Report from Dagstuhl Seminar 19361

Logic and Learning

Edited by

Michael Benedikt[1], Kristian Kersting[2], Phokion G. Kolaitis[3], and Daniel Neider[4]

1 **University of Oxford, GB,** michael.benedikt@gmail.com
2 **TU Darmstadt, DE,** kersting@cs.tu-darmstadt.de
3 **University of California – Santa Cruz and IBM Almaden Research Center – San Jose, US,** kolaitis@soe.ucsc.edu
4 **MPI-SWS – Kaiserslautern, DE,** neider@mpi-sws.org

Abstract

The goal of building truly intelligent systems has forever been a central problem in computer science. While logic-based approaches of yore have had their successes and failures, the era of machine learning, specifically deep learning is also coming upon significant challenges. There is a growing consensus that the inductive reasoning and complex, high-dimensional pattern recognition capabilities of deep learning models need to be combined with symbolic (even programmatic), deductive capabilities traditionally developed in the logic and automated reasoning communities in order to achieve the next step towards building intelligent systems, including making progress at the frontier of hard problems such as explainable AI. However, these communities tend to be quite separate and interact only minimally, often at odds with each other upon the subject of the "correct approach" to AI. This report documents the efforts of Dagstuhl Seminar 19361 on "Logic and Learning" to bring these communities together in order to: (i) bridge the research efforts between them and foster an exchange of ideas in order to create unified formalisms and approaches that bear the advantages of both research methodologies; (ii) review and analyse the progress made across both communities; (iii) understand the subtleties and difficulties involved in solving hard problems using both perspectives; (iv) make attempts towards a consensus on what the hard problems are and what the elements of good solutions to these problems would be.

The three focal points of the seminar were the strands of "Logic for Machine Learning", "Machine Learning for Logic", and "Logic vs. Machine Learning". The seminar format consisted of long and short talks, as well as breakout sessions. We summarise the motivations and proceedings of the seminar, and report on the abstracts of the talks and the results of the breakout sessions.

Seminar September 1–6, 2019 – http://www.dagstuhl.de/19361
2012 ACM Subject Classification Computing methodologies → Artificial intelligence, Computing methodologies → Machine learning approaches, Computing methodologies → Logical and relational learning, Theory of computation → Constraint and logic programming, Theory of computation → Models of learning, General and reference → Surveys and overviews, Theory of computation → Logic
Keywords and phrases Artificial Intelligence, Automated reasoning, Databases, Deep Learning, Inductive Logic Programming, Logic, Logic and Learning, Logic for Machine Learning, Logic vs. Machine Learning, Machine Learning, Machine Learning for Logic, Neurosymbolic methods
Digital Object Identifier 10.4230/DagRep.9.9.1
Edited in cooperation with Adithya Murali

1 Executive Summary

Michael Benedikt (University of Oxford, GB)
Kristian Kersting (TU Darmstadt, DE)
Phokion G. Kolaitis (University of California – Santa Cruz & IBM Almaden Research Center – San Jose, US)
Adithya Murali (University of Illinois – Urbana-Champaign, US)
Daniel Neider (MPI-SWS – Kaiserslautern, DE)

Motivation

Logic and learning are central to Computer Science, and in particular to AI research and allied areas. Alan Turing envisioned, in his paper "Computing Machinery and Intelligence" [1], a combination of statistical (*ab initio*) machine learning and an "unemotional" symbolic language such as logic. However, currently, the interaction between research in logic and research in learning is far too limited; in fact, they are often perceived as being completely distinct or even opposing approaches.

While there has been interest in using machine learning methods within many application areas of logic, the investigation of these interactions has usually been carried out within the confines of a single problem area. We believe that an interaction involving a broader perspective is needed. It would be fruitful to look for common techniques in applying learning to logic-related tasks, which requires looking across a wide spectrum of applications. It is also important to consider the ways that logic and learning, deduction and induction, can work together.

Design of the Seminar

The main aim of this Dagstuhl Seminar was to address the above problems by bring researchers from the logic and learning communities together and to create bridges between the two fields via the exchange of ideas ranging between the (seemingly) polar possibilities of the injection of declarative methods in machine learning and the use and applications of learning technologies in logical contexts. This included creating an understanding of the work in different applications, an increased understanding of the formal connections between these applications, and the development of a more unified view of the current attempts to organically reconcile deductive and inductive approaches. In order to structure these explorations, the focal points of the seminar were the following three distinct strands of interaction between logic and learning:

1. *Machine Learning for Logic*, including the learning of logical artifacts, such as formulas, logic programs, database queries and integrity constraints, as well as the application of learning to tune deductive systems.
2. *Logic for Machine Learning*, including the role of logics in delineating the boundary between tractable and intractable learning problems, the construction of formalisms that allow learning systems to take advantage of specified logical rules, and the use of logic as a declarative framework for expressing machine learning constructs.
3. *Logic vs. Machine Learning*, including the study of problems that can be solved using either logic-based techniques or via machine learning, an exploration of the trade-offs between these techniques, and the development of benchmarks for comparing these methods.

Summary of seminar activities

The seminar was attended by 41 researchers across various communities including logic, databases, Inductive Logic Programming (ILP), formal verification, machine learning, deep learning, and theorem proving. The membership consisted of senior and junior researchers, including graduate students, post-doctoral researchers, and industry experts. The seminar was conducted through talks and breakout sessions, with breaks for discussion between the attendees. There were three long talks, 21 short talks, and three breakout sessions on the discussion of open problems in logic and learning.

The talks consisted of: (i) presentation of recent advances in research questions and methodologies relating to the motivations discussed above; (ii) surveys of the state of research on various problems requiring the combination of deductive and inductive reasoning as well as methodologies developed to address fundamental hurdles in this space; (iii) new perspectives on the organic combination of logical formulations and methods with machine learning in specific application domains; (iv) theoretical formulations and results on problems in learning logical representations; (v) demonstrations of state-of-the art tools combining logic and learning for applications such as theorem proving or entity resolution; (vi) presentation of research on challenge problems for the field of AI and intelligent reasoning.

