applications in learning theory, communication complexity etc. In the last two years, there has been a sequence of results in the area of threshold functions (eg., improved pseudorandom generators for various types of threshold functions, better bounds on certain quantitative functional parameters such as noise sensitivity, average sensitivity etc). Most of these results have been inspired by a better understanding of invariance principles from probability theory. In this talk, I'll give a survey of some these results related to threshold functions and the corresponding invariance principles.

3.2 Bounding the Average and Noise Sensitivity of Polynomial Threshold Functions

Prahladh Harsha (TIFR Mumbai, IN)

Joint work of Harsha, Prahladh; Klivans, Adaml; Meka, Raghu
Main reference In Proc. 42nd ACM Symp. on Theory of Computing (STOC) (Cambridge, Massachusetts, 6-8 June), pages 543-552, 2010
URL http://arxiv.org/abs/0909.5175

In 1994, Gotsman and Linial posed the following question: "what is the maximum number of edges of the n-dimensional Boolean hypercube cut by any degree d polynomial surface?" Generalizing from the degree one case, they conjectured that the symmetric function slicing the middle d layers of the Boolean hypercube achieves this maximum. A restating of this conjecture is that the average sensitivity of a degree d polynomial threshold function (PTF) is at most $O(d\sqrt{n})$. A closely related (and in fact, equivalent) conjecture is that the noise sensitivity of a degree d PTF (for noise rate δ) is at most $O(d\sqrt{\delta})$.

In this work, we give the first nontrivial upper bounds on the average sensitivity and noise sensitivity of polynomial threshold functions. More specifically, we show that

- The average sensitivity of a d-PTF is at most $O(n^{1-1/(4d+6)})$
- The noise sensitivity of f with noise rate δ is at most $O(\delta^{1/(4d+6)})$.

The Gotsman-Linial conjecture itself remains unresolved. Nevertheless, these bounds immediately yield (the first) polynomial time agnostic learning algorithms for the class of degree d PTFs.

3.3 Isomorphism and Canonization of Bounded Treewidth Graphs

Fabian Wagner (Universität Ulm, DE)

Main reference B. Das, J. Torán, F. Wagner: Restricted Space Algorithms for Isomorphism on Bounded Treewidth Graphs, in proceedings of STACS, pp. 227-238, 2010.
F. Wagner: Graphs of Bounded Treewidth can be Canonized in AC^1, to appear in proceedings of CSR 2011, technical report: ECCC TR-032, 2011.

In two recent results we show that isomorphism testing of bounded treewidth graphs is in LogCFL [Das-Torán-Wagner'10] and canonization is in AC^1 [Wagner'11], improving the previously known upper bounds of TC^1 [Grohe-Verbitsky'06] and TC^2 [Köbler-Verbitsky'08].

Both results extend in two different ways the techniques of canonization algorithms developed for restricted graph classes.

Most notably Lindells logspace algorithm for trees [Lindell'92] which was generalized to logspace algorithms for partial 2-trees [Arvind-Das-Köbler'08], k-trees [Köbler-Kuhnert'09], interval graphs [Köbler-Kuhnert-Laubner-Verbitsky'09], planar graphs [Datta-Limaye-Nimbhorkar-Thierauf-Wagner'09], $K_{3,3}$- and K_5-minor free graphs [Datta-Nimbhorkar-Thierauf-Wagner'09].

In the talk I will present the proof techniques from both results [Das-Torán-Wagner'10] and [Wagner'11].

References

1 B. Das, J. Toran, F. Wagner: Restricted Space Algorithms for Isomorphism on Bounded Treewidth Graphs, Proceedings of STACS'10, LIPIcs Vol. 5, pp. 227–238, 2010.
2 F. Wagner: Graphs of Bounded Treewidth can be Canonized in AC^1, to appear in proceedings of CSR 2011, technical report: ECCC TR-032, 2011.

3.4 Counting Classes and the Fine Structure between NC^1 and DLOG

Meena Mahajan (The Institute of Mathematical Sciences – Chennai, IN)

Joint work of Mahajan, Meena; Datta, Samir; Rao, B. V. Raghavendra; Thomas, Michael; Vollmer, Heribert
Main reference Datta, Samir; Mahajan, Meenal; Rao, B. V. Raghavendra; Thomas, Michael; Vollmer, Heribert; "Counting classes and the fine structure between $NC^1$and L," Proc. 35th MFCS 2010
URL http://dx.doi.org/10.1007/978-3-642-15155-2_28

The class NC^1 of problems solvable by bounded fan-in circuit families of logarithmic depth is known to be contained in logarithmic space DLOG, but not much about the converse is known. In this paper we examine the structure of classes in between NC^1 and DLOG based on counting functions or, equivalently, based on arithmetic circuits. The classes PNC^1 and CNC^1, defined by a test for positivity and a test for zero, respectively, of arithmetic circuit families of logarithmic depth, sit in this complexity interval. We study the landscape of Boolean hierarchies, constant-depth oracle hierarchies, and logarithmic-depth oracle hierarchies over PNC^1 and CNC^1. We provide complete problems, obtain the upper bound DLOG for all these hierarchies, and prove partial hierarchy collapses. In particular, the constant-depth oracle hierarchy over PNC^1 collapses to its first level PNC^1, and the constant-depth oracle hierarchy over CNC^1 collapses to its second level.

3.5 The Hardness of the Noncommutative Determinant

Srikanth Srinivasan (Institute for Advanced Study – Princeton, US)

Joint work of Chien, Steve; Harsha, Prahladh; Sinclair, Alistair; Srinivasan, Srikanth
Main reference Steve Chien, Prahladh Harsha, Alistair Sinclair, Srikanth Srinivasan; "Almost Settling the Hardness of Noncommutative Determinant," STOC 2011, to appear.
URL http://arxiv.org/abs/1101.1169

We consider the complexity of the determinant over noncommutative domains and show that even in simple noncommutative settings such as the setting of 2X2 matrices, this problem is #P-hard. This has applications to the analysis of a natural extension to the Godsil-Gutman estimator for the 0-1 permanent and to arithmetic circuit complexity.

3.6 IP = PSPACE using Error Correcting Codes

Or Meir (Weizmann Institute – Rehovot, IL)

Main reference Or Meir; "IP = PSPACE using Error Correcting Codes," ECCC TR10-137
URL http://eccc.hpi-web.de/report/2010/137/

The IP theorem, which asserts that IP = PSPACE (Lund et. al., and Shamir, in J. ACM 39(4)), is one of the major achievements of complexity theory. The known proofs of the theorem are based on the arithmetization technique, which transforms a quantified Boolean formula into a related polynomial. The intuition that underlies the use of polynomials is commonly explained by the fact that polynomials constitute good error correcting codes. However, the known proofs seem tailored to the use of polynomials, and do not generalize to arbitrary error correcting codes.

In this work, we show that the IP theorem can be proved by using general error correcting codes. We believe that this establishes a rigorous basis for the aforementioned intuition, and sheds further light on the IP theorem.

3.7 Efficient Probabilistically Checkable Debates

Andrew Drucker (MIT – Cambridge, US)

Main reference Andrew Drucker; "Efficient Probabilistically Checkable Debates," Electronic Colloquium on Computational Complexity (ECCC), TR11-073, 2011.
URL http://eccc.hpi-web.de/report/2011/073/

Probabilistically checkable debate systems (PCDSs) are debates between two competing provers, which a polynomial-time verifier inspects in $O(1)$ bits.

Condon et al. showed all PSPACE languages have PCDSs. This implies that the approximation versions of some natural PSPACE-complete problems are also PSPACE-complete.

We give an improved construction: for a language L with an ordinary debate system defined by uniform circuits of size $s = s(n)$, we give a PCDS for L with debate bitlength

s·polylog(s). This tightens the connection between the time complexities of PSPACE-complete problems and their approximation versions.

Our key ingredient is a novel application of error-resilient communication protocols (specifically, the Braverman-Rao protocol). By requiring ordinary debates to be resiliently encoded, we endow them with a useful "stability" property, which lets them be converted into PCDSs. Our main technical challenge is to enforce error-resilient encoding by the debaters.

3.8 Limits on the Computational Power of Random Strings

Eric Allender (Rutgers University – Piscataway, US)

Joint work of Allender, Eric; Friedman, Luke; Gasarch, William
Main reference Eric Allender, Luke Friedman, William Gasarch; "Limits on the Computational Power of Random Strings," ICALP, 2011
URL http://eccc.hpi-web.de/report/2010/139/

R, the set of Kolmogorov-random strings, is a central notion in the study of algorithmic information theory, and in recent years R has increasingly been studied in relation to computational complexity theory. This talk takes as its starting point three strange inclusions that have been proved since 2002: 1. NEXP is contained in the class of problems NP-Turing-reducible to R. 2. PSPACE is contained in the class of problems poly-time Turing-reducible to R. 3. BPP is contained in the class of problems poly-time truth-table-reducible to R.

(These inclusions hold for both of the most widely-studied variants of Kolmogorov complexity: the plain complexity C(x) and the prefix-complexity K(x).

They also hold no matter which "universal" Turing machine is used in the definitions of the functions C and K.)

These inclusions are "strange" since R is not even computable! Thus it is not at all clear that these are meaningful upper bounds on the complexity of BPP, PSPACE, and NEXP, and indeed it is not at all clear that it is very interesting to consider efficient reductions to noncomputable sets such as R.

In this talk, I will try to convince you that the class of problems efficiently reducible to R is, indeed, a complexity class. The main theorems are that, if we restrict attention to prefix complexity K and the corresponding set of random strings R_K, then the class of decidable problems that are in NP relative to R_K (no matter which universal machine is used to define K) lies in EXPSPACE, and the class of decidable problems that are poly-time truth-table reducible to R_K (no matter which universal machine is used to define K) lies in PSPACE.

Thus we can "sandwich" PSPACE between the class of problems truth-table- and Turing-reducible to R_K, and the class of decidable problems that are in NP relative to R_K lies between NEXP and EXPSPACE. The corresponding questions for plain Kolmogorov complexity C are wide open; no upper bounds are known at all for the class of decidable problems efficiently reducible to R_C.

These results also provide the first quantitative limits on the applicability of uniform derandomization techniques.

3.9 How well do AC^0 Circuits Approximate Parity? Approximating AC^0 by "Small" Height Decision Trees

Paul Beame (University of Washington – Seattle, US)

Joint work of Beame, Paul; Impagliazzo, Russell; Srinivasan, Srikanth

We show that for every $C = C(n) > 1$ and every family F of k-DNF formulas in n variables one can build a single decision tree T of height $h = n/C$ such that for all but a $2^{-h/(2^{O(k)} \log |F|)}$ fraction of paths p in T, every formula in F reduces to a $(2^k C \log |F|)^{O(k^3)}$-junta on assignments consistent with p. This improves a construction, due to Ajtai, which had an exponentially larger dependence on k.

As a consequence of this construction we show that every n-input AC^0 circuit of size S and depth d can be approximated by a decision tree of height at most $n - B\{S, d\}n$, where $B\{S, d\} = 2^{-2dlog_2^{4/5}} S$, that is always correct but may not produce an answer on at most a $2^{-B\{S,d\}n}$ fraction of branches. It follows that any such AC^0 circuit has correlation at most $2^{-B\{S,d\}n}$ with the n-bit Parity function.

Our proof is constructive and yields a deterministic algorithm running in time $2^{n-B\{S,d\}n} S^{O(1)}$ that exactly counts the number of satisfying assignments of any n-input AC^0 circuit of size S and depth d. Indeed, in the same running time we can deterministically construct a decision tree of size at most $2^{n-B\{S,d\}n}$ that exactly computes the function given by such a circuit.

(Our constructions more naturally yield randomized algorithms for these problems but deterministic algorithms follow because only limited independence is required in their analysis.)

3.10 Randomized OBDDs for the Most Significant Bit of Multiplication Need Exponential Size

Beate Bollig (Technische Universität Dortmund, DE)

Joint work of Bollig, Beate; Gille, Marc
Main reference Beate Bollig and Marc Gille; "Randomized OBDDs for the most significant bit of multiplication need exponential size," Information Processing Letters 111, 151-155, 2011.
URL http://dx.doi.org/10.1016/j.ipl.2010.11.013

The basic arithmetic functions have been in the middle of several complexity theoretical investigations and ordered binary decision diagrams (OBDDs) are a popular dynamic data structure for Boolean functions.

Only in 2008 it has been shown that the size of deterministic OBDDs for the most significant bit of integer multiplication is exponential.

Since probabilistic methods have turned out to be useful in almost all areas of computer science, one may ask whether randomization can help to represent the most significant bit of multiplication in smaller size.

Here, it is proved that the randomized OBDD complexity is also exponential.

