Seminar Künstliche Intelligenz – Sommer-Semester 2020

Themen und Abstracts

Hier eine kurze Übersicht der Themen mit Abstracts.

Thema 1: Learning Plannable Representations with Causal InfoGAN

In recent years, deep generative models have been shown to 'imagine' convincing high-dimensional observations such as images, audio, and even video, learning directly from raw data. In this work, we ask how to imagine goal-directed visual plans – a plausible sequence of observations that transition a dynamical system from its current configuration to a desired goal state, which can later be used as a reference trajectory for control. We focus on systems with high-dimensional observations, such as images, and propose an approach that naturally combines representation learning and planning. Our framework learns a generative model of sequential observations, where the generative process is induced by a transition in a low-dimensional planning model, and an additional noise. By maximizing the mutual information between the generated observations and the transition in the planning model, we obtain a low-dimensional representation that best explains the causal nature of the data. We structure the planning model to be compatible with efficient planning algorithms, and we propose several such models based on either discrete or continuous states. Finally, to generate a visual plan, we project the current and goal observations onto their respective states in the planning model, plan a trajectory, and then use the generative model to transform the trajectory to a sequence of observations. We demonstrate our method on imagining plausible visual plans of rope manipulation.

9 Seiten.

Kurutach, T., Tamar, A., Yang, G., Russell, S. J., and Abbeel, P. (2018). Learning plannable representations with causal infogan. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 8733–8744. Curran Associates, Inc

Thema 2: The price of query rewriting in ontology-based data access

We give a solution to the succinctness problem for the size of first-order rewritings of conjunctive queries in ontology-based data access with ontology languages such as *OWL2QL*,

linear $\texttt{DataLog}^{\pm}$ and sticky $\texttt{DataLog}^{\pm}$. We show that positive existential and nonrecursive datalog rewritings, which do not use extra non-logical symbols (except for intensional predicates in the case of datalog rewritings), suffer an exponential blowup in the worst case, while first-order rewritings can grow superpolynomially unless $\mathsf{NP} \subseteq \mathsf{P/poly}$. We also prove that nonrecursive datalog rewritings are in general exponentially more succinct than positive existential rewritings, while first-order rewritings can be superpolynomially more succinct than positive existential rewritings. On the other hand, we construct polynomial-size positive existential and nonrecursive datalog rewritings under the assumption that any data instance contains two fixed constants.

16 Seiten.

Gottlob, G., Kikot, S., Kontchakov, R., Podolskii, V., Schwentick, T., and Zakharyaschev, M. (2014a). The price of query rewriting in ontology-based data access. *Artificial Intelligence*, 213:42 – 59

Thema 3: Flexible Representative Democracy: An Introduction with Binary Issues

We introduce Flexible Representative Democracy (FRD), a novel hybrid of Representative Democracy (RD) and Direct Democracy (DD), in which voters can alter the issue-dependent weights of a set of elected representatives. In line with the literature on Interactive Democracy, our model allows the voters to actively determine the degree to which the system is direct versus representative. However, unlike Liquid Democracy, FRD uses strictly non-transitive delegations, making delegation cycles impossible, preserving privacy and anonymity, and maintaining a fixed set of accountable elected representatives. We present FRD and analyze it using a computational approach with issues that are independent, binary, and symmetric; we compare the outcomes of various democratic systems using Direct Democracy with majority voting and full participation as an ideal baseline. We find through theoretical and empirical analysis that FRD can yield significant improvements over RD for emulating DD with full participation.

9 Seiten.

Abramowitz, B. and Mattei, N. (2019). Flexible representative democracy: An introduction with binary issues. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 3–10. International Joint Conferences on Artificial Intelligence Organization

Thema 4: Exact Query Reformulation over Databases with First-order and Description Logics Ontologies

We study a general framework for query rewriting in the presence of an arbitrary first-order logic ontology over a database signature. The framework supports deciding the existence of a safe-range first-order equivalent reformulation of a query in terms of the database signature, and if so, it provides an effective approach to construct the reformulation based on interpolation using standard theorem proving techniques (e.g., tableau). Since the reformulation is a safe-range formula, it is effectively executable as an SQL query. At the end, we present a non-trivial application of the framework with ontologies in the very expressive ALCHOIQ description logic, by providing effective means to compute safe-range first-order exact reformulations of queries.