The breakout sessions were conducted in three continuing parts, each spanning one session. The first part involved all the participants in a discussion of the current (small and large) open problems in AI, challenge problems for the field of intelligent systems, and research questions about defining specific goals representing a successful combination of inductive and deductive reasoning. This involved a deliberation of what problems were relevant, which problems could be potentially related to or dependent upon each other, and various suggestions to formalise commonly desired research goals. This session resulted in the choice of three broad areas for further specific discussion: (i) Explainable AI (ii) Injecting symbolic knowledge or constraints into neural networks, and (iii) Learning of logical formulae (first-order logic) from satisfaction on structures in a differentiable manner. The second part consisted of parallel thematic sessions on these three areas. Each thematic session was conducted in the form of a round-table discussion and was led by one or two participants who championed the theme. The third session brought all the participants together again to conclude with a summary of the ideas exchanged during the parallel sessions.

Conclusion

We consider the seminar a success. There is a growing need to enable the disparate communities of logic and learning to interact with each other, and we noted from the seminar that researchers from each community appreciated the perspective offered by the other, often identified techniques used by the other community that could be imported into their own, and, interestingly, were in agreement about the relevant and important problems of the day. The format of the seminar including ample time for discussions and breakout sessions received positive feedback from the participants.

References

1 A. M. Turing, "Computing machinery and intelligence", Mind, vol. LIX, pp. 433–460, October 1950

2 Table of Contents

3 Overview of Talks

3.1 Six perspectives on logic & learning (in infinite domains)

Vaishak Belle (University of Edinburgh, GB)

The unification of low-level perception and high-level reasoning is a long-standing problem in artificial intelligence, and among other approaches, the integration of logic and learning potentially offers the most general solution to that problem. Although there has been considerable progress on this integration, models in practise continue to make the finite domain assumption, and so models are essentially propositional, programs are loop-free, and so on. In this talk, we discuss a number of different ways in which the infinite is embraced. In recent work, for example, we have looked at the problems of inference and (parameter and structure) learning in continuous domains, that is, where logical atoms model continuous properties. In other work, we report on the synthesis of plans with loops in the presence of probabilistic nondeterminism. Finally, we touch on proposals for declaratively modelling logical reasoning, probabilistic inference and learning problems in continuous domains.

This talk reports on joint work with a number of collaborators and is drawn from the following papers: [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11].

References

1 Amelie Levray, Vaishak Belle, “Learning Tractable Probabilistic Models in Open Worlds”, CoRR, abs/1901.05847, 2019.
2 Laszlo Treszkai, Vaishak Belle, “A Correctness Result for Synthesizing Plans With Loops in Stochastic Domains”, CoRR, abs/1905.07028, 2019.
3 Vaishak Belle, “Abstracting Probabilistic Models”, CoRR, abs/1810.02434, 2018.
4 Vaishak Belle, Luc De Raedt, “Semiring Programming: A Framework for Search, Inference and Learning”, CoRR, abs/1609.06954, 2016.
5 Vaishak Belle, Brendan Juba, “Implicitly Learning to Reason in First-Order Logic”, NeurIPS 2019.
6 Stefanie Speichert, Vaishak Belle, “Learning Probabilistic Logic Programs in Continuous Domains”, 29th International Conference on Inductive Logic Programming (ILP), 2019.
7 Andreas Bueff, Stefanie Speichert, Vaishak Belle, “Tractable Querying and Learning in Hybrid Domains via Sum-Product Networks”, Workshop on Hybrid Reasoning and Learning (HRL), KR, 2018.
8 Samuel Kolb, Martin Mladenov, Scott Sanner, Vaishak Belle, Kristian Kersting, “Efficient Symbolic Integration for Probabilistic Inference”, Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18), Main track, Pages 5031-5037, July 2018.
9 Vaishak Belle, “Open-Universe Weighted Model Counting”, Thirty-First AAAI Conference on Artificial Intelligence, 2017.
10 Davide Nitti, Vaishak Belle, Tinne Laet, Luc De Raedt, Machine Learning, Volume 106 Issue 12, Pages 1905-1932, December 2017.
11 Vaishak Belle, Andrea Passerini, Guy Van Den Broeck, “Probabilistic inference in hybrid domains by weighted model integration”, Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI’15), Pages 2770-2776, July 2015.

3.2 Neural Model Counting

Ismail Ilkan Ceylan (University of Oxford, GB)

Joint work of Ralph Abboud, Ismail Ilkan Ceylan, Thomas Lukasiewicz
Main reference Ralph Abboud, İsmail İlkan Ceylan, Thomas Lukasiewicz: "Learning to Reason: Leveraging Neural Networks for Approximate DNF Counting", CoRR, Vol. abs/1904.02688, 2019.
URL https://arxiv.org/abs/1904.02688

Weighted model counting (WMC) has emerged as a prevalent approach for probabilistic inference. In its most general form, WMC is #P-hard and, as a result, solving real-world WMC instances is intractable. Weighted DNF counting (weighted #DNF) is a special case where approximations with probabilistic guarantees can be tractably obtained, but this requires time O(mn), where m denotes the number of variables and n the number of clauses of the input DNF. In this talk, I will present a novel approach for weighted #DNF that combines approximate model counting with deep learning and accurately approximates model counts in just O(m + n). Our experiments show that our model learns and generalizes very well to large-scale #DNF instances.

3.3 Learning Constraints from Examples

Luc De Raedt (KU Leuven, BE)

Joint work of Luc De Raedt, Andreas Passerini, Stefano Teso
Main reference Luc De Raedt, Andrea Passerini, Stefano Teso: "Learning Constraints From Examples", in Proc. of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018, pp. 7965–7970, 2018.
URL https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17229

While constraints are ubiquitous in artificial intelligence and constraints are also commonly used in machine learning and data mining, the problem of learning constraints from examples has received less attention. In this talk I shall discuss the problem of constraint learning in detail, indicate some subtle differences with standard machine learning problems, sketch some applications and summarize the state-of-the-art.

3.4 Query Learning of Omega Regular Languages

Dana Fisman (Ben Gurion University – Beer Sheva, IL)

Joint work of Dana Fisman, Dana Angluin, Timos Antonopoulos, Udi Boker
Main reference Dana Fisman: "Inferring regular languages and ω-languages", J. Log. Algebr. Meth. Program., Vol. 98, pp. 27–49, 2018.
URL https://doi.org/10.1016/j.jlamp.2018.03.002

Omega languages, i.e. languages of infinite words (or of infinite trees), play an important role in modeling, verification and synthesis of reactive systems. While query learning of regular languages of finite words can be done in polynomial time using a polynomial number of membership and equivalence queries, there is no known polynomial learning algorithm for the full class of omega regular languages. In this talk we will discuss the obstacles in obtaining a polynomial learning algorithm and go through state-of-the art results on learning of regular languages of infinite words and of infinite trees.