3.11 The Size and Depth of Layered Boolean Circuits

Anna Gál (University of Texas at Austin, US)

Joint work of Gál, Anna; Jang, Jing-Tang

We consider the relationship between size and depth for layered Boolean circuits, synchronous circuits and planar circuits as well as classes of circuits with small separators. In particular, we show that every layered Boolean circuit of size s can be simulated by a layered Boolean circuit of depth $O(\sqrt{s \log s})$. For planar circuits and synchronous circuits of size s, we obtain simulations of depth $O(\sqrt{s})$. The best known result so far was by Paterson and Valiant, and Dymond and Tompa, which holds for general Boolean circuits and states that $D(f) = O(C(f)/\log C(f))$, where $C(f)$ and $D(f)$ are the minimum size and depth, respectively, of Boolean circuits computing f. The proof of our main result uses an adaptive strategy based on the two-person pebble game introduced by Dymond and Tompa. Improving any of our results by polylog factors would immediately improve the bounds for general circuits.

3.12 Preimage Resistance Beyond the Birthday Barrier – The Case of Blockcipher Based Hashing

Matthias Krause (Universität Mannheim, DE)

Joint work of Krause, Matthias; Armknecht, Frederik; Fleischmann, Ewan

Security proofs are an essential part of modern cryptography. Often the challenge is not to come up with appropriate schemes but rather to technically prove that these satisfy the desired security properties. We provide for the first time techniques for proving asymptotically optimal preimage resistance bounds for block cipher based double length, double call hash functions. More precisely, we consider for some blocklength n compression functions H from $\{0,1\}^{3n} \to \{0,1\}^{2n}$ using two calls to an ideal block cipher with an n-bit block size. Optimally, an adversary trying to find a preimage for H should require ?(2^{2n}) queries to the underlying block cipher. As a matter of fact there have been several attempts to prove the preimage resistance of such compression functions, but no proof did go beyond the ?(2^n) "birthday" barrier, therefore leaving a huge gap when compared to the optimal bound.

In this paper, we introduce two new techniques on how to lift this bound to ?(2^{2n}). We demonstrate our new techniques for a simple and natural design of H, being the concatenation of two instances of the well-known Davies-Meyer compression function.

3.13 Pseudorandom Generators for Group Products

Pavel Pudlák (Acad. of Sciences – Prague, CZ)

Joint work of Koucký, Michal; Nimbhorkar, Prajakta; Pudlák, Pavel
Main reference M. Koucký, P. Nimbhorkar, P. Pudlák; "Pseudorandom Generators for Group Products," Proc. of the 2011 ACM STOC.
URL http://www.math.cas.cz/ pudlak/version-g.pdf

We proved that the pseudorandom generator introduced by Impagliazzo et al. in 1994 with proper choice of parameters fools group products of a given finite group G. The seed length is $O(|G|^{O(1)} \cdot \log n + \log n/\delta)$, where n is the length of the word and δ is the allowed error. The result is equivalent to the statement that the pseudorandom generator with seed length $O(2^{O(w \log w)} \cdot \log n + \log n/\delta)$ fools read-once permutation branching programs of width w.

In the lecture I give a short sketch of the main ideas focusing on a geometric lemma that is one of the main technical parts.

3.14 Pseudorandom Generators for Combinatorial Checkerboards

Thomas W. Watson (University of California – Berkeley, US)

We define a combinatorial checkerboard to be a function $f : \{1, \ldots, m\}^d \to \{1, -1\}$ of the form $f(u_1, \ldots, u_d) = \prod_{i=1}^{d} f_i(u_i)$ for some functions $f_i : \{1, \ldots, m\} \to \{1, -1\}$. This is a variant of combinatorial rectangles, which can be defined in the same way but using $\{0, 1\}$ instead of $\{1, -1\}$.

We consider the problem of constructing explicit pseudorandom generators for combinatorial checkerboards. This is a generalization of small-bias generators, which correspond to the case $m = 2$.

We construct a pseudorandom generator that ϵ-fools all combinatorial checkerboards with seed length $O\big(\log m + \log d \cdot \log\log d + \log^{3/2} \frac{1}{\epsilon}\big)$. Previous work by Impagliazzo, Nisan, and Wigderson implies a pseudorandom generator with seed length $O\big(\log m + \log^2 d + \log d \cdot \log \frac{1}{\epsilon}\big)$. Our seed length is better except when $\frac{1}{\epsilon} \geq d^{\omega(\log d)}$.

3.15 Pseudorandomness for Read-Once Formulas

Andrej Bogdanov (Chinese University of Hong Kong, HK)

Joint work of Bogdanov, Andrej; Papakonstantinou, Periklis; Wan, Andrew
Main reference manuscript (unpublished)

We give an explicit construction of a pseudorandom generator for read-once formulas whose inputs can be read in arbitrary order. For formulas in n inputs and arbitrary gates of fan-in at most $d = O(n/\log n)$, the pseudorandom generator uses $(1 - \Omega(1))n$ bits of randomness and produces an output that looks $2^{-\Omega(n)}$-pseudorandom to all such formulas.

Our analysis is based on the following lemma. Let $pr = Mz + e$, where M is the parity-check matrix of a sufficiently good binary error-correcting code of constant rate, z is a random string, e is a small-bias distribution, and all operations are modulo 2. Then for every pair of functions $f, g\colon B^{n/2} \to B$ and every equipartition (I, J) of $[n]$, the distribution pr is pseudorandom for the pair $(f(x|_I), g(x|_J))$, where $x|_I$ and $x|_J$ denote the restriction of x to the coordinates in I and J, respectively.

3.16 Testing Assignments for Satisfying a Monotone Formula

Eldar Fischer (Technion – Haifa, IL)

Joint work of Fischer, Eldar; Lachish, Oded; Nimbhorkar, Prajakta

Property Testing deals with the following question: Distinguish using as few queries to the input as possible, ideally with a number of queries independent of the input size, between inputs that satisfy a given property and inputs that are far from any possible input satisfying the property. In the massively parametrized model, a fixed part of the input is fully given to the algorithm in advance, on which the algorithm has to be exact (i.e. the approximation of "not being far from a satisfying input" can only be made for the input not given in advance).

In this talk we consider properties that relate to tree-like structures that are given in advance, and in particular to read-once monotone Boolean formulas. Such formulas are representable as trees with And/Or labels on their nodes, and the part of the input not given in advance is the assignment, i.e. the values at the leaves.

The main result is a test for the property of an assignment satisfying the given formula, whose number of queries does not depend on the input size at all. We also discuss some related questions, such as making the test tolerant, or alternatively reducing its number of queries to be quasi-polynomial in the approximation parameter.

3.17 On the Semantics of Local Characterizations for Linear-Invariant Properties

Jakob Nordström (KTH Stockholm, SE)

Joint work of Bhattacharyya, Arnab; Grigorescu, Elena; Nordström, Jakob; Xie, Ning
Main reference Arnab Bhattacharyya, Elena Grigorescu, Jakob Nordström, Ning Xie; "Separations of Matroid Freeness Properties," ECCC TR10-136
URL http://eccc.hpi-web.de/report/2010/136/

A property of functions on a vector space is said to be linear-invariant if it is closed under linear transformations of the domain. Linear-invariant properties are some of the most well-studied properties in the field of property testing. Testable linear-invariant properties can always be characterized by so-called local constraints, and of late there has been a rapidly developing body of research investigating the testability of linear-invariant properties in terms of their descriptions using such local constraints. One problematic aspect that has been largely ignored in this line of research, however, is that syntactically distinct local characterizations need not at all correspond to semantically distinct properties. In fact, there are known fairly dramatic examples where seemingly infinite families of properties collapse into a small finite set that was already well-understood.

In this work, we therefore initiate a systematic study of the semantics of local characterizations of linear-invariant properties. For such properties the local characterizations have an especially nice structure in terms of forbidden patterns on linearly dependent sets of vectors, which can be encoded formally as matroid constraints. We develop techniques for determining, given two such matroid constraints, whether these constraints encode identical or distinct properties, and show for a fairly broad class of properties that these techniques provide necessary and sufficient conditions for deciding between the two cases. We use these tools to show that recent (syntactic) testability results indeed provide an infinite number of infinite strict hierarchies of (semantically) distinct testable locally characterized linear-invariant properties.

3.18 Quantum Query Complexity of State Generation

Troy Lee (National University of Singapore, SG)

Joint work of Lee, Troy; Mittal, Rajat; Reichardt, Ben; Spalek, Robert; Szegedy, Mario
Main reference Troy Lee, Rajat Mittal, Ben W. Reichardt, Robert Spalek; "An adversary for algorithms," arXiv:1011.3020v1
URL http://arxiv.org/abs/1011.3020

State-conversion generalizes query complexity to the problem of converting between two input-dependent quantum states by making queries to the input. We characterize the complexity of this problem by introducing a natural information-theoretic norm that extends the Schur product operator norm. The complexity of converting between two systems of states is given by the distance between them, as measured by this norm.

In the special case of function evaluation, the norm is closely related to the general adversary bound, a semi-definite program that lower-bounds the number of input queries needed by a quantum algorithm to evaluate a function. We thus obtain that the general adversary bound characterizes the quantum query complexity of any function whatsoever. This generalizes and simplifies the proof of the same result in the case of Boolean input and output. Also in the case of function evaluation, we show that our norm satisfies a remarkable composition property, implying that the quantum query complexity of the composition of two functions is at most the product of the query complexities of the functions, up to a constant. Finally, our result implies that discrete and continuous-time query models are equivalent in the bounded-error setting, even for the general state-conversion problem.

3.19 Quantum One-Way Communication can be Exponentially Stronger Than Classical Communication

Oded Regev (Tel Aviv University, IL)

Joint work of Regev, Oded; Klartag, Bo'az
Main reference Bo'az Klartag, Oded Regev, "Quantum One-Way Communication is Exponentially Stronger Than Classical Communication," STOC 2011, to appear.
URL http://eccc.hpi-web.de/report/2010/143/

In STOC 1999, Raz presented a (partial) function for which there is a quantum protocol communicating only $O(\log n)$ qubits, but for which any classical (randomized, bounded-error) protocol requires poly(n) bits of communication.

That quantum protocol requires two rounds of communication. Ever since Raz's paper it was open whether the same exponential separation can be achieved with a quantum protocol that uses only one round of communication. In other words, can quantum one-way communication be exponentially stronger than classical two-way communication? Here we settle this question in the affirmative.

Note: This talk is about lower bounds for *classical* communication complexity; no knowledge of quantum communication complexity is assumed or required.

3.20 An Optimal Lower Bound on the Communication Complexity of Gap-Hamming-Distance

Oded Regev (Tel Aviv University, IL)

Main reference Amit Chakrabarti, Oded Regev; "An Optimal Lower Bound on the Communication Complexity of Gap-Hamming-Distance," STOC 2011, to appear.
URL http://arxiv.org/abs/1009.3460

We consider the Gap Hamming Distance problem in communication complexity. Here, Alice receives an n-bit string x, and Bob receives an n-bit string y.

They are promised that the Hamming distance between x and y is either at least $n/2+\sqrt{n}$ or at most $n/2-\sqrt{n}$, and their goal is to decide which is the case.

The naive protocol requires n bits of communication and it was an open question whether this is optimal. This was shown in several special cases, e.g., when the communication is deterministic [Woodruff'07] or when the number of rounds of communication is limited [Indyk-Woodruff'03, Jayram-Kumar-Sivakumar'07, Brody-Chakrabarti'09, Brody-Chakrabarti-Regev-Vidick-deWolf'09].

Here we settle this question by showing a tight lower bound of $\Omega(n)$ on the randomized communication complexity of the problem. The bound is based on a new geometric statement regarding correlations in Gaussian space, related to a result of C. Borell from 1985, which is proven using properties of projections of sets in Gaussian space.

3.21 From Affine to Two-Source Extractors via Approximate Duality

Eli Ben-Sasson (Technion – Haifa, IL)

Joint work of Ben-Sasson, Eli; Zewi, Noga
Main reference Eli Ben-Sasson, Noga Zewi ; "From Affine to Two-Source Extractors via Approximate Duality," Proceedings of STOC 2011.
URL http://eccc.hpi-web.de/report/2010/144/

Two-source and affine extractors and dispersers are fundamental objects studied in the context of derandomization. (Two-source dispersers are equivalent to bipartite Ramsey graphs.) We show how to construct two-source extractors and dispersers (i.e., bipartite Ramsey graphs) for arbitrarily small min-entropy rate in a black-box manner from affine extractors with sufficiently good parameters. Our analysis relies on the study of approximate duality, a concept related to the polynomial Freiman-Ruzsa conjecture from additive combinatorics.

3.22 Improved Constructions of Three Source Extractors

Xin Li (University of Texas – Austin, US)

Main reference Xin Li; "Improved Constructions of Three Source Extractors," CCC 2011
URL http://eccc.hpi-web.de/report/2010/190/

This talk presents recent constructions of extractors for three independent weak random sources on n bits with min-entropy $n^{1/2+\alpha}$, for any arbitrary constant $\alpha > 0$. This improves the previous best result where at least one source is required to have min-entropy $n^{0.9}$. Other results include extractors for three independent sources with uneven lengths and extractors for two independent affine sources on n bits with entropy $n^{1/2+\alpha}$.