3.5 Bounds in Query Learning

James Freitag (University of Illinois – Chicago, US)

Joint work of Hunter Chase, James Freitag
Main reference Hunter Chase, James Freitag: "Bounds in Query Learning", CoRR, Vol. abs/1904.10122, 2019.
URL https://arxiv.org/abs/1904.10122

I will discuss some bounds in query learning related to combinatorial quantities isolated in model theory, namely, Littlestone dimension and consistency dimension. These quantities are related to exact learning by equivalence queries and learning by equivalence queries and membership queries. Both quantities were also isolated in model theory (with different names), but can be formulated in a purely combinatorial manner. I will also discuss other potential connections between combinatorial notions from model theory and various settings of learning.

3.6 Learning Logically Specified Problems

Martin Grohe (RWTH Aachen, DE)

Joint work of Martin Grohe, Gyorgy Turan, Martin Ritzert
Main reference Martin Grohe, Martin Ritzert: "Learning first-order definable concepts over structures of small degree", in Proc. of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017, pp. 1–12, IEEE Computer Society, 2017.
URL https://doi.org/10.1109/LICS.2017.8005080

After some general remarks about learning frameworks for logical specifications, applications scenarios, and practical challenges, I will focus on a declarative model theoretic learning framework. Within this framework, I will talk about recent positive and negative learnability results that we obtained for learning models specified in first-order and monadic second-order logic.

3.7 On Learning to Prove

Daniel Huang (University of California – Berkeley, US)

Main reference Daniel Huang: "On Learning to Prove", CoRR, Vol. abs/1904.11099, 2019.
URL https://arxiv.org/abs/1904.11099

In this talk, we consider the problem of learning a first-order theorem prover that uses a representation of beliefs in mathematical claims to construct proofs. The inspiration for doing so comes from the practices of human mathematicians where"plausible reasoning" is applied in addition to deductive reasoning to find proofs. Towards this end, we introduce a representation of beliefs that assigns probabilities to the exhaustive and mutually exclusive first-order possibilities found in Hintikka's theory of distributive normal forms. The representation supports Bayesian update, induces a distribution on statements that does not enforce that logically equivalent statements are assigned the same probability, and suggests an embedding of statements into an associated Hilbert space. We then examine conjecturing as model selection and an alternating-turn game of determining consistency. The game is amenable (in principle) to self-play training to learn beliefs and derive a prover that is complete when logical omniscience is attained and sound when beliefs are reasonable. The representation has super-exponential space requirements as a function of quantifier depth so the ideas in this paper should be taken as theoretical. We will comment on how abstractions can be used to control the space requirements at the cost of completeness.

3.8 Counterexample-Guided Strategy Improvement for POMDPs Using Recurrent Neural Networks

Nils Jansen (Radboud University Nijmegen, NL)

Joint work of Steven Carr, Nils Jansen, Ralf Wimmer, Alexandru Serban, Bernd Becker, Ufuk Topcu
Main reference Steven Carr, Nils Jansen, Ralf Wimmer, Alexandru Constantin Serban, Bernd Becker, Ufuk Topcu: "Counterexample-Guided Strategy Improvement for POMDPs Using Recurrent Neural Networks", in Proc. of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019, pp. 5532–5539, ijcai.org, 2019.
URL https://doi.org/10.24963/ijcai.2019/768

We study strategy synthesis for partially observable Markov decision processes (POMDPs). The particular problem is to determine strategies that provably adhere to (probabilistic) temporal logic constraints. This problem is computationally intractable and theoretically hard. We propose a novel method that combines techniques from machine learning and formal verification. First, we train a recurrent neural network (RNN) to encode POMDP strategies. The RNN accounts for memory-based decisions without the need to expand the full belief space of a POMDP. Secondly, we restrict the RNN-based strategy to represent a finite-memory strategy and implement it on a specific POMDP. For the resulting finite Markov chain, efficient formal verification techniques provide provable guarantees against temporal logic specifications. If the specification is not satisfied, counterexamples supply diagnostic information. We use this information to improve the strategy by iteratively training the RNN. Numerical experiments show that the proposed method elevates the state of the art in POMDP solving by up to three orders of magnitude in terms of solving times and model sizes.

3.9 Implicitly Learning to Reason in First-Order Logic

Brendan Juba (Washington University – St. Louis, US)

Joint work of Vaishak Belle, Brendan Juba
Main reference Vaishak Belle, Brendan Juba: "Implicitly learning to reason in first-order logic", pp. 3376–3386, Curran Associates, Inc., 2019.
URL http://papers.nips.cc/paper/8599-implicitly-learning-to-reason-in-first-order-logic.pdf

We consider the problem of answering queries about formulas of first-order logic based on background knowledge partially represented explicitly as other formulas, and partially represented as examples independently drawn from a fixed probability distribution. PAC semantics, introduced by Valiant, is one rigorous, general proposal for learning to reason in formal languages: although weaker than classical entailment, it allows for a powerful model theoretic framework for answering queries while requiring minimal assumptions about the form of the distribution in question. To date, however, the most significant limitation of that approach, and more generally most machine learning approaches with robustness guarantees, is that the logical language is ultimately essentially propositional, with finitely many atoms. Indeed, the theoretical findings on the learning of relational theories in such generality have been resoundingly negative. This is despite the fact that first-order logic is widely argued to be most appropriate for representing human knowledge. In this work, we present a new theoretical approach to robustly learning to reason in first-order logic, and consider universally quantified clauses over a countably infinite domain. Our results exploit symmetries exhibited by constants in the language, and generalize the notion of implicit learnability to show how queries can be computed against (implicitly) learned first-order background knowledge.

3.10 DeepProbLog: Integrating Logic, Probability and Neural Networks

Angelika Kimmig (Cardiff University, GB)

Joint work of Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, Luc De Raedt
Main reference Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, Luc De Raedt: "DeepProbLog: Neural Probabilistic Logic Programming", CoRR, Vol. abs/1907.08194, 2019.
URL https://arxiv.org/abs/1907.08194

ProbLog is a probabilistic programming language that extends the logic programming language Prolog. As other probabilistic programming and statistical relational AI techniques, it supports inference and learning. It has recently been extended to incorporate also neural networks in the framework of DeepProbLog. The resulting framework tightly integrates logic, probability and neural networks and supports both learning and reasoning and the symbolic and subsymbolic level.