3.23 Extractors for Circuit Sources

Emanuele Viola (Northeastern Univ. – Boston, US)

Main reference Emanuele Viola; "Extractors for circuit sources," ECCC Technical Report 2011-056
URL http://www.eccc.uni-trier.de/report/2011/056/

We obtain the first deterministic extractors for sources generated (or sampled) by small circuits of bounded depth. Our main results are:

(1) We extract $k(k/nd)^{O(1)}$ bits with exponentially small error from n-bit sources of min-entropy k that are generated by functions $f : \{0,1\}^\ell \to \{0,1\}^n$ where each output bit depends on $\leq d$ input bits. In particular, we extract from NC^0 sources, corresponding to $d = O(1)$.

(2) We extract $k(k/n^{1+\gamma})^{O(1)}$ bits with super-polynomially small error from n-bit sources of min-entropy k that are generated by poly(n)-size AC^0 circuits, for any $\gamma > 0$.

As our starting point, we revisit the connection by Trevisan and Vadhan (FOCS 2000) between circuit lower bounds and extractors for sources generated by circuits. We note that such extractors (with very weak parameters) are equivalent to lower bounds for generating distributions (FOCS 2010; with Lovett, CCC 2011).

Building on those bounds, we prove that the sources in (1) and (2) are (close to) a convex combination of high-entropy "bit-block" sources. Introduced here, such sources are a special case of affine ones. As extractors for (1) and (2) one can use the extractor for low-weight affine sources by Rao (CCC 2009).

Along the way, we exhibit an explicit Boolean function $b : \{0,1\}^n \to \{0,1\}$ such that poly(n)-size AC^0 circuits cannot generate the distribution $(x, b(x))$, solving a problem about the complexity of distributions.

Independently, De and Watson (ECCC TR11-037) obtain a result similar to (1) in the special case $d = o(\lg n)$.

3.24 What Binary Codes can be Obtained by Concatenating AG Codes with Hadamard?

Amnon Ta-Shma (Tel Aviv University, IL)

Joint work of Ta-Shma, Amnon; Ben-Aroya, Avraham
Main reference Avraham Ben-Aroya, Amnon Ta-Shma; "Constructing Small-Bias Sets from Algebraic-Geometric Codes", pp. 191–197, FOCS 2009
URL http://dx.doi.org/10.1109/FOCS.2009.44

The currently best explicit constructions of binary error correcting codes of distance close to half work by concatenating a good outer code with Hadamard. When the outer code is RS the obtained code has length $n = O(k^2/eps^2)$, while taking an AG code of degree larger than the genus gives $n = O(k/eps^3)$, where k is the code dimension, $1/2 - eps$ is the distance and the O notation also hides logarithmic factors.

I show that the result can be improved by taking AG codes with degree *smaller* than the genus. Specifically, we obtain codes of length $n = O(k/eps^2)^{5/4}$ which improves upon previous explicit constructions when eps is roughly (ignoring logarithmic factors) in the range $[k^{-1.5}, k^{-0.5}]$. We also discuss an argument by Voloch exposing the limitations of the technique.

3.25 Derandomizing Polynomial Identity Testing for Multilinear Constant-Read Formulae

Matthew Anderson (University of Wisconsin – Madison, US)

Joint work of Anderson, Matthew; van Melkebeek, Dieter; Volkovich, Ilya
Main reference M. Anderson, D. van Melkebeek, and I. Volkovich; "Derandomizing Polynomial Identity Testing for Multilinear Constant-Read Formulae," Electronic Colloquium on Computational Complexity, Technical Report ECCC-TR 10-188, 2010.
URL http://eccc.hpi-web.de/report/2010/188/

We present a polynomial-time deterministic algorithm for testing whether constant-read multilinear arithmetic formulae are identically zero. In such a formula each variable occurs only a constant number of times and each subformula computes a multilinear polynomial. Our algorithm runs in time $s^{O(1)} \cdot n^{k^{O(k)}}$, where s denotes the size of the formula, n denotes the number of variables, and k bounds the number of occurrences of each variable. Before our work no subexponential-time deterministic algorithm was known for this class of formulae. We also present a deterministic algorithm that works in a blackbox fashion and runs in time $n^{k^{O(k)}+O(k \log n)}$ in general, and time $n^{k^{O(k^2)}+O(k\delta)}$ for depth δ. Finally, we extend our results and allow the inputs to be replaced with sparse polynomials. Our results encompass recent deterministic identity tests for sums of a constant number of read-once formulae, and for multilinear depth-four formulae.

3.26 Randomness Efficient Testing of Sparse Blackbox Identities of Unbounded Degree over the Reals

Markus Bläser (Universität des Saarlandes, DE)

Joint work of Bläser, Markus; Engels, Christian
Main reference Markus Blaeser and Christian Engels; "Randomness Efficient Testing of Sparse Black Box Identities of Unbounded Degree over the Reals," pp. 555–566, LIPIcs Vol. 9, STACS 2011.
URL http://dx.doi.org/10.4230/LIPIcs.STACS.2011.555

We construct a hitting set generator for sparse multivariate polynomials over the reals. The seed length of our generator is $O(\log^2(mn/\epsilon))$ where m is the number of monomials, n is number of variables, and $1-\epsilon$ is the hitting probability. The generator can be evaluated in time polynomial in $\log m$, n, and $\log 1/\epsilon$. This is the first hitting set generator whose seed length is independent of the degree of the polynomial.

The seed length of the best generator so far by Klivans and Spielman depends logarithmically on the degree.

From this, we get a randomized algorithm for testing sparse blackbox polynomial identities over the reals using $O(\log^2(mn/\epsilon))$ random bits with running time polynomial in $\log m$, n, and $\log \frac{1}{\epsilon}$.

We also design a deterministic test with running time $\tilde{O}(m^3n^3)$.

Here, the $\tilde{O}$-notation suppresses polylogarithmic factors.

The previously best determininstic test by Lipton and Vishnoi has a running time that depends polynomially on $\log \delta$, where δ is the degree of the blackbox polynomial.

3.27 Learning Read-Constant Polynomials of Constant Degree Modulo Composites

Kristoffer Arnsfelt Hansen (Aarhus University, DK)

Joint work of Chattopadhyay, Arkadev; Gavaldà, Ricard; Hansen, Kristoffer Arnsfelt; Thérien, Denis

Boolean functions that have constant degree polynomial representation over a fixed finite ring form a natural and strict subclass of the complexity class ACC^0. They are also precisely the functions computable efficiently by *programs* over fixed and finite nilpotent groups. This class is not known to be learnable in any reasonable learning model.

In this paper, we provide a deterministic polynomial time algorithm for learning Boolean functions represented by polynomials of constant degree over arbitrary finite rings from membership queries, with the additional constraint that each variable in the target polynomial appears in a constant number of monomials. Our algorithm extends to superconstant but low degree polynomials and still runs in quasipolynomial time.

3.28 Locality of AC^0-Computable Graph Queries

Nicole Schweikardt (Goethe-Universität Frankfurt am Main, DE)

Joint work of Anderson, Matthew; van Melkebeek, Dieter; Schweikardt, Nicole; Segoufin, Luc
Main reference Matthew Anderson, Dieter van Melkebeek, Nicole Schweikardt, Luc Segoufin, "Locality of queries definable in invariant first-order logic with arbitrary built-in predicates," pp. 368–379, LNCS Vol. 6756, ICALP 2011.
URL http://dx.doi.org/10.1007/978-3-642-22012-8_29

Our main theorem states that AC^0-computable graph queries are f-local, for every function f growing faster than a polylogarithmic function. Here, f-local means that for any graph of size n and any two tuples of nodes a and b whose $f(n)$-neighborhoods are isomorphic, b belongs to the query result if and only if a does.

Our proof makes use of the known tight lower bounds for *parity* on constant-depth circuits. The size $f(n)$ of the neighborhoods is optimal, since for every polylogarithmic function g we can find an AC^0-computable graph query that is not g-local.

3.29 A Survey of Exponential Algorithms for Constraint Satisfaction Problems

Robin Moser (ETH Zürich, CH)

In this talk, I surveyed the two most important currently competitive algorithms for k-SAT: Schöning's Algorithm and the PPSZ algorithm due to Paturi, Pudlák, Saks and Zane. There have been very recent news about both algorithms. In the case of Schöning, we were recently able with Scheder to provide a deterministic variant of essentially the same running time as the randomized one. In the case of PPSZ, a recent breakthrough result due to Timon Hertli shows that the success probability of the algorithm is no worse in the general case than in the case of a uniquely satisfying assignment, demonstrating that PPSZ is dramatically faster than was previously known.

3.30 A Full Derandomization of Schöning's k-SAT Algorithm

Robin Moser (ETH Zürich, CH)

Joint work of Moser, Robin; Scheder, Dominik
Main reference Robin A. Moser, Dominik Scheder; "A Full Derandomization of Schoening's k-SAT Algorithm," submitted, available at arXiv:1008.4067
URL http://arxiv.org/abs/1008.4067

Schöning [1] presents a simple randomized algorithm for k-SAT with running time $a_k^n \cdot \mathrm{poly}(n)$ for $a_k = 2(k-1)/k$. We give a deterministic version of this algorithm running in time $a_k^{n+o(n)}$.

References

1 U. Schöning. A probabilistic algorithm for k-SAT and constraint satisfaction problems. In *FOCS '99: Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, page 410, Washington, DC, USA, 1999. IEEE CS.

3.31 Algorithmic Meta Theorems Inside Logspace and Their Applications

Michael Elberfeld (Universität Lübeck, DE)

Joint work of Jakoby, Andreas; Stockhusen, Christoph; Tantau, Till
Main reference M. Elberfeld, A. Jakoby, and T. Tantau; "Logspace versions of the Theorems of Bodlaender and Courcelle," Proceedings of FOCS 2010, pp. 143–152, 2010.
URL http://eccc.hpi-web.de/report/2010/062/

Bodlaender's Theorem states that for every k there is a linear-time algorithm that decides whether an input graph has tree width k and, if so, computes a width-k tree decomposition. Courcelle's Theorem sets up on Bodlaender's Theorem and states that for every monadic second-order (MSO) formula ϕ and for every k there is a linear-time algorithm that decides whether a given logical structure $\mathcal{A}$ of tree width at most k satisfies ϕ.

Recently we proved that the Theorems of Bodlaender and Courcelle also hold when "linear time" is replaced by "logarithmic space", yielding an algorithmic meta theorem that can be used to show deterministic logarithmic-space solvability for MSO-definable problems on logical structures of bounded tree width.

In the talk I outline the above results, and discussed current work on (1) their applications and (2) refined algorithmic meta theorems for circuit classes inside L.

3.32 Triangular Rank and Sampling With Multiplicative Errors

Ilan Newman (Haifa University, IL)

Joint work of Newman, Ilan; Rabinovich, Yuri

It is well known that if a family F of subsets of a universe U has VC-dimension d, then for every probability distribution D over U there exists a sample $S \subset U$ of size $d * poly(1/\epsilon)$

which faithfully represents P on F up to an ϵ additive error. Namely, for which $|(S \cap T)/|S| - Pr_D(T)| \le \epsilon$ for every set T in F.

We ask whether there exists a property of F analogous to VC-dimension which would ensure the existence of a small sample faithfully representing D on F up to $(1+\epsilon)$ *multiplicative* error. The answer turns out to be positive, and the key parameter is the triangular rank of F. In particular, we show that if the triangular rank of F is t, then there exists a sample of size $min\{r \log |F| * poly(1/\epsilon), r^2 \log r * poly(1/\epsilon)\}$ which faithfully represents D on F up to $(1+\epsilon)$ multiplicative factor.

One application of this is to show some upper bounds on the dimension reduction possibility for ℓ_1 metrics.

3.33 Tight Lower Bounds for Greedy Routing in Uniform Small World Rings

Philipp Woelfel (University of Calgary, CA)

Joint work of Dietzfelbinger, Martin; Woelfel, Philipp
Main reference Martin Dietzfelbinger and Philipp Woelfel; "Tight Lower Bounds for Greedy Routing in Uniform Small World Rings," Proc. of 41st STOC, pp. 591–600, 2009.
URL http://dx.doi.org/10.1145/1536414.1536494

Consider the following game played on a board with $2n+1$ spaces labeled $-n, ..., -1, 0, 1, ..., n$ (from left to right). Fix some probability distribution μ over $\{1, ..., n\}^k$, where $k = O(logn)$. In the beginning, a token is placed on the board at a position chosen uniformly at random. Then, in each step the token is moved closer to its target, space 0, according to the following random experiment: Suppose after the i-th step the token is located at position x. In the $(i+1)$-th step, k distances $d_1, ..., d_k$ are chosen at random according to the probability distribution μ. Then, the distance d in $\{0, d_1, ..., d_k\}$ is applied that moves the token closest to its target. I.e., the token is moved position $x - d$, where $d \in \{0, d_1, ..., d_k\}$ minimizes $|x - d|$.

We are interested in the number of steps, T, it takes to move the token from its random start position to position 0, for the best probability distribution μ. The problem has been first studied in the context of Small World Graph routing [1, 2] and has applications in P2P Networks [3, 4] and Black Box Optimization [5].

Probability distributions μ for which $E[T] = O((logn)^2/k)$ are well known. In this talk I sketch a proof that this upper bound is asymptotically tight. Previous lower bounds were only known for restricted versions of the problem.

References

1 J.M. Kleinberg. Navigation in a small world. Nature, p. 845, 2000.
2 J.M. Kleinberg. The small-world phenomenon: An algorithmic perspective. 32nd STOC, 2000.
3 P. Ganesan and G. S. Manku. Optimal routing in chord. 15th SODA, 2004.
4 I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup protocol for internet applications. IEEE/ACM Transactions on Networking, 11:17–32, 2003.
5 M. Dietzfelbinger, J. E. Rowe, I. Wegener, and P. Woelfel. Tight bounds for blind search on the integers. 25th STACS, 2008.

3.34 On the Correlation of Parity with Small Depth Circuits

Johan Håstad (KTH Stockholm, SE)

We prove that the correlation of the depth d circuit of size S with parity is bounded by $2^{-\Omega(n/(\log S)^{d-1})}$. The result follows from applying the classical switching lemma together with some deterministic post-processing.

Participants

- Eric Allender
Rutgers Univ. – Piscataway, US
- Matthew Anderson
University of Wisconsin – Madison, US
- Paul Beame
University of Washington – Seattle, US
- Eli Ben-Sasson
Technion – Haifa, IL
- Markus Bläser
Universität des Saarlandes, DE
- Andrej Bogdanov
Chinese Univ. of Hong Kong, HK
- Beate Bollig
Technische Universität Dortmund, DE
- Holger Dell
HU Berlin, DE
- Martin Dietzfelbinger
TU Ilmenau, DE
- Andrew Drucker
MIT – Cambridge, US
- Michael Elberfeld
Universität Lübeck, DE
- Eldar Fischer
Technion – Haifa, IL
- Lance Fortnow
Northwestern University – Evanston, US
- Anna Gál
University of Texas – Austin, US
- Martin Grohe
HU Berlin, DE
- Kristoffer Arnsfelt Hansen
Aarhus University, DK
- Prahladh Harsha
TIFR Mumbai, IN
- Johan Hastad
KTH Stockholm, SE
- Stasys Jukna
Univ. Frankfurt am Main, DE
- Valentine Kabanets
Simon Fraser University – Burnaby, CA
- Michal Koucký
Acad. of Sciences – Prague, CZ
- Matthias Krause
Universität Mannheim, DE
- Troy Lee
National University of Singapore, SG
- Xin Li
University of Texas - Austin, US
- Meena Mahajan
The Institute of Mathematical Sciences – Chennai, IN
- Pierre McKenzie
Université de Montréal, CA
- Or Meir
Weizmann Inst. – Rehovot, IL
- Peter Bro Miltersen
Aarhus University, DK
- Robin Moser
ETH Zürich, CH
- Ilan Newman
University of Haifa, IL
- Jakob Nordström
KTH Stockholm, SE
- Pavel Pudlák
Acad. of Sciences – Prague, CZ
- Oded Regev
Tel Aviv University, IL
- Rüdiger Reischuk
Universität Lübeck, DE
- Rahul Santhanam
University of Edinburgh, GB
- Georg Schnitger
Univ. Frankfurt am Main, DE
- Uwe Schöning
Universität Ulm, DE
- Nicole Schweikardt
Univ. Frankfurt am Main, DE
- Srikanth Srinivasan
Institute for Advanced Study – Princeton, US
- Amnon Ta-Shma
Tel Aviv University, IL
- Till Tantau
Universität Lübeck, DE
- Dieter van Melkebeek
University of Wisconsin – Madison, US
- Emanuele Viola
Northeastern Univ. – Boston, US
- Heribert Vollmer
Leibniz Univ. Hannover, DE
- Fabian Wagner
Universität Ulm, DE
- Thomas W. Watson
University of California – Berkeley, US
- Philipp Wölfel
University of Calgary, CA

Report from Dagstuhl Seminar 11131

Exploration and Curiosity in Robot Learning and Inference

Edited by
Jeremy L. Wyatt[1], Peter Dayan[2], Ales Leonardis[3], and Jan Peters[4]

1 **University of Birmingham, GB,** j.l.wyatt@cs.bham.ac.uk
2 **University College London, GB,** dayan@gatsby.ucl.ac.uk
3 **University of Ljubljana, SI,** alesl@fri.uni-lj.si
4 **MPI für biologische Kybernetik – Tübingen, DE,** jan.peters@tuebingen.mpg.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11131 "Exploration and Curiosity in Robot Learning and Inference". This seminar was concerned with answering the question: *how should a robot choose its actions and experiences so as to maximise the effectiveness of its learning?*. The seminar brought together workers from three fields: machine learning, robotics and computational neuroscience. The seminar gave an overview of active research, and identified open research problems. In particular the seminar identified the difficulties in moving from theoretically well grounded notions of curiosity to practical robot implementations.

Seminar 27. March – 1. April, 2011 – www.dagstuhl.de/11131
1998 ACM Subject Classification I.2.9 Robotics
Keywords and phrases Artificial Intelligence, Robotics, Learning, Exploration, Curiosity
Digital Object Identifier 10.4230/DagRep.1.3.67

1 Executive Summary

Jeremy L. Wyatt
Peter Dayan
Ales Leonardis
Jan Peters

Aims and background to the seminar

This seminar was concerned with answering the question: *how should a robot choose its actions and experiences so as to maximise the effectiveness of its learning?*.

This seminar was predicated on the assumption that to make significant progress in autonomous robotics, systems level theories of how robots should operate will be required. In recent years methods from machine learning have been employed with great success in robotics, in fields as diverse as visual processing, map building, motor control and manipulation. The machine learning algorithms applied to these problems have included statistical machine learning approaches, such as EM algorithms and density estimation, as well as dimensionality reduction, reinforcement learning, inductive logic programming, and other supervised learning approaches such as locally weighted regression. Most of these robot learning solutions currently require a good deal of supervised learning, or structuring of the learning data for a specific task. As robots become more autonomous these learning algorithms will have to be embedded in algorithms which choose the robot's next learning experience. The problems become particularly challenging in the context of robotics, and even worse for a robot that is faced with many learning opportunities. A robot that can perform both manipulation, language use, visual learning, and mapping may have several quite different learning opportunities facing it at any one time. How should a robot control its curiosity in a principled way in the face of such a variety of choices? How should it choose data on the basis of how much it knows, or on how surprising it finds certain observations? What rational basis should robot designers choose to guide the robot's choice of experiences to learn from? There has been initial progress in several fields, including *machine learning*, *robotics* and also in *computational neuroscience*. In this seminar we brought together these three communities to shed light on the problem of how a robot should select data to learn from, how it should explore its environment, and how it should control curiosity when faced with many learning opportunities.

Machine Learning

This problem of how to explore is one that has been studied both in the context of reinforcement learning (exploration vs. exploitation) and supervised learning (active learning). Within the former the language of the sequential decision making community is that of MDPs and POMDPs. These are related to Bayesian perspectives on supervised Machine Learning in that we can think of a posterior over hypotheses that results from data that may be seen. In statistics this is related to the idea of conducting pre-posterior analysis. However, the theories from sequential decision making and active learning are currently unintegrated and partial, and it is not clear how they should apply in robotics. The current machine learning preference for Bayesian methods suggests that the ways that model uncertainty can be captured and exploited will be critical. During the seminar we looked for suggestions from this community as to how problems of exploration and curiosity in robotics can be formalised, especially at a systems level.

Robotics

Within areas like SLAM (Simultaneous Localisation and Mapping) the problem of how to select data has been addressed, but heuristic measures of exploratory worth are typically employed. Again, the principal formalism is that of Bayesian filtering, within which a POMDP is posed, but typically only the belief filter part is used. Rather than look ahead at all possibilities, heuristics such as information gain are used. There are also other approaches necessary in areas where for example robot learners are learning associations between data from multiple modalities, from time series, and where there maybe limited intervention from humans. Again, learning approaches in contemporary robotics are typically statistical, but there are other approaches. Techniques are also adapted to the domain, such as in the community working on the use of robotics in scientific discovery in the laboratory, where the robot has the ability to determine which experiments to perform, and the methods used employ a great deal of structure and prior knowledge about the domain. There are also challenges for exploration and curiosity from the use of robots for scientific exploration, such as in planetary missions, and in subsea exploration. We looked at how these problems are currently posed, and the challenges they pose for machine learning approaches to data selection.

Animal Cognition and Computational Neuroscience

The field of computational neuroscience has many insights to offer roboticists. Animals are forced to consider at each moment how they select data. There are examples of this in studies of motor learning, animal foraging, studies of neurotransmitters, and particular learning circuits, as well as in the study of areas of the brain concerned with action selection. Computational neuroscience has also been strongly influenced by statistical, particularly Bayesian approaches to inference and learning. For example much recent work strongly suggests a Bayesian underpinning to learning in the motor system, and other work has investigated possible neural bases for learning in the face of suprise and uncertainty. Work on reinforcement learning has linked with studies of brain areas such as the Basal ganglia, and there is debate as to whether or not the purpose of certain neurons is to provide a cue for learning in the face of novelty. This is related to the idea of infotaxis as a general mechanism for exploration control in some animal. The connection to the statistical theories from machine learning and optimal control are intriguing. This gives us a strong basis for the hope that a common framework for exploration and curiosity might emerge as a consequence of this seminar.

Summary of objectives

In summary the objectives of this seminar were to:
- Identify the different formulations of exploration and curiosity control, and to categorise robot problems into appropriate classes.
- Share statistical and non-statistical representations suitable for control of curiosity and exploration across communities.
- Identify the links between studies of learning control and motivation in computational neuroscience and formalisations from robotics and machine learning.
- Discuss possible formalisations of the problem of learning one of many possible tasks.
- Identify whether solution classes are heuristic or optimal.

Summary of the seminar program

The seminar was grouped into three themes, roughly according to Marr's levels of description: computational, algorithmic and implementational. Many talks crossed more than one level, but within these themes we were able to organize talks around more specific research areas. These areas were:

1. Ideas from neuroscience about the implementation of exploration and action in the brain.
2. Evidence from the ethology and psychology about the requirements for exploration, and algorithmic frameworks that fit the data on human behaviour.
3. Computational frameworks for intrinsic motivation and the evolution of extrinsic reward functions.
4. Algorithms and properties for specific sub-problems within curiosity and exploration: such as visual object search or the behaviour of greedy algorithms for solving sub-modular problems.
5. Robot implementations of algorithms for control of exploration and curiosity in real tasks.

Summary of the fundamental results

The main findings presented can be grouped into four parts. It is worth stating from the outset that a very large number of the talks, though by no means all, employed a reward based framework. It is not possible in this summary to mention all of the thirty talks given, instead we mention talks that illustrate the common themes of the seminar.

First a tutorial on the Basal Ganglia and it's role in action selection, including for exploration was given by Humphries. In this field there are now a range of computational models that simulate some of the internal workings of the Basal Ganglia. It was clear, however, that there are numerous structures about which little is known, and that many details of the models remain to filled in, or to be tested. Dayan provided evidence that was broadly negative with respect to a Bayesian view of exploration in humans. Dayan showed that human behaviour in a non-stationary bandit task is not better explained by a Bayesian view than by a simple soft-max reinforcement learning model. Sloman argued that the requirements to support exploration include the need to decompose domains into reuseable patterns. Contrary to a large number of the speakers Sloman argued against a statistical approach to exploration control. In the workshop as a whole statistical methods, with rewards, and often Bayesian inference were dominant, but these talks from biology present evidence that was not always supportive of this dominant approach.

Several different frameworks for intrinsic motivation were given. In several of these (Schmidhuber, Polani, Auer) the idea that exploration is driven by curiosity to enable greater understanding and ability to exploit the environment was central. These approaches can be contrasted with those that are ultimately driven by the need to maximise extrinsic rewards (Tishby, Starzyk). There seemed little question that all of these frameworks are quite general, but no clear unifying account is available. Others (Uchibe, Barto, Elfwing) showed ways to evolve extrinsic reward functions that ultimately contribute to overall agent fitness. Overall the division seems to be between information seeking and value seeking frameworks for self-motivation.