3.11 Learning Description Logic Concepts: Complexity and (Un)decidability

Carsten Lutz (Universität Bremen, DE)

Joint work of Maurice Funk, Jean Christoph Jung, Carsten Lutz, Hadrien Pulcini, Frank Wolter
Main reference Maurice Funk, Jean Christoph Jung, Carsten Lutz, Hadrien Pulcini, Frank Wolter: "Learning Description Logic Concepts: When can Positive and Negative Examples be Separated?", in Proc. of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019, pp. 1682–1688, ijcai.org, 2019.
URL https://doi.org/10.24963/ijcai.2019/233

Learning description logic (DL) concepts from positive and negative examples given in the form of labeled data items in a KB has received significant attention in the literature. We study the question of when a separating DL concept exists and provide useful model-theoretic characterizations as well as complexity results for the associated decision problem. For expressive DLs such as ALC and ALCQI, our characterizations show a surprising link to the evaluation of ontology-mediated conjunctive queries. We exploit this to determine the combined complexity and data complexity of separability, including a surprising undecidability result for a common DL with rather modest expressive power.

3.12 Intuitive Mathematics: Building a Proof System with Deep Reinforcement Learning

Mateusz Malinowski (Google DeepMind – London, GB)

Joint work of Alhussein Fawzi, Mateusz Malinowski, Hamza Fawzi, Omar Fawzi
Main reference Alhussein Fawzi, Mateusz Malinowski, Hamza Fawzi, Omar Fawzi: "Learning dynamic polynomial proofs", CoRR, Vol. abs/1906.01681, 2019.
URL https://arxiv.org/abs/1906.01681

Deep reinforcement learning that combines two learning paradigms is a promising method to solve complex problems that often escape the traditional formalism. With minimal domain specification and many data points, it has been shown to work effectively in the domain of complex video games. The same learning paradigm can also be used to improve the search for suitable tactics in the existing proof systems. However, I believe we can step even further and think of an end-to-end proof system. I also believe that not only theorem provers can benefit from deep reinforcement learning, but also that the former can be an excellent testbed for the latter. This talk is divided into two parts. In the first part, I will share my experience from making an algebraic proof system (I attach the corresponding paper) with deep reinforcement learning. I will mostly pay attention to 1) the reinforcement learning part, 2) incorporating inductive biases into a learned representation. The second part of my talk is more speculative. Here, I share my thoughts on building such a system that is more inspired by a human development process. The core idea relies not only on using deep reinforcement learning for more efficient search, but also to encapsulate "mathematical intuition" in the learned model.

3.13 Learning Models over Relational Databases

Dan Olteanu (University of Oxford, GB)

Joint work of Dan Olteanu, Maximilian Schleich (Oxford), Mahmoud Abo-Khamis, Ryan Curtin, Hung Q. Ngo (RelationalAI), Ben Moseley (CMU), XuanLong Nguyen (Michigan)

I will make the case for a first-principles approach to machine learning over relational databases that exploits recent development in database systems and theory. The input to learning classification and regression models is defined by feature extraction queries over relational databases. The mainstream approach to learning over relational data is to materialize the training dataset, export it out of the database, and then learn over it using statistical software packages. These three steps are expensive and unnecessary. Instead, one can cast the machine learning problem as a database problem by decomposing the learning task into a batch of aggregates over the feature extraction query and by computing this batch over the input database. Ongoing results show that the performance of this approach benefits tremendously from structural properties of the relational data and of the feature extraction query; such properties may be algebraic (semi-ring), combinatorial (hypertree width), or statistical (sampling). It also benefits from factorized query evaluation and query compilation. For a variety of models, including factorization machines, decision trees, and support vector machines, this approach may come with lower computational complexity than the materialization of the training dataset used by the mainstream approach. This translates to several orders-of-magnitude speed-up over state-of-the-art systems such as TensorFlow, R, Scikit-learn, and mlpack. While these results are promising, there is much more awaiting to be discovered.

This work is part of the FDB project.

3.14 Learning ontologies: a question-answer game

Ana Ozaki (Free University of Bozen-Bolzano, IT)

Joint work of Ricardo Duarte, Boris Konev, Carsten Lutz, Ana Ozaki, Frank Wolter

Ontologies have been applied to integrate and abstract information from multiple data sources; to describe knowledge in various domains, in particular, those related life sciences; among others. Building an ontology often requires the interaction between experts in a domain of interest and experts in modelling ontologies, called ontology engineers. We treat the problem of building an ontology as a learning problem. An ontology engineer, playing the role of the learner, attempts to build an ontology that reflects the knowledge of a domain expert (the teacher) by posing questions. This setting can be seen as an instance of Angluin's exact learning model with membership and equivalence queries. We investigate polynomial learnability for different ontology languages within this learning model and show non-polynomial learnability for ontologies formulated in the ontology language EL, and polynomial learnability for fragments of this language. We also present an implementation of an (exponential) algorithm for learning EL ontologies.

The talk will be primarily based upon the work in [1, 2].

References

1. Boris Konev, Carsten Lutz, Ana Ozaki, Frank Wolter, "Exact Learning of Lightweight Description Logic Ontologies", Journal of Machine Learning Research, Vol. 18(201), Pages 1–63, 2018.
2. Mario Ricardo Cruz Duarte, Boris Konev, Ana Ozaki, "ExactLearner: A Tool for Exact Learning of EL Ontologies", Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR, Pages 409–414, 2018.

3.15 Learning Logics, Program Synthesis, and Neural Nets

Madhusudan Parthasarathy (University of Illinois – Urbana-Champaign, US)

This talk will survey three fields that have slightly different emphasis: learning logics (learning formulas from data), program synthesis (especially using learning), and neural nets (for recognizing patterns). I will try to explore these areas, their applications we have pursued (synthesizing programs, synthesizing inductive, mining specifications), and new synergies that suggest a new kind of intelligence that combines neural inductive learning and symbolic learning of interpretable concepts that can be used for reasoning with applications to a more general artificial intelligence.