In algorithms the important findings concerned cases where problems that in the general case are intractable can be tackled much more effectively in special cases. Tsotsos showed as part of his talk that some visual search problems are tractable even though in the general case they are not. Krause showed that where problems have a sub-modular property that greedy algorithms can be close to optimal. Dearden showed how for a particular search task that entropic heuristics perform close to the level of more computational expensive information lookahead methods. While the most general algorithmic frameworks to exploration are based on the solution of POMDPs, each of these talks showed that a solution to a simpler problem can often provide very good performance.

In robotic tasks some approaches were necessarily more pragmatic, and this meant that many moved away from a purely reward based framework. Several showed ways of approximating solutions to POMDPs in real robot systems. These included using hierarchical approaches, sampling methods, limited horizon lookahead, or methods that split the problem into parts with, and without state uncertainty (Wyatt, Peters, Martinez-Cantin). While some advocated implementation of the principled frameworks for intrinsic motivation (Pape), problems were often moved away from the common statistical, reward based framework to enable solutions. The benefit of heuristic goal selection methods on top of precise planning approaches to achieving selected goals was demonstrated (Hanheide, Skocaj). A variety of robotic tasks were shown to be tackled with active learning, including motor control and social learning (Peters, Lopes).

The main themes that emerged were that while the dominant paradigm was one that was statistical and reward based, there were alternatives. While there were theoretically rigorous frameworks based on rewards, these were actually not much used by roboticists, who preferred pragmatic approaches. In the middle sit those exploring algorithms that while still approximate, offer some performance bounds relative to that which is optimal, howsoever defined.

2 Table of Contents

Executive Summary
Jeremy L. Wyatt, Peter Dayan, Ales Leonardis, Jan Peters 67

Overview of Talks

A Model for Evaluating Autonomous Exploration
Peter Auer 74

Learning Qualitative Spatio-temporal models of Activities and Objects from Video
Anthony G. Cohn 74

Polar Exploration
Peter Dayan 74

Robot Inference in Ocean Exploration
Richard W. Dearden 75

Supporting exploration in children with disabilities through lifelong robotic assistants
Yiannis Demiris 75

Embodied Evolution of Learning Ability and the Emergence of Different Mating Strategies in a Small Robot Colony
Stefan Elfwing 75

Hierarchical object inference and exploration in mobile whiskered robots
Charles Fox 76

Learning-based modeling and stability design of interactive human locomotion patterns
Martin A. Giese 76

Dora: explore and exploit probabilistic knowledge for object search
Marc Hanheide 77

Neural substrates for action selection: the basal ganglia
Mark D. Humphries 77

Exploration in Relational Worlds
Kristian Kersting 78

Adaptive Submodularity: A New Approach towards Active Learning and Stochastic Optimization
Andreas Krause 78

Active Exploration in Social Learning
Manuel Lopes 79

Non-convex optimization for active robot learning
Ruben Martinez-Cantin 79

Towards practical implementations of curious robots
Leo Pape 79

Exploration in Learning of Motor Skills for Robotics
Jan Peters 80

Potential Information Flows as Driving Principle for Agents
Daniel Polani 80

Formal Theory of Fun and Creativity for Curious Agents
Juergen Schmidhuber 81

Exploration in parameter space - a practical evaluation
Frank Sehnke 81

Interactive learning in dialogue with a tutor
Danijel Skocaj 82

How to combine AI and Piaget's genetic epistemology (understanding possibility and necessity in multiple exploration domains)
Aaron Sloman 82

A Mechanism for Learning, Attention Switching, and Cognition
Janusz Starzyk 84

Information flows and discounting in planning and learning
Naftali Tishby 85

How should Robbie search for my lost cup?
John K. Tsotsos 85

Finding Intrinsic Rewards by Embodied Evolution and Constrained Reinforcement Learning
Eiji Uchibe 86

Goal-Directed Action Generalization
Ales Ude 86

Robot Visual Learning
Markus Vincze 86

Active information gathering in vision and manipulation
Jeremy L. Wyatt 87

Embodied Learning of Qualitative Models
Jure Zabkar 87

Working Groups and Open Problems

Group 1: Self-motivation 88

Group 2: Scaling methods for exploration and curiosity 89

Group 3: Optimality and sub-optimality in exploration 91

Group 4: Unifying themes 93

Participants 95

3 Overview of Talks

3.1 A Model for Evaluating Autonomous Exploration

Peter Auer (Montan-Universität Leoben, AT)

I consider the problem of evaluating exploration strategies without predefined goals or rewards. Instead of evaluating only the behavioural complexity of the system after exploration, I propose to measure what the system can accomplish.

In a rather simplistic scenario this could be measured by the number of states the system can get into efficiently.

3.2 Learning Qualitative Spatio-temporal models of Activities and Objects from Video

Anthony G. Cohn (University of Leeds, GB)

In this talk I will present ongoing work at Leeds on building models of the world from observation, concentrating on unsupervised learning. The representations exploit qualitative spatio-temporal relations. A novel method for robustly transforming video data to qualitative relations will be presented.

I will present results from several domains including a kitchen scenario and an aircraft apron.

3.3 Polar Exploration

Peter Dayan (University College London, GB)

I will discuss the exploratory and exploitative behavior of human subjects in a four-armed restless bandit task. Despite our best analytical efforts, we could find no evidence that subjects awarded exploration bonuses to options they hadn't tried for a while. Instead, their exploratory behaviour was well-captured by a form of softmax choice in a conventional reinforcement learning model.

Fronto-polar cortex, a large and poorly understood area of the human brain, was specifically activated on trials that this model classified as exploratory, to a degree that depended on the requirement for cognitive control associated with those trials.

3.4 Robot Inference in Ocean Exploration

Richard W. Dearden (University of Birmingham, GB)

We look at the problem of finding hydrothermal vents using an autonomous underwater vehicle. This is an exploration problem where reward is only associated with finding a few specific states, rather than the quality of the map discovered. The problem can be formulated as a partially observable Markov decision process, but is far too large to be solved exactly. We examine two approaches, one based on approximating the solution to the POMDP using forward search in belief space to take account the information gained through each observation made, and the other based on treating the problem as a mapping problem and using entropy reduction. We show that both approaches perform much better than the state of the art, and that statistically the two approaches perform very similarly, suggesting that entropy reduction is a useful heuristic even for problems like this where it is clearly optimising the wrong criterion.

3.5 Supporting exploration in children with disabilities through lifelong robotic assistants

Yiannis Demiris (Imperial College London, GB)

Children and adults with sensorimotor disabilities can significantly increase their autonomy through the use of assistive robots. As the field progresses from short-term, task- specific solutions to long-term, adaptive ones, new challenges are emerging. In this talk a lifelong methodological approach is presented, that attempts to balance the immediate context-specific needs of the user, with the long-term effects that the robot's assistance can potentially have on the user's developmental trajectory. I will use examples from adaptive robotic wheelchairs assisting young children and adults to illustrate the methodology, and I will discuss the underlying computational learning mechanisms and robotic infrastructure.

3.6 Embodied Evolution of Learning Ability and the Emergence of Different Mating Strategies in a Small Robot Colony

Stefan Elfwing (Okinawa Institute of Science and Technology, JP)

In this study, we use a framework for performing embodied evolution with a limited number of robots, by utilizing time-sharing in subpopulations of virtual agents hosted in each robot. Within this framework, we explore the combination of within-generation learning of basic survival behaviors by reinforcement learning, and evolutionary adaptations over the generations of the basic behavior selection policy, the reward functions, and meta-parameters for reinforcement learning. We apply a biologically inspired selection scheme, in which there is no explicit communication of the individuals' fitness information. The individuals can only reproduce offspring by mating-a pair-wise exchange of genotypes-and the probability that an individual reproduces offspring in its own subpopulation is dependent on the individual's "health," i.e., energy level, at the mating occasion. In addition, we investigate the emergence of different mating strategies, i.e., different basic behavior selection policies.

In the experiments, we observed two individual mating strategies: 1) Roamer strategy, where an agent never waits for potential mating partners and 2) Stayer strategy, where an agent waits for potential mating partners depending on the agent's current state. The most interesting finding was that in some simulations the evolution produced a mixture of mating strategies within a population, typically with a narrow roamer subpopulation and a broader stayer subpopulation with distinct differences in genotype, phenotype, behavior, and performance between the subpopulations.

3.7 Hierarchical object inference and exploration in mobile whiskered robots

Charles Fox (University of Sheffield, GB)

I will describe work on object recognition with autonomous mobile whiskered robots. Whiskers are highly localised sensors and the question of "where to look next" to gather useful information is important. I will describe a Bayesian Blackboard approach to recognising hierarchical objects, using Gibbs sampling and annealing together with blackboard system heuristics. Entropy in the Gibbs sampler provides a natural measure of saliency which may be used as a basis to guide exploration.

3.8 Learning-based modeling and stability design of interactive human locomotion patterns

Martin A. Giese (Universitätsklinikum Tübingen, DE)

Joint work of M.A. Giese, A. Mukovskiy, J.J. Slotine L. Omlor

The online synthesis of interactive human body movements is a key problem for different technical applications, such as robotics and computer graphics. We present a biologically inspired algorithm for the modeling of complex body movements based on dynamic primitives by a combination of unsupervised learning and methods from nonlinear dynamics. We illustrate how this method can be applied for the synthesis of collective behaviors of groups of locomoting agents, taking the full nonlinearity of the kinematics into account. For selected cases we show how Contraction Theory can be applied to analyze and design the stability properties of the emerging collective behavioral patterns.

Supported by the DFG, EU FP6 Project COBOL, and FP7 Projects TANGO and AMARSI.

3.9 Dora: explore and exploit probabilistic knowledge for object search

Marc Hanheide (University of Birmingham, GB)

I am presenting Dora, our robot systems that explores its environment in a task-driven way to find an object. Besides taking topological and spatial knowledge pre-acquired as part of a mapping process into account it is also equipped with a probabilistic ontology that enables it to exploit common-sense knowledge. This probabilistic model comprises information about the likelihood of finding certain objects in particular categories of rooms (e.g. that cornflakes are usually found in kitchens) and observation models of successfully detecting an object if it is present. By reasoning within this probabilistic spatial representation that is incorporates perceived instance knowledge a probabilistic belief state of the world is maintained. Planning within this probabilistic representation using a switching planner, switching between sequential and contingent planning sessions, we can exploit the common-sense knowledge to find objects more efficiently in many cases. However, the system can still cope with violated assumptions, i.e. when common-sense knowledge is not applicable, and sensing errors. I discuss the limitations and caveats of our approach so far and briefly talk about the challenge to combine task-driven with curiosity driven exploration in our world. The accompanying video can be viewed at
http://www.youtube.com/watch?v=0QcmSDZR-c4.

3.10 Neural substrates for action selection: the basal ganglia

Mark D. Humphries (ENS – Paris, FR)

All animals must continuously sequence and co-ordinate behaviors appropriate to both their context and internal milieu if they are to survive. It is natural to wonder what parts of the nervous system - the neural substrate - evolved to carry out this action selection process. I will discuss the proposal that the vertebrate brain has co-evolved (or co-opted) specialised and centralised neural systems for action selection, handling both the competition between systems accessing the final motor pathway and the open-ended nature of a flexible, extensible behavioural repertoire. In particular, I will focus on the set of fore- and mid-brain nuclei called the basal ganglia that seem to implement a repeated circuit design ideal for computing input selection. Of particular note for understanding spatial exploration are the circuits implemented by the ventral basal ganglia, which integrate information on current position, overall strategy, and reward. The talk will first lay out this circuit, sketching the unique features of both its components and overall architecture that have given clues to its function. I will then address the breadth of computational modelling of the generic basal ganglia circuit, considering both mechanism-mining (build biologically accurate model, search for insights) and mechanism-mapping (take algorithm, fit to extant neural circuit) approaches to understanding neural computation.

3.11 Exploration in Relational Worlds

Kristian Kersting (Fraunhofer IAIS – St. Augustin, DE)

One of the key problems in model-based reinforcement learning is balancing exploration and exploitation. Another is learning and acting in large relational domains, in which there is a varying number of objects and relations between them. We provide a solution to exploring large relational Markov decision processes by developing relational extensions of the concepts of the Explicit Explore or Exploit (E3) algorithm. A key insight is that the inherent generalization of learnt knowledge in the relational representation has profound implications also on the exploration strategy: what in a propositional setting would be considered a novel situation and worth exploration may in the relational setting be an instance of a well-known context in which exploitation is promising. Our experimental evaluation shows the effectiveness and benefit of relational exploration over several propositional benchmark approaches on noisy 3D simulated robot manipulation problems.

3.12 Adaptive Submodularity: A New Approach towards Active Learning and Stochastic Optimization

Andreas Krause (ETH Zürich, CH)

Joint work of Krause, Andreas; Golovin, Daniel; Ray, Debaivoti

Many information gathering problems require us to adaptively select observations to obtain the most useful information. These problems involve sequential stochastic optimization under partial observability – a fundamental but notoriously difficult challenge. Fortunately, often such problems have a structural property that makes them easier than general sequential stochastic optimization. In this talk, I will introduce this structural property – a new concept that we call adaptive submodularity – which generalizes submodular set functions to adaptive policies.

In many respects adaptive submodularity plays the same role for adaptive problems such as sequential experimental design as submodularity plays for nonadaptive problems (such as placing a fixed set of sensors). Specifically, just as many nonadaptive problems with submodular objectives have efficient algorithms with good approximation guarantees, so too do adaptive problems with adaptive submodular objectives. I will illustrate the usefulness of the concept by giving several examples of adaptive submodular objectives arising in diverse applications including sensor selection, viral marketing and active learning.

Proving adaptive submodularity for these problems allows us to recover existing results in these applications as special cases and handle natural generalizations. In an application to Bayesian experimental design, we show how greedy optimization of a novel adaptive submodular criterion outperforms standard myopic heuristics such as information gain and value of information.

3.13 Active Exploration in Social Learning

Manuel Lopes (INRIA – Bordeaux, FR)

In the work we present a system to learn task representations from ambiguous feedback. We consider an inverse reinforcement learner that receives feedback from a user with an unknown and noisy protocol. The system needs to estimate simultaneously what the task is, and how the user is providing the feedback. We further explore the problem of ambiguous protocols by considering that the words used by the teacher have an unknown relation with the action and meaning expected by the robot. We present computational results in a large scale problem and on a simplified robotic one. We show that it is possible to learn the task under a noisy and ambiguous feedback. Using an active learning approach, the system is able to reduce the length of the training period.

3.14 Non-convex optimization for active robot learning

Ruben Martinez-Cantin (CUD – Zaragoza, ES)

Active learning provides the optimal trade-off between statistical learning and decision making. Traditionally, both fields have been characterized for relying in convex functions to solve high dimensional problems. However, many robotics problems can be formulated as a low dimensional problem in continuos state and action spaces. Thus, non-convex optimization methods such as EI can be applied efficiently.

3.15 Towards practical implementations of curious robots

Leo Pape (IDSIA – Lugano, CH)

Schmidhuber's theory of compression-driven progress considers limited computational agents that try to represent their observations in an efficient manner. Finding efficient representations entails identifying regularities that allow the observer to compress the original observations and predict future observations. Compression progress is achieved when the agent discovers previously unknown regularities that allow for increased compression of observations. The compression progress of the compression mechanism is monitored by an action generation method that generates actions for which it expects further improvement in the compressor. This allows the agent to focus on collecting interesting observations, that is, observations that are novel, yet learnable by the compressor. In my talk I will discuss how this general principle of compression progress is used for the implementation of curious robots.

3.16 Exploration in Learning of Motor Skills for Robotics

Jan Peters (MPI für biologische Kybernetik – Tübingen, DE)

Intelligent autonomous robots that can assist humans in situations of daily life have been a long standing vision of robotics, artificial intelligence, and cognitive sciences. A elementary step towards this goal is to create robots that can learn tasks triggered by environmental context or higher level instruction. However, learning techniques have yet to live up to this promise as only few methods manage to scale to high-dimensional manipulator or humanoid robots. In this talk, we investigate a general framework suitable for learning motor skills in robotics which is based on the principles behind many analytical robotics approaches. It involves generating a representation of motor skills by parameterized motor primitive policies acting as building blocks of movement generation, and a learned task execution module that transforms these movements into motor commands. We discuss learning on three different levels of abstraction, i.e., learning for accurate control is needed to execute, learning of motor primitives is needed to acquire simple movements, and learning of the task-dependent "hyperparameters" of these motor primitives allows learning complex tasks. We discuss task-appropriate learning approaches for imitation learning, model learning and reinforcement learning for robots with many degrees of freedom.

Empirical evaluations on a several robot systems illustrate the effectiveness and applicability to learning control on an anthropomorphic robot arm. A large number of real-robot examples will be demonstrated ranging from Learning of Ball-Paddling, Ball-In-A-Cup, Darts, Table Tennis to Grasping.

3.17 Potential Information Flows as Driving Principle for Agents

Daniel Polani (University of Hertfordshire, GB)

In recent years, models for self-motivation in robots and agents have attracted significant interest in the research community. One challenge is to formulate such intrinsic drives with as little "engineer's bias" as possible. Many important approaches have been suggested in the last decades, which one can roughly classify e as process-oriented and structure-oriented. We consider the structure-oriented quantity of "empowerment", which is the potential Shannon information flow in the external perception-action loop of agents. Embedded in the framework of information theory, it displays a number of properties that make it promising for the use as an intrinsic motivation drive in a structured environment. Empowerment probes salient aspects of the structure of an agent's environment, and displays intuitively plausible solutions. The talk will introduce the concept, demonstrate a number of examples and demonstrate several properties of empowerment which may make it relevant not only for constructing artificial drives but also for speculations on biological ones.

3.18 Formal Theory of Fun and Creativity for Curious Agents

Juergen Schmidhuber (IDSIA – Lugano, CH)

Main reference J. Schmidhuber, "Formal Theory of Creativity, Fun, and Intrinsic Motivation (1990-2010)." IEEE Transactions on Autonomous Mental Development, 2(3):230-247,2010.
URL http://www.idsia.ch/~juergen/ieeecreative.pdf

Our fast deep / recurrent neural nets recently achieved numerous 1st ranks in many pattern recognition competitions and benchmarks, without any unsupervised pre-training. The future, however, will belong to active systems learning to sequentially shift attention towards informative inputs, not only solving externally posed tasks, but also their own self-generated tasks designed to improve their understanding of the world according to our Formal Theory of Fun and Creativity, which requires two interacting modules: (1) an adaptive (possibly neural) predictor or compressor or model of the growing data history as the agent is interacting with its environment, and (2) a (possibly neural) reinforcement learner. The learning progress of (1) is the FUN or intrinsic reward of (2). That is, (2) is motivated to invent skills leading to interesting or surprising novel patterns that (1) does not yet know but can easily learn (until they become boring). We discuss how this principle explains science and art and music and humor.

Further details are available at

- Formal Theory of Fun & Creativity: www.idsia.ch/~juergen/creativity.html
- Artificial Curiosity: www.idsia.ch/~juergen/interest.html

3.19 Exploration in parameter space - a practical evaluation

Frank Sehnke (ZSW – Stuttgart, DE)

Parameter Exploring Policy Gradients like NES, SDE-PG, PGPE and NPGPE have drawn some attention lately. We will summarise why this is the case, and focus on the question why exploring in parameter space is so special. Furthermore, learning of real-world complex RL tasks with these methods has shown that subtle deviations from the mathematically correct implementation of the exploration part lead to significant improvements to both robustness of the algorithm and quality of the solution found. The changes introduced are strikingly similar to the underlying principles of exploration in evolutionary methods, and as a whole give rise to a provocative question: Are we optimising the wrong parameters for exploration in PG?

3.20 Interactive learning in dialogue with a tutor

Danijel Skocaj (University of Ljubljana, SI)

Joint work of Danijel Skocaj, Matej Kristan, Alen Vrecko, Marko Mahnic, Barry Ridge, Ales Leonardis, Pierre Lison, Miroslav Janicek, Geert-Jan M. Kruijff , Ivana Kruijff-Korbayova, Raveesh Meena, Michael Zillich, Kai Zhou, Moritz Göbelbecker, Thomas Keller, Marc Hanheide, Nick Hawes
URL http://cogx.eu/results/george/

In this talk we will present representations and mechanisms that facilitate continuous learning of visual concepts in dialogue with a tutor and show the implemented robot system. We will present how the beliefs about the world are created by processing visual and linguistic information and how they are used for planning the system behaviour aimin at satisfying its internal drive - to extend its knowledge. The system facilitates different kinds of learning initiated by the human tutor or by the system itself. We demonstrate these principles in the case of learning about object colours and basic shapes and present the experimental results that justify such mixed-initiative learning process.

3.21 How to combine AI and Piaget's genetic epistemology (understanding possibility and necessity in multiple exploration domains)

Aaron Sloman (University of Birmingham, GB)

It is not widely known that shortly before he died Jean Piaget and his collaborators produced a pair of books on Possibility and Necessity, exploring questions about how two linked sets of abilities develop:

(a) The ability to think about how things might be, or might have been, different from the way they are.
(b) The ability to notice limitations on possibilities, i.e. what is necessary or impossible.

I believe Piaget had deep insights into important problems for cognitive science that have largely gone unnoticed, and are also important for research on intelligent robotics, or more generally Artificial Intelligence (AI), as well as for studies of animal cognition and how various animal competences evolved and develop.

The topics are also relevant to understanding biological precursors to human mathematical competences and to resolving debates in philosophy of mathematics, e.g. between those who regard mathematical knowledge as purely analytic, or logical, and those who, like Immanuel Kant, regard it as being synthetic, i.e. saying something about reality, despite expressing necessary truths that cannot be established purely empirically, even though they may be initially discovered empirically (as happens in children).

It is not possible in one seminar to summarise either book, but I shall try to present an overview of some of the key themes and will discuss some of the experiments intended to probe concepts and competences relevant to understanding necessary connections.

In particular, I hope to explain: (a) The relevance of Piaget's work to the problems of designing intelligent machines that learn the things humans learn.

(Most researchers in both Developmental Psychology and AI/Robotics have failed to notice or have ignored most of the problems Piaget identified.) (b) How a deep understanding of AI, and especially the variety of problems and techniques involved in producing machines that can learn and think about the problems Piaget explored, could have helped Piaget describe and study those problems with more clarity and depth, especially regarding the forms of representation required, the ontologies required, the information processing mechanisms required and the information processing architectures that can combine those mechanisms in a working system – especially architectures that grow themselves.

That kind of computational or "design-based" understanding of the problems can lead to deeper clearer specifications of what it is that children are failing to grasp at various stages in the first decade of life, and what sorts of transitions can occur during the learning. I believe the problems, and the explanations, are far more complex than even Piaget thought. The potential connection between his work and AI was appreciated by Piaget himself only very shortly before he died.

One of the key ideas implicit in Piaget's work (and perhaps explicit in something I have not read) is that the learnable environment can be decomposed into explorable domains of competence that are first investigated by finding useful, reusable patterns, describing various fragments. Then eventually a large scale reorganisation is triggered (per domain) which turns the information about the domain into a more economical and more powerful generative system that subsumes most of the learnt patterns and, through use of compositional semantics in the internal representation, allows coping with much novelty – going far beyond what was learnt.

(I think this is the original source of human mathematical competences.)

Language learning seems to use a modified, specialised, version of this more general (but not totally general) mechanism, but the linguistic mechanisms were both a later product of evolution and also get turned on later in young humans than the more general domain learning mechanisms. The linguistic mechanisms also require (at a later stage) specialised mechanisms for learning, storing and using lots of exceptions to the general rules induced (the syntactic and semantic rules).

The language learning builds on prior learning of a variety of explorable domains, providing semantic content to be expressed in language. Without that prior development, language learning must be very shallow and fragmentary – almost useless.

When two or more domains of exploration have been learnt they may be combinable, if their contents both refer to things and processes in space time. E.g. a domain of actions on sand and a domain of actions on water could be combined, producing things like sandcastles with moats, mud, etc.

I think Piaget was trying to say something like this but did not have the right concepts, though his experiments remain instructive.

Space-time is the the great bed in which many things can lie together and produce novelty. One of the features is a kind of continuity of interaction of structures (e.g. levers, gear wheels, pulleys, string, etc.) that I don't think current forms of representation in AI are well suited to.

Moreover, I have the impression that attempts to deal with uncertainty by using probabilities are completely mistaken, and that biological evolution produced something far more powerful, mainly concerned with use of non-metrical or semi-metrical (partially ordered) forms of representation.

Producing working demonstrations of these ideas in a functional robot able to manipulate things as a child does will require major advances in AI, though there may already be more work of this type than I am aware of.

A version of this talk presented at Birmingham on 21st Feb 2011 is available here (PDF): http://www.cs.bham.ac.uk/research/projects/cogaff/misc/talks/#talk90

(Comments and criticisms welcome, before, during or after Dagstuhl!)

Since I originally gave the talk I have discovered that some of the ideas are also in Annette Karmiloff-Smith's 1992 book, Beyond Modularity, though she does not stress, as Piaget does, the importance of being able to think about possibility and necessity.

3.22 A Mechanism for Learning, Attention Switching, and Cognition

Janusz Starzyk (Ohio University, US)

This talk would directly relate to questions of how should a robot control its curiosity when faced with many learning opportunities and a variety of choices, how should it choose data on the basis of how much it knows and what are his objectives, or on how surprising it finds certain observations? In this talk a new machine learning method called motivated learning (ML) will be presented.

ML applies to autonomous embodied intelligence agents who build perceptions and learn their actions so as to maximize the effectiveness of their mental development through learning. Motivated learning drives a machine to develop abstract motivations and choose its own goals. ML also provides a self-organizing system that controls a machine's behavior based on competition between dynamically-changing motivations, perceptions and internal thoughts.