3.16 Entity Resolution: A Case for Logic and Learning

Lucian Popa (IBM Almaden Center – San Jose, US)

Joint work of Douglas Burdick, Ronald Fagin, Sairam Gurajada, Jungo Kasai, Phokion Kolaitis, Yunyao Li, Lucian Popa, Kun Qian, Prithviraj Sen, Wang-Chiew Tan

Main reference Kun Qian, Lucian Popa, Prithviraj Sen: "Active Learning for Large-Scale Entity Resolution", in Proc. of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 – 10, 2017, pp. 1379–1388, ACM, 2017.

URL https://doi.org/10.1145/3132847.3132949

Entity resolution is a key form of reasoning over data that allows to establish explicit connections among entities across diverse datasets. In this talk, I will make the case that building good abstractions and tools for entity resolution requires a combination of logic-based methods and machine learning techniques. I will briefly describe a declarative approach that uses constraints and provides a logical foundation for reasoning about entity resolution specifications and their expressive power. This also forms the theoretical underpinning for a concrete high-level language that is used in production by IBM. I will then talk about learning techniques to facilitate the generation of good entity resolution programs using the logic-based language as the target. Of particular importance are active learning techniques where the machine and the human-expert cooperate in order to reach high-accuracy entity resolution algorithms in concrete application scenarios.

3.17 Synthesizing Datalog Programs using Numerical Relaxation

Xujie Si (University of Pennsylvania – Philadelphia, US)

Joint work of Xujie Si, Mukund Raghothaman, Kihong Heo, Mayur Naik
Main reference Xujie Si, Mukund Raghothaman, Kihong Heo, Mayur Naik: "Synthesizing Datalog Programs using Numerical Relaxation", in Proc. of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019, pp. 6117–6124, ijcai.org, 2019.
URL https://doi.org/10.24963/ijcai.2019/847

The problem of learning logical rules from examples arises in diverse fields, including program synthesis, logic programming, and machine learning. Existing approaches either involve solving computationally difficult combinatorial problems or performing parameter estimation in complex statistical models. In this paper, we present DIFFLOG, a technique to extend the logic programming language Datalog to the continuous setting. By attaching real-valued weights to individual rules of a Datalog program, we naturally associate numerical values with individual conclusions of the program. Analogous to the strategy of numerical relaxation in optimization problems, we can now first determine the rule weights which cause the best agreement between the training labels and the induced values of output tuples, and subsequently recover the classical discrete-valued target program from the continuous optimum. We evaluate DIFFLOG on a suite of 34 benchmark problems from recent literature in knowledge discovery, formal verification, and database query-by-example, and demonstrate significant improvements in learning complex programs with recursive rules, invented predicates, and relations of arbitrary arity.

3.18 Information Theory and Data Management

Dan Suciu (University of Washington – Seattle, US)

Joint work of Dan Suciu, Batya Kenig
Main reference Batya Kenig, Dan Suciu: "Integrity Constraints Revisited: From Exact to Approximate Implication", CoRR, Vol. abs/1812.09987, 2018.
URL https://arxiv.org/abs/1812.09987

I will describe three applications of Information Theory to Data Management: upper bounds on query size, approximate constraints, and containment of queries with bag semantics.

3.19 Higher Order Theorem Proving by Deep Learning

Christian Szegedy (Google Inc. – Mountain View, US)

Joint work of Aditya Paliwal, Christian Szegedy, Dennis Lee, Kshitij Bansal, Markus Rabe, Sarah Loos, Stewart Wilcox, Viktor Toman, Kshitij Bansal
Main reference Kshitij Bansal, Sarah M. Loos, Markus N. Rabe, Christian Szegedy, Stewart Wilcox: "HOList: An Environment for Machine Learning of Higher Order Logic Theorem Proving", in Proc. of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, Proceedings of Machine Learning Research, Vol. 97, pp. 454–463, PMLR, 2019.
URL http://proceedings.mlr.press/v97/bansal19a.html

I give an overview of the HOList benchmark and the DeepHOL system for fully automated theorem proving for higher order logic in large theories using a tactic based prover trained by deep reinforcement learning. I will discuss recent results on exploration based strategies for theorem proving. This avoids the necessity of imitation learning on human prooflogs. It is also demonstrates how the choice of suitable deep learning model architecture affects the overall proving performance significantly.

3.20 Machine Learning and Knowledge Graphs

Balder Ten Cate (Google Inc. – Mountain View, US)

In the light of recent developments in (deep) Machine Learning (ML) involving attention mechanisms, and pretraining and finetuning of models, there is a question of whether knowledge bases can add any value, since it has already been shown that these ML models can already learn to answer queries directly from documents. In this talk, I will discuss several advantages that Knowledge Graphs (KG) still offer, such as interoperability, stability over time, and controllability. I will then turn to the question of how ML and KGs can work together, and how ML models can learn to use the data present in a KG. I discuss some approaches to solving these problems and present a high level overview of different possible interfaces between ML and KG.

3.21 Combining Learning and Reasoning over Large Formal Math Corpora

Josef Urban (Czech Technical University – Prague, CZ)

The talk will start with a brief motivation for building strong AI for math and science via combining learning and reasoning over large formal mathematical corpora created with proof assistants such as Mizar, Isabelle, HOL and Coq. I will then describe several tasks in this area such as learning of premise selection over large libraries, learning to guide saturation-style and tableau-style automated theorem provers (ATPs), learning to guide tactical interactive theorem provers, learning of theorem proving strategies, conjecturing, etc. I will also mention various feedback loops between proving and learning in some of these settings, and show some of our autoformalization experiments.

3.22 Statistical Relational Learning

Guy Van den Broeck (UCLA, US)

This talk discusses the role of logical reasoning in statistical machine learning. While their unification has been a long-standing and crucial open problem, automated reasoning and machine learning are still disparate fields within artificial intelligence. I will describe recent progress towards their synthesis in several facets. I start with a very practical question:

how can we enforce logical constraints on the output of deep neural networks to incorporate symbolic knowledge? Second, I explain how circuits developed for tractable logical reasoning can be turned into statistical models. When brought to bear on a variety of machine learning tasks, including discrete density estimation and simple image classification, these probabilistic and logistic circuits yield state-of-the-art results. Finally I give a brief overview of statistical relational learning.