This provides interplay of externally driven and internally generated signals that control a machine's behavior.

ML method can be combined with artificial curiosity and reinforcement learning.

It enhances their versatility and learning efficiency, particularly in changing environments with complex dependencies between environment parameters. It has been demonstrated that ML not only yields a more sophisticated learning mechanism and system of values than reinforcement learning (RL), but is also more efficient in learning complex relations and delivers better performance than RL in dynamically changing environments.

In addition, ML provides a much needed mechanism for switching a machine's attention to new motivations and implementation of internal goals. A motivated learning machine develops and manages its own motivations and selects goals using continuous competition between various levels of abstract pain signals (including curiosity and possible attention switching signals). This form of distributed goal management and competing motivations is a core element of central executive control that may govern the cognitive operation of intelligent machines.

This talk will present basic properties and underlying philosophy of ML. In addition, it will present the basic neural network structures used to create abstract motivations, higher level goals, and subgoals. It will show some simulation results to compare ML and RL in environments of gradually increasing sophistication and levels of difficulty.

Subsequently, an organization of central executive unit for a cognitive system will be described. This unit uses distributed and competing signals that represent goals, motivations, emotions, and attention. Only after the winner of such competition is established, it drives a focus point of a conscious experience. This computational model of consciousness mimics the biological systems functionally and retains a well-defined architecture necessary for implementing consciousness in machines. In addition a new concept of mental saccades will be presented to explain the attention switching and focusing in this computational model of consciousness.

The model uses competition among three different types of signals in a cognitive cycle of the agent. Mental saccades may not relate to a similar mechanism in human mind, however, such mechanism is useful for the computational implementation of machine consciousness.

3.23 Information flows and discounting in planning and learning

Naftali Tishby (The Hebrew University of Jerusalem, IL)

We argue that information seeking and reward seeking behavior should be optimized together in a "free-energy" minimization, that combines physical rewards with information gains and costs. The local information gain includes both the control/decision complexity and information provided by the environment responses (Tishby & Polani, 2009). One interesting interpretation of this framework is that information terms can be considered as internal rewards while physical (standard) rewards are external to the organism. Perfect adaptation to the environment happens when these two reward systems are proportional to each other. This can happen either through adaptation of the (subjective, belief states) transitions probabilities or of the external rewards (or both).

In this talk I will discuss another new application of this framework – information discounting. The Bellman equations for information have no explicit natural discounting. But we know that information about future event is sub-extensive for any stationary environment (Bialek, Nemenman, Tishby 2002).

This suggests that information gains must also be discounted, either exponentially (for finite dimensional environment, when predictive information grows logarithmically) or hyperbolically for infinite dimensional environments (when predictive information grows like a power-law) to be consistent with the predictive information decay. This can provide a new principled approach for general reward discounting in exploration/exploitation models.

3.24 How should Robbie search for my lost cup?

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

I begin with the workshop themes of curiosity, exploration and learning. I present a view of curiosity as attention to the relevant and ignoring the irrelevant. The model of Selective Turning is briefly presented. For exploration, I show an algorithm to search for an object in an unknown room.

Examples of the performance of the algorithm are included. A comparative experiment testing performance characteristics given different search policies, different starting points and different object placement, all with or without prior knowledge is presented. In all cases, the policy of maximising the probability of finding the object while minimizing distance travelled is superior. The presentation ends by pointing to a number of areas where learning procedures would be helpful and lead to a more robust process.

3.25 Finding Intrinsic Rewards by Embodied Evolution and Constrained Reinforcement Learning

Eiji Uchibe (Okinawa Institute of Science and Technology, JP)

Finding the design principle of reward functions is a big challenge both in artificial intelligence and neuroscience. Successful acquisition of a task usually requires not only the rewards for goals, but also for intermediate states to promote effective exploration. We propose a method to design 'intrinsic' rewards of autonomous robots by combining constrained policy gradient reinforcement learning and embodied evolution. To validate the method, we use the Cyber Rodent robots, in which collision avoidance, recharging from battery pack, and 'mating' by software reproduction are three major 'extrinsic' rewards. We show in hardware experiments that the robots can find appropriate 'intrinsic' rewards for the vision of battery packs and potential mating partners to promote approach behaviors.

3.26 Goal-Directed Action Generalization

Ales Ude (Jozef Stefan Institute – Ljubljana, SI)

Acquisition of new sensorimotor knowledge by imitation is a promising paradigm for robot learning. To be effective, action learning should not be limited to direct replication of movements obtained during training but must also enable the generation of actions in situations a robot has never encountered before. In this talk I will present a methodology that enables the generalization of the available sensorimotor knowledge. New actions are synthesized by the application of statistical methods, where the goal and other characteristics of an action are utilized as queries to create a suitable control policy, taking into account the current state of the world. Nonlinear dynamic systems are employed as a motor representation. The proposed approach enables the generation of a wide range of policies without requiring an expert to modify the underlying representations to account for different task-specific features and perceptual feedback.

3.27 Robot Visual Learning

Markus Vincze (TU Wien, AT)

Allowing a robot to acquire 3D object models autonomously not only requires robust feature detection and learning methods but also mechanisms for guiding learning and assessing learning progress. In this talk we presented probabilistic measures for observed detection success, predicted detection success and the completeness of learned models, where learning is incremental and online.

This allows the robot to decide when to add a new keyframe to its view-based object model, where to look next in order to complete the model, predicting the probability of successful object detection given the model trained so far as well as knowing when to stop learning.

3.28 Active information gathering in vision and manipulation

Jeremy L. Wyatt (University of Birmingham, GB)

We can think of vision as a process by which an agent gathers information which is relevant to a task. In this approach the value of perceptual information is anchored in the additional reward the information allows you to get in a task. I describe two pieces of work. The first is the use of POMDPs to plan task specific visual processing. In this we show that by planning to reduce uncertainty we get visual routines for specific tasks that are more reliable and faster than simply running all visual operators. In the second piece of work we pose gaze allocation in a similar reward based framework. In this work camera movements generate images that reduce uncertainty about the locations of objects and places involved in a manipulation task. In a reward-based framework the robot picks the camera movements that generate the most additional expected reward for the motor behaviours (here grasps) by reducing the relevant uncertainty. Thus both these pieces of work show how perception can be viewed as an information gathering process that can be grounded via the notion of task rewards.

3.29 Embodied Learning of Qualitative Models

Jure Zabkar (University of Ljubljana, SI)

People most often reason qualitatively and if the goal of AI is to mimic human intelligence, the robots should also be equiped with qualitative reasoning (QR) capabilities. QR is based on approximate understanding of functional relationships between the quantities. Qualitative models are less precise than numerical ones but more robust (noise resistant) and high-level (therefore simpler and easier to understand). In this talk, I will present some algorithms for learning qualitative models and their applications in embodied systems. Most of this work was carried out in collaboration with Janez Demsar, Martin Mozina and Ivan Bratko and supported by FP6 project XPERO.

4 Working Groups and Open Problems

4.1 Group 1: Self-motivation

The membership of this group was: Andy Barto, Daniel Polani, Tali Tishby, Aaron Sloman, and Jürgen Schmidhuber. This group was concerned with frameworks for understanding intrinsic motivation. The key questions addressed included:

- What is necessary for computational accounts of intrinsic motivation?
- How do different computational theories overlap and how do they differ?
- Are there universal principles?
- How do computational theories relate to what psychologists have said about intrinsic motivation and what behavioral data show?

The core discussion centered around low-level (and essentially information-theoretic) concepts to characterize intrinsically motivated behaviour - we did not discuss higher-level concepts. Sloman wondered if key issues were not addressable using information theory due to its apparent lack of semantics. Tishby argued that this view arose from a widespread misunderstanding of the concept of the concept of information. In fact, semantics, so the argument goes, can be put in via the concept of "relevant" and "valuable" information which carries information about the task to do. Additional structure can be added via the embodiment and the ensuing preferred information channels by which an agent interacts with the environment. This physical structure imposes constraints on the form of the information processing/propagation – thus, "not all information is created equal".

Particular topics and questions that were identified as being of future interest to the community included:

1. Empowerment: projecting potential action sequence into states that can be reached in the near future. Andy suggested a match with the ''drive for mastery" hypothesized in psychology.
2. Predictive Information: Entropy is related to predictability. Der and Ay use this for generating sensitive and predictive behaviour online. Complexity arises because predictive information makes the future predictive from the past, but the past needs to be rich at the same time. A 'poor' past will have no information to predict about the future.
3. Other intrinsic motivation measures: there are many more heuristic measures such as learning progress and the autotelic principle. These capture aspects of what one expects from self-motivated behaviour (including Csíkszentmihályi's notion of 'flow'). They still, however, contain still some arbitrary aspects which one may want to avoid.
4. Compression: in Schmidhuber's compression scheme the machine attempts to compress the entire past. The expected compression progress is used as a driver for 'artificial curiosity' and 'satisfaction' of the success criterion. One open question is how does the machine provide any externally visible abstraction? A second issue is whether the machine is to be limited - this is plausible in view of realistic machines and it is a Kolmogorov analogue of Shannon-type constraints and thus fits the general philosophy well: abstraction as a response to constraints: "making a virtue out of necessity".
5. Valuable information: Tishby has proposed the idea of valuable information. This brings us back to the Shannon picture of information. In this framework the machine needs to remember or compress only what is valuable: this is a very key point. There is no need to take in all information, only what will contribute to future reward. This is bounded by what can at all be predicted in the future, i.e. predictive information.

There was a more extensive discussion around the idea of valuable information, and it's relation to other concepts. If we consider the entropy of a series of a system through time, there is an extensive part (that grows linearly with time and is "uninteresting") and a subextensive part (convex, less-than-linear) component. Predictive information, the information that the past at some point has about the future, happens to be the same as the subextensive part of the entropy of a system. However, compressed info may not be unpackable, i.e. the information may be compressed in such a way that some external agent may unpack it (this must be possible, or it could not be called compression), but it may no longer be accessible to the agent itself which may be not complex enough to unpack the message (at least that's how Polani understands it). Barto thought that this is the same idea that comes up in the deterministic case as the idea of a homomorphic image of a dynamical system: Given some variables that are of interest, what is the minimal state representation that can provide enough information to 1) determine itself at the next time step, and 2) determine the values of the variables of interest. On that level, this seems to largely agree with Tishby's relevant/valuable information (via information bottleneck), in the deterministic case. Again, however, it is important to note that in Tishby's framework only part of the past is important. Barto wondered if indefinite memory systems might violate this: e.g., setting a bit in the arbitrary past can influence the present. Tishby and Polani argued that that is still just a part of the past. Polani says it is either known to the agent (then just part of the memory) or will play a role in the future, and then discovered to be an "inconsistency" that has not yet been observed. The concept of natural cut-offs was then discussed. These are the finite life-spans of organisms that force exploitation and they create natural time-scales and cut-offs for learning phases. Finally it was asked what would an agent do using Tishby's scheme if it had no value: i.e. we only use information gain as a source of reward? After some discussion, Tali concluded that it would act to efficiently gain information about state transitions.

Other areas of related work were also identified. Mike Duff's work on Bayesian optimal exploration and Bellman's "curse of the expanding grid" was identified as relevant. In this case exploration is posed as a reasoning in a meta-MDP whose states consist of physical states composed with beliefs states about model parameters (e.g., transition probabilities). Duff used the decreases in the variance of belief states as a reward function for doing RL in the meta state. How is this related to Tishby's ideas? It was agreed that while there were similarities that Tishby's approach is distinguished from Bayesian approaches.

Summary: In conclusion much of the discussion boiled down to the issue of whether we desire information-seeking behaviour or value-seeking behaviour. Two important open questions concern whether these are really similar? If not then when not?

4.2 Group 2: Scaling methods for exploration and curiosity

This break out group (Jeremy Wyatt, Eiji Uchibe, Kristian Kersting, Marc Hanheide, Stefan Elfwing, Frank Sehnke, Leo Pape, Jure Zakbar) was concerned with the issue of how to scale existing methods. It seemed to us that there were several challenges to scaling techniques for exploration and uncertainty to larger problem domains and more complex robot systems. Each of these defines a open challenge for the research community. Each is detailed in turn.

A. How to integrate common sense and structured knowledge with uncertainty?

It is an old AI chestnut that common sense knowledge can speed up reasoning and make action more efficient by restricting our space of reasonable actions. However, since many of the approaches for exploration control are statistical, capturing common sense knowledge means we need to be able to express it in some similar manner. In the past five years ontology based reasoning has been applied in robots (e.g. the work of Saffioti (Orebro), Beetz et al. (TU Munich)). This work has mostly been applied in making inferences about objects and rooms in topological and semantic SLAM. Work by Kuipers et al uses a map with several levels of abstraction (such as metric, topological, and semantic, sometimes referred to as the spatial semantic hierarchy). The top level typically represents the properties of objects and places and the relations between them as a labeled graph, where the labels can be inferred partly based on an ontology capturing common sense knowledge by using is-a and has-a relationships. Jensfelt et al have now shown how to attach probabilities to such a map. The representation then becomes a chain graph. Hanheide et al showed in the seminar how to reason with such a map, both to infer the probabilities of particular labels and also to plan courses of exploratory action. An alternative approach by Zender et al is to use non-monotonic reasoning to draw and withdraw inferences in a logical framework.