3.23 Towards Finding Longer Proofs

Zsolt Zombori (Alfréd Rényi Institute of Mathematics – Budapest, HU)

Joint work of Zsolt Zombori, Adrián Csiszárik, Henryk Michalewski, Cezary Kaliszyk, Josef Urban
Main reference Zsolt Zombori, Adrián Csiszárik, Henryk Michalewski, Cezary Kaliszyk, Josef Urban: "Towards Finding Longer Proofs", CoRR, Vol. abs/1905.13100, 2019.
URL https://arxiv.org/abs/1905.13100

I present a reinforcement learning (RL) based guidance system for automated theorem proving geared towards Finding Longer Proofs (FLoP). FLoP focuses on generalizing from short proofs to longer ones of similar structure. To achieve that, FLoP uses state-of-the-art RL approaches that were previously not applied in theorem proving. In particular, we show that curriculum learning significantly outperforms previous learning-based proof guidance on a synthetic dataset of increasingly difficult arithmetic problems.

4 Breakout Sessions

4.1 Differentiable FOL Learning from Structures

Adithya Murali (University of Illinois – Urbana-Champaign, US)

Joint work of All participants

Logical structures possess many advantages: they are highly interpretable (and can therefore be inspected or studied in detail), highly compositional, and have many data-efficient learning algorithms. Most importantly, in the context of many AI problems such as analogy reasoning, policy learning or even simple classification, logic offers many excellent modelling choices that can abstract higher-order patterns over primitives. These primitives could correspond to complex non-logical entities such as visual inputs [1] or other signals. While all of these characteristics are well-known and were used in AI systems many years ago, the common criticism is that they are extremely intolerant to noise and, traditionally, offer no way of expressing something like *approximate satisfiability* with respect to a concept. As rightly observed by the authors in [2], they also typically cannot naturally (directly) handle ambiguous non-symbolic data such as raw pixel inputs.

In the last few years, differentiable programming has emerged as a framework for program induction. In this setup, the class of programs is defined by low-level differentiable function families that can be combined by simple higher-level programmatic combinators, and the form of the program is learnt using gradient descent. These programmatic combinators could be higher-order functions like *map* or *fold* [3], but could also be tape head movement or memory updates as in the work on Neural Turing Machines and Differentiable Neural Computers [4, 5, 6]. The learnt concept usually possesses some programmatic structure that can indicate the logic behind the learnt solution, but the crucial inductive generalisations and patterns are learnt from data and are contained in the weights of a neural network. There is therefore a natural interest to apply this philosophy to learning logics, and the breakout session on differentiable FOL (First-Order Logic) learning from structures was centered around this interest. The research question was whether first-order logic formulae could be learnt – as classifiers discriminating between a few first-order structures classified as positive or negative examples – in a *differentiable* manner.

The session began with an introduction to the context of the problem and the desired goals as described above. The first point of discussion was the possibility of being able to embed formulae into a real vector space. The argument against this was that meaningful embeddings (for example, those that mapped semantically equivalent formulae to the same vector) seemed difficult to obtain in a general way. The suggestion was to then obtain these embeddings using a neural network, similar to word embeddings [7, 8] and code embeddings [9]. This was deemed an unsatisfactory solution towards the desired goals as it was unclear how one would obtain a classifier at the end using this method.

Two contrary positions were offered. The first one was that one could embed the structures instead into a (perhaps high-dimensional) space and search for particular classes of formulae that would correspond to tractable structures in that space. However, it was unclear how the learning would be done in a differentiable manner in that setting. The second suggestion received more interest and essentially posited that the semantics of the logic itself could be *lifted* to a continuous or differentiable setting. This would give us the desired 'approximate satisfiability' and could be used to define a metric representing the ability of the formula to discriminate between the given examples, which could then be used to learn the formula that minimises that metric using gradient descent. There is work in this direction [2], including some that were part of talks at the seminar [10, 11]. Some merits and demerits of these works were discussed.

The next question that was addressed was where the data would come from to train such a system, and several suggestions were offered:

- Databases (where the task would be to learn SQL queries).
- Knowledge Graphs. The criticism was that such graphs are large and few, and would therefore not be a source for enough examples. It was not clear whether one could use subgraphs of these graphs to generate more data.
- Randomly generated structures and discriminators. However, it was not clear whether this would generate sufficiently many examples since First-Order Logic follows a Zero-One law and it could happen that either the generated formula would hold on either structure or neither one.
- Datasets like the Visual Genome image dataset [12]. This suggestion was well-received since it would provide enough examples and was an interesting application domain.

The final segment of the breakout session was on suitable problems and ambitious research questions towards differentiably learning FOL formulae. One example provided was that of problems in the verification domain such as invariant synthesis or precondition generation, since the problem statements already require the output to be formulae. Another example was that of generalisation to unseen combinations of properties [13], where a logical representation would naturally be better at expressing such combinations and differentiable learning would remove the need to craft the individual properties by hand. The session concluded with the following question: the methods and works indicated so far do not handle quantifiers; how would one differentiably learn FOL formulae that have quantifiers?

References

1 Adithya Murali, P. Madhusudan, "Augmenting Neural Nets with Symbolic Synthesis: Applications to Few-Shot Learning", CoRR, abs/1907.05878, July 2019.
2 Richard Evans, Edward Grefenstette, "Learning explanatory rules from noisy data", Journal of Artificial Intelligence Research archive, Volume 61 Issue 1, Pages 1-64, January 2018.
3 Lazar Valkov, Dipak Chaudhari, Akash Srivastava, Charles A. Sutton, Swarat Chaudhuri, "Houdini: Lifelong Learning as Program Synthesis", Neural Information Processing Systems (NeurIPS), 2018.
4 Alex Graves, Greg Wayne, Ivo Danihelka, "Neural Turing Machines", CoRR, abs/1410.5401, December 2014.
5 Mark Collier, Joeran Beel, "Implementing Neural Turing Machines", 27th International Conference on Artificial Neural Networks (ICANN), 2018.
6 Alex Graves et al., "Hybrid computing using a neural network with dynamic external memory", Nature, Vol. 538, Pages 471–476, October 2016.
7 Tomas Mikolov, Kai Chen, Greg Corrado, Jeffrey Dean, "Efficient estimation of word representations in vector space", CoRR, abs/1301.3781, 2013.
8 Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, Jeffrey Dean, "Distributed representations of words and phrases and their compositionality", Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013, pages 3111–3119, December 2013.
9 Uri Alon, Meital Zilberstein, Omer Levy, Eran Yahav, "code2vec: Learning Distributed Representations of Code", Proceedings of the ACM on Programming Languages, Volume 3 Issue POPL, January 2019.
10 Xujie Si, Mukund Raghothaman, Kihong Heo, Mayur Naik, "Synthesizing Datalog Programs Using Numerical Relaxation", Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19), Pages 6117-6124, August 2019.
11 Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, Luc De Raedt, "DeepProbLog: Neural Probabilistic Logic Programming", NeurIPS 2018.
12 Ranjay Krishna et al., "Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations", International Journal of Computer Vision, Volume 123 Issue 1, Pages 32-73, May 2017.
13 Junhyuk Oh, Satinder Singh, Honglak Lee, Pushmeet Kohli, "Zero-Shot Task Generalization with Multi-Task Deep Reinforcement Learning", International Conference on Machine Learning (ICML), 2017.