Open questions include:

1. What are the advantages and disadvantages of the logical and probabilistic approaches, particularly concerning the time complexity of inference and the reliability of the results? Does probabilistic inference scale badly with the size of the graph representing default knowledge?
2. How should probabilities be incorporated into such representations of categorical knowledge in a principled manner? Is it feasible to expect us to be able to populate such complex structures with probabilities grounded in real data prior to operation? What do google – vision/machine learning type approaches have to offer us here?

B. Are Partially Observable models unnecessarily hard?

Partial observability as modelled by the stochastic process community, e.g. POMDPs is both a blessing and a curse. It provides us with a rather general framework for reasoning under uncertainty about the value of information. But even the most benign classes of POMDPs are at best in NP, and even practical approximate algorithms for them scale only to hundreds of states. In addition learning in POMDPs is hard. This begs the question are POMDPs too general a formulation, Here we felt there were two issues. What proofs can we provide about existing simple approaches such as greedy approaches to information gathering? Second what alternatives are there to POMDPs as a formalism, including special cases that are tractable. We deal with each in turn:

Open questions include:

1. Performance bounds for simple methods: Work on topics like sub-modularity and adaptive sub-modularity shows that it is sometimes possible to prove that we can do well with simple algorithms. There are two main open questions with sub-modularity for exploration:
   a. Does some form of sub-modularity apply where the experiments depend on the state, and where the state can be changed by the robot's actions?
   b. Does some form of sub-modularity apply where the robot physically alters the external world state, e.g. by moving objects, or alternatively alters the belief states of other agents (e.g. via dialogue)?
2. More restrictive models: Some workers (e.g. Kaelbling) have proposed dividing their problems up into a probabilistic bit and a non-probabilistic bit, e.g. just using POMDPs to model small patches of state uncertainty in an otherwise observable world. Is this an approach in which it is possible to derive bounds on how sub-optimal performance is? Work by Hanheide et al presented at the seminar also touched on algorithms that split

the exploration problem into a probabilistic and a non-probabilistic part. What kinds of algorithms are there that take broadly this approach?

C. Will data based machine learning make the problem go away?

One major motivation for autonomous exploration is because of the limited experience of the robots that we use in practice. However, there is a lot of work on learning for vision from the vast amounts of data available on the internet. Is it actually possible that very large available data sets, e.g. google for vision or very large scale SLAM will make much of the exploration problem go away? In other words we asked the question whether in the long run there are solutions other than autonomous exploration for filling out the knowledge bases of our robots sufficiently well to enable them to perform a wide variety of tasks.

D. How hard can open worldness be in practice?

Modern planning approaches are very efficient for closed worlds. One difficulty when applying such planning approaches to exploration problems is that for many exploration problems the robot has to reason about open worldness. This could be reasoning about the addition of new objects or places as they are found, or reasoning about changes to the types of entities it may have in its knowledge base. There are two ways to deal with this:

1. Soft open worldness: We try to retain the benefits of efficient algorithms for reasoning in closed worlds by allowing only a finite amount of open worldness, e.g. a finite number of additional objects or categories in the planning domain, each of which have some unknown properties. These are essentially slots that we can fill via planning actions.
2. Hard open worldness: we essentially abandon current efficient methods for reasoning in closed worlds and allow the planning domain to be modified on the fly during the planning process. This would require a whole new branch of planning research.

E. How do we fairly compare apples and oranges as the varieties of exploration opportunities rise in number?

One of the major themes emerging from the seminar is the issue of how intrinsic and extrinsic rewards are fairly compared. Tishby presented one way to approach this problem at the seminar. However, there is also the issue of how various sources of intrinsic reward should be fairly compared while leading to efficient learning. The frameworks of Polani and Schmidhuber address these problems, but since the only extant empirical results are for grid worlds or similar, we are not yet clear about how the data efficiency of these as the number of possible learning activities increases in real robot worlds. In robot implementations at the moment the comparison is done heuristically using fudge factors. This is unsatisfactory. This lead us to the general open question:

1. As possibilities for learning activities increase do existing ways of comparing value break. How can we fairly, and yet tractably compare very different learning goals as the number of different types of goal rises? It seemed to us there were two possible answers:
   a. Ground rewards for information types in a known distribution over future tasks.
   b. The notions of empowerment and compression are perfectly sufficient and efficient.

4.3 Group 3: Optimality and sub-optimality in exploration

Group members were Peter Auer, Peter Dayan, Richard Dearden, Charles Fox, Mark Humphries, Andreas Krause, Manuel Lopes, and Ruben Martinez-Cantin.

Solved problems: Although more can always be done, there are a number of fairly strong results for bandit problems and MDPs:

1. Bandits: there are Gittins indices in a (brittle) Bayesian case; regret bounds with stochastic arms (at log t cost), or with adversarial arms ($\sqrt{t}$ cost); good ideas for continuous bandit problems with smoothness assumptions (e.g. Gaussian process payoffs) or contextual bandits (with useful 'input' information; the regret bounds being norm dependent). Most methods are based on forms of optimism. Of course, not everything about bandits is solved: particularly for complex action spaces.
2. Reinforcement learning in MDPs: the two main successful models are regret-like bounds, which achieve regret of $\sqrt{t}$ against the optimal policy in the non-adversarial setting, and phasic (E^3) or continuously exploring R_{max} , which provide guarantees about how long it takes to provide a near-optimal policy, without worrying about the regret along the way. There are also many results on MDPs with function approximation – key here is whether the optimal value function is in the approximation class; there are then results on the loss compared to an optimal policy. Most methods require knowledge of the maximum reward or policy payoff (or tricks to get around this).
3. Continuous time/space linear quadratic stochastic control: these problems are also easy.
4. Information gathering POMDPs: in which actions do not affect the state at all, may have special solutions (like the SPRT and cases of submodularity); extensions to planning informative sampling trajectories might be possible.

Unsolved problems: These fall into different categories. First are problems that are well-accepted by the community as being hard, and as having attracted substantial thought, if not results:

1. POMDPs: we couldn't think of any generally successful approach to POMDPs. Indeed, it was reported that not only are POMDPs NP-hard, but also approximations to POMDPs are NP-hard; there is no simple, but non-trivial, subclass of POMDP that is easy, apart perhaps for information gathering POMDPs; and methods such as factored representations that work well for MDPs, are much less effective for POMDPs. Note that even the planning problem is hard for general POMDPs, given knowledge of the model.
2. Continuous time problems: most work has been done in discrete time problems. There is much less known about exploration in continuous time problems, notably semi-Markov decision-problems (SMDPs), for instance when agents can choose when to act. Problems that haven't been fully explored or perhaps completely precisely formulated:
3. Problem classes: we agreed that there is a need for a suite of benchmark problems that probe different aspects of the difficulty of exploration in various ways.
4. Multi-task learning: many issues surrounding curiosity or the early acquisition of competence could be illuminated in a multi-task setting, with the statistical structure of the commonality between tasks determining what the benefits could be of the competence. We therefore need models for this statistical structure and generalization.
5. Satisficing: we had some discussion about the effect of resource limitations in space and/or time on finding and executing optimal policies. There are some models for speed-accuracy tradeoffs in simple information-gathering tasks (as in the SPRT), but they don't seem to have been generalized to MDPs. Actual bounds on resources (e.g. different polynomial orders) can depend sensitively on the model of computation supported by the hardware, and so are hard to assess. There is work in general on parallelizability (circuit complexity), but we didn't know of work on MDPs.

6. Risk-sensitivity: there is some work on risk-sensitive solutions – using different norms, or penalizing variance and higher order moments as well as means (Singh; Kappen). This could also be formulated in terms of multiple objectives ('maximize reward, but don't get near to the following states'). Observations
7. Computationally simple heuristics like entropy reduction can work very well, at least as well as much more complex methods such as depth-limited model-search. Characterizing when simple heuristics work well is a key, if somewhat amorphous, research question.
8. For problems satisfying submodularity, being greedy about choice can be near-optimal. In some applications, the main objective function of interest may not be submodular. In these applications, it may be of interest to identify surrogate objectives, which are submodular and thus can be greedily optimized, while still approximating the true objective. The relevant submodular functions may not be as simple as entropy reduction.
9. There was some debate about the evolutionary importance of certain sorts of learning and planning. Animals are endowed with extremely strong heuristics (rats act as if things above them are threats; things below them are edible; we have complex ways of assigning salience to sensory inputs) – perhaps these, and other sources of priors such as environmental affordances of object classes, could provide a powerful framework for avoiding doing too much learning.
10. Nobody knew of quantum algorithms for fast (PO)MDP solving . . .

4.4 Group 4: Unifying themes

The working group consisted of John Tsotsos, Tony Cohn, Janusz Starzyk, Danijel Skočaj, and Aleš Leonardis.

Summary: The working group discussed more general issues such as:

1. What do the approaches described in the talks have in common?
2. Is there a unifying framework or topology for the approaches described in the talks?
3. What are the open questions and solution classes?

A lively exchange of diverse opinions converged on a consensus that two themes have been prevalent at the seminar, namely, a rather theoretical treatment of reinforcement learning and, on the other hand, presentations of real application scenarios. While the former excel in abstract models and rigorous treatment of computational issues, the examples shown to support the theoretical findings are often based on unrealistic assumptions and limited to graphical animations and simulations. On the other hand, real applications demonstrated how noise and various constraints, stemming from the real-world, can be tackled, but then most of the approaches relied on heuristics and provided almost no theoretical bounds or proofs on the performance. To bring closer together the theory and the practice should be one of the main goals of the field over the next years.

In a similar vein the working group also discussed a dichotomy between general theoretical approaches on the one hand, and specific (cased-based) applications on the other. A translation of general theoretical principles (such as information gain as the driving force for award driving curiosity and exploration) into specific application domains (such as manipulation, navigation, spatial mapping) appears to be hard and significant efforts towards more generalized solutions should be sought.

Finally, the working group also acknowledged a positive influence that interdisciplinary research and scientific insights in neuroscience and psychology can have on the field. In particular, developmental psychology could provide some indications towards the answers to the questions such as “how / when / under which circumstances does the curiosity arise?”, “what are the requirements?”, “what is the role of scaffolding?”.

In summary, the following open questions have been identified:

1. How do we bridge the gap between general theory and specific real-world applications?
2. How do we bridge the gap between low-level (signal, sensor driven approaches) methods and high-level (symbolic, semantic based methods)?
3. How do we achieve generalization and scalability (e.g. shareability among tasks)?
4. What is the relationship between the theories in engineering science and the neuroscience?

Participants

- Peter Auer
Montan-Universität Leoben, AT
- Andrew Barto
University of Massachusets – Amherst, US
- Anthony G. Cohn
University of Leeds, GB
- Peter Dayan
University College London, GB
- Richard W. Dearden
University of Birmingham, GB
- Yiannis Demiris
Imperial College London, GB
- Stefan Elfwing
Okinawa Institute of Science and Technology, JP
- Charles Fox
University of Sheffield, GB
- Martin A. Giese
Univ.klinikum Tübingen, DE
- Marc Hanheide
University of Birmingham, GB
- Mark D. Humphries
ENS – Paris, FR
- Kristian Kersting
Fraunhofer IAIS – St. Augustin, DE
- Andreas Krause
ETH Zürich, CH
- Ales Leonardis
University of Ljubljana, SI
- Manuel Lopes
INRIA - Bordeaux, FR
- Ruben Martinez-Cantin
CUD – Zaragoza, ES
- Leo Pape
IDSIA - Lugano, CH
- Jan Peters
MPI für biologische Kybernetik – Tübingen, DE
- Daniel Polani
University of Hertfordshire, GB
- Jürgen Schmidhuber
IDSIA – Lugano, CH
- Frank Sehnke
ZSW – Stuttgart, DE
- Danijel Skocaj
University of Ljubljana, SI
- Aaron Sloman
University of Birmingham, GB
- Janusz Starzyk
Ohio University, US
- Naftali Tishby
The Hebrew University of Jerusalem, IL
- John K. Tsotsos
York University – Toronto, CA
- Eiji Uchibe
Okinawa Institute of Science and Technology, JP
- Ales Ude
Jozef Stefan Institute – Ljubljana, SI
- Patrick van der Smagt
DLR Oberpfaffenhofen, DE
- Markus Vincze
TU Wien, AT
- Jeremy L. Wyatt
University of Birmingham, GB
- Jure Zabkar
University of Ljubljana, SI