4.2 Explainable AI

Adithya Murali (University of Illinois – Urbana-Champaign, US)

Joint work of All participants

Supervised and semi-supervised deep learning systems such as AlphaGo Zero [1] have surpassed expectations in many fields, including gameplay, self-driving cars, and medical diagnosis. This has prompted the concerns of individuals and governments alike on the decisions made by so-called “black-box” models, or more generally *Explainable AI* (XAI). There have been many attempts through the years [2, 3, 4] to define and characterise in various application domains the concept of an *explanation*, which was the prime focus of the breakout session on XAI.

The primary and perhaps simplistic suggestion was that explanations could be formulae, programs, or constraints – or more generally symbolic objects that are highly interpretable. This of course presents numerous problems, including the fact that such objects, while highly compositional, do not usually offer forms of “approximate satisfaction” with respect to a concept. A different suggestion was to relax the symbolic requirement by instead using objects that are sub-symbolic but still possess some compositional structure. An example of this would be high-level programmatic structures over non-symbolic primitives as in the spirit of works such as [5, 6]. A contrary suggestion posited that explanations could be more complex objects such as dialogue, using even a “black-box” internal representation to continuously clarify and detail the output in the context of questions asked by an agent seeking explanation. However, the argument was made that this would not be an XAI system, but rather an AI system that produces explanations. This distinction indicated the possibility of XAI being an AI-complete problem.

The discussion then turned to some alternative strategies to further the task of exploring XAI. Since attempts to find an all-encompassing definition seemed to fail or were too trivial, one suggestion was to instead attempt a characterisation of explanations by finding a desired set of properties or axioms. Combinations of these axioms would help determine structures that could act as explanations for various applications. One could then use this setup to abstractly study the properties of such “explanation structures” and prove theorems about them. This is in the spirit of works such as [7], which defined a family of fairness measures using axioms.

The last segment of the session was on possible applications or problems where explanations might be necessary. One such example that was suggested was that of chess engines, where one could require not merely gameplay but the extraction of concrete strategies spanning across many moves or explanations for positions that the engines would rate as advantageous for one player. A more interesting suggestion pertained to games such as Angry Birds™ that are played in rounds, where the player would modify their strategy slightly based on the successes or failures of their attempts in earlier rounds. The discussion concluded with a reading of the main points made during the session.

References

1 David Silver et al., “Mastering the game of Go without human knowledge”, Nature, Vol. 550, October 2017.

2 Michael van Lent, William Fisher, Michael Mancuso, “An explainable artificial intelligence system for small-unit tactical behavior”, Proceedings of the National Conference on Artificial Intelligence, pp. 900-907, July 2004.

3 Mustafa Bilgic, Raymond J. Mooney, “Explaining Recommendations: Satisfaction vs. Promotion”, Proceedings of Beyond Personalization 2005, the Workshop on the Next Stage of Recommender Systems Research (IUI2005), January 2005.

4 Nazneen Fatema Rajani, Raymond J. Mooney, “Ensembling Visual Explanations for VQA”, Proceedings of the NeurIPS 2017 workshop on Visually-Grounded Interaction and Language (ViGIL), December 2017.

5 Brenden M. Lake, Ruslan Salakhutdinov, Joshua B. Tenenbaum, “Human-level concept learning through probabilistic program induction”, Science, Vol. 350, Issue 6266, pp. 1332-1338, December 2015.

6 Adithya Murali, P. Madhusudan, "Augmenting Neural Nets with Symbolic Synthesis: Applications to Few-Shot Learning", CoRR, abs/1907.05878, July 2019.
7 Tian Lan, David Kao, Mung Chiang, Ashutosh Sabharwal, "An axiomatic theory of fairness in network resource allocation", Proceedings of the 29th conference on Information communications (INFOCOM'10), pages 1343-1351, March 2010.

4.3 Injecting Symbolic Knowledge/Constraints into Neural Networks

Adithya Murali (University of Illinois – Urbana-Champaign, US)

Joint work of All participants

While connectionist and computationalist methods are often spoken about as being at odds with each other, there have always been efforts to find techniques that reconcile them in some manner. One of the strategies proposed often is to *constrain* the training or output of neural networks using some logical theories or constraints. This was the position that was taken by the participants of the breakout session on injecting symbolic knowledge/constraints into neural networks.

The discussion began with the mention of early works such as Knowledge-Based Artificial Neural Networks (KBANN) [1], which incorporated certain "domain theories" represented in propositional logic into neural networks, building *hybrid learning systems* that used two kinds of information sources: structured knowledge in the form of logic as well as a set of classified examples. The authors of that work showed that these models could learn to generalise to unseen examples better than those using only one kind of information source. Recent works such as the work on Adviceptron [2] or Knowledge-Based Probabilistic Logic Learning [3] have developed on this philosophy by relaxing the hard symbolic constraints to soft constraints, or generalising it to noisy domains.

Then, the desired goals of such an injection of symbolic knowledge into neural networks were discussed. There were at least three clear goals:

1. To speed up or robustify the learning process of neural networks by using symbolic knowledge, similar to the goals of works in [2, 4].
2. To obtain a model at the end of the learning/training process that either incorporates, remembers, or in some way satisfies the given constraints. The choice of defining the appropriate constraints could, however, be domain-specific. For example, the constraints could enforce a particular hypothesis class or representation for the output (say, programs of some bounded measure) as in the work of [5]. They could also be soft constraints that essentially encode an objective function that measures the learnt model's ability to satisfy the real (hard) constraints as in the work of [6].
3. To attempt at building mechanisms of interaction between neural networks and symbolic knowledge modules, perhaps by providing the neural network access to symbolic reasoning engines (similar to the work in [7]).

The rest of the discussion essentially focused on the appropriate language or *medium* to express constraints. There were a few different positions expressed upon this subject. Experts from the Machine Learning community were divided on the position of injecting knowledge using the loss function of the model training phase. Some experts argued that finding the right loss function that approximates (in limit) the desired constraints would be enough.

While this is fairly standard practice, careful analysis of the loss functions using theorems about their properties does not appear to be standard practice. It was also argued that such a methodology was currently still a creativity-based approach rather than a systematic one. There are certainly exceptions as in the work of [6] (which also formed a central theme of one of the seminar talks) where the authors defined a 'semantic loss' that lifts logical constraints with Boolean satisfiability into a continuous form where the constraints could be satisfied fairly well or poorly. However, a crucial argument of those experts against entirely depending on loss functions was that if the data had different patterns or correlations, or in some sense had opposing conclusions to that of the loss function, it would be quite impossible for the training to successfully bridge the two. In general it was possible that the training would jump back forth between the two possibilities and not really converge.

The discussion then turned to a few orthogonal approaches, suggesting that constraints could be richer than simple input/output examples or logically expressed properties. For instance, they could express templates for neural networks as in the work of [8, 9]. Yet another suggestion was that one could, instead of targeting an injunction of knowledge, target extraction of knowledge by coming up with a general formal framework for verifying machine learning models. There is work in this direction [10], but it appeared that most participants disagreed with extraction as against injunction for the purposes of the goals illustrated above.

The session concluded with some thoughts about situations or challenge problems where injecting knowledge would be essential.

References

1 Geoffrey G. Towell, Jude W. Shavlik, "Knowledge-based artificial neural networks", Artificial Intelligence archive, Volume 70 Issue 1-2, Pages 119 – 165, October 1994.
2 G. Kunapuli, K. P. Bennett, R. Maclin, J. W. Shavlik, "The Adviceptron: Giving Advice To The Perceptron", Twentieth Conference on Artificial Neural Networks In Engineering (ANNIE'10), November 2010.
3 Phillip Odom, Tushar Khot, Reid Porter, Sriraam Natarajan, "Knowledge-based probabilistic logic learning", Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15), Pages 3564-3570, January 2015.
4 Jialin Wu, Raymond J. Mooney, "Self-Critical Reasoning for Robust Visual Question Answering", Proceedings of Neural Information Processing Systems (NeurIPS), December 2019.
5 Abhinav Verma, Vijayaraghavan Murali, Rishabh Singh, Pushmeet Kohli, Swarat Chaudhuri, "Programmatically Interpretable Reinforcement Learning", Proceedings of the 35th International Conference on Machine Learning (ICML), 2018.
6 Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, Guy Van den Broeck, "A Semantic Loss Function for Deep Learning with Symbolic Knowledge", Proceedings of the 35th International Conference on Machine Learning (ICML), 2018.
7 Kevin Ellis, Daniel Ritchie, Armando Solar-Lezama, Joshua B. Tenenbaum, "Learning to Infer Graphics Programs from Hand-Drawn Images", Advances in Neural Information Processing Systems 31 (NeurIPS 2018), 2018.
8 Gustav Sourek, Vojtech Aschenbrenner, Filip Zelezny, Ondrej Kuzelka, "Lifted Relational Neural Networks", COCO'15 Proceedings of the 2015th International Conference on Cognitive Computation: Integrating Neural and Symbolic Approaches – Volume 1583, Pages 52-60, December 2015.
9 Gustav Sourek, Vojtech Aschenbrenner, Filip Zelezny, Steven Schockaert, Ondrej Kuzelka, "Lifted relational neural networks: efficient learning of latent relational structures", Journal of Artificial Intelligence Research archive, Volume 62 Issue 1, Pages 69-100, May 2018.
10 Guy Katz et al., "Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks", Computer Aided Verification – 29th International Conference (CAV 2017) July 2017.

Participants

- Isolde Adler
University of Leeds, GB
- Molham Aref
relational AI – Berkeley, US
- Vaishak Belle
University of Edinburgh, GB
- Michael Benedikt
University of Oxford, GB
- Ismail Ilkan Ceylan
University of Oxford, GB
- Victor Dalmau
UPF – Barcelona, ES
- Luc De Raedt
KU Leuven, BE
- Dana Fisman
Ben Gurion University – Beer Sheva, IL
- James Freitag
University of Illinois – Chicago, US
- Ivan Gavran
MPI-SWS – Kaiserslautern, DE
- Martin Grohe
RWTH Aachen, DE
- Barbara Hammer
Universität Bielefeld, DE
- Daniel Huang
University of California – Berkeley, US
- Nils Jansen
Radboud University Nijmegen, NL
- Brendan Juba
Washington University – St. Louis, US
- Kristian Kersting
TU Darmstadt, DE
- Sandra Kiefer
RWTH Aachen, DE
- Angelika Kimmig
Cardiff University, GB
- Phokion G. Kolaitis
University of California – Santa Cruz & IBM Almaden Research Center – San Jose, US
- Egor Kostylev
University of Oxford, GB
- Paul Krogmeier
University of Illinois – Urbana Champaign, US
- Luis C. Lamb
Federal University of Rio Grande do Sul, BR
- Carsten Lutz
Universität Bremen, DE
- Mateusz Malinowski
Google DeepMind – London, GB
- Henryk Michalewski
University of Warsaw, PL
- Adithya Murali
University of Illinois – Urbana-Champaign, US
- Sriraam Natarajan
University of Texas – Dallas, US
- Daniel Neider
MPI-SWS – Kaiserslautern, DE
- Dan Olteanu
University of Oxford, GB
- Ana Ozaki
Free University of Bozen-Bolzano, IT
- Madhusudan Parthasarathy
University of Illinois – Urbana-Champaign, US
- Lucian Popa
IBM Almaden Center – San Jose, US
- Martin Ritzert
RWTH Aachen, DE
- Xujie Si
University of Pennsylvania – Philadelphia, US
- Dan Suciu
University of Washington – Seattle, US
- Christian Szegedy
Google Inc. – Mountain View, US
- Balder Ten Cate
Google Inc. – Mountain View, US
- Josef Urban
Czech Technical University – Prague, CZ
- Steffen van Bergerem
RWTH Aachen, DE
- Guy Van den Broeck
UCLA, US
- Zsolt Zombori
Alfréd Rényi Institute of Mathematics – Budapest, HU