34 Seiten.

Franconi, E., Kerhet, V., and Ngo, N. (2013). Exact query reformulation over databases with first-order and description logics ontologies. *Journal of Artificial Intelligence Research*, 48:885–922

Thema 5: Fusing uncertain knowledge and evidence for maritime situational awareness via Markov Logic Networks

The concepts of event and anomaly are important building blocks for developing a situational picture of the observed environment. We here relate these concepts to the JDL fusion model and demonstrate the power of Markov Logic Networks (MLNs) for encoding uncertain knowledge and compute inferences according to observed evidence. MLNs combine the expressive power of first-order logic and the probabilistic uncertainty management of Markov networks. Within this framework, different types of knowledge (e.g. a priori, contextual) with associated uncertainty can be fused together for situation assessment by expressing unobservable complex events as a logical combination of simpler evidences. We also develop a mechanism to evaluate the level of completion of complex events and show how, along with event probability, it could provide additional useful information to the operator. Examples are demonstrated on two maritime scenarios of rules for event and anomaly detection.

12 Seiten.

Snidaro, L., Visentini, I., and Bryan, K. (2015). Fusing uncertain knowledge and evidence for maritime situational awareness via markov logic networks. *Information Fusion*, 21:159 – 172

Thema 6: BLOG: Probabilistic Models with Unknown Objects

We introduce BLOG, a formal language for defining probability models with unknown objects and identity uncertainty. A BLOG model describes a generative process in which some steps add objects to the world, and others determine attributes and relations on these objects. Subject to certain acyclicity constraints, a BLOG model specifies a unique

probability distribution over first-order model structures that can contain varying and unbounded numbers of objects. Furthermore, inference algorithms exist for a large class of BLOG models.

9 Seiten.

Milch, B., Marthi, B., Russell, S., Sontag, D., Ong, D. L., and Kolobov, A. (2006). Blog: Probabilistic models with unknown objects. In Raedt, L. D., Dietterich, T., Getoor, L., and Muggleton, S. H., editors, *Probabilistic, Logical and Relational Learning - Towards a Synthesis*, number 05051 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany

Thema 7: A Contribution to the Critique of Liquid Democracy

Liquid democracy, which combines features of direct and representative democracy has been proposed as a modern practice for collective decision making. Its advocates support that by allowing voters to delegate their vote to more informed voters can result in better decisions. In an attempt to evaluate the validity of such claims, we study liquid democracy as a means to discover an underlying ground truth. We revisit a recent model by Kahng et al. [2018] and conclude with three negative results, criticizing an important assumption of their modeling, as well as liquid democracy more generally. In particular, we first identify cases where natural local mechanisms are much worse than either direct voting or the other extreme of full delegation to a common dictator. We then show that delegating to less informed voters may considerably increase the chance of discovering the ground truth. Finally, we show that deciding delegations that maximize the probability to find the ground truth is a computationally hard problem

8 Seiten.

Caragiannis, I. and Micha, E. (2019). A contribution to the critique of liquid democracy. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 116–122. International Joint Conferences on Artificial Intelligence Organization

Thema 8: Planning by Rewriting

Domain-independent planning is a hard combinatorial problem. Taking into account plan quality makes the task even more difficult. This article introduces Planning by Rewriting (PbR), a new paradigm for efficient high-quality domain-independent planning. PbR exploits declarative plan-rewriting rules and efficient local search techniques to transform an easy-to-generate, but possibly suboptimal, initial plan into a high-quality plan. In addition to addressing the issues of planning efficiency and plan quality, this framework offers

a new anytime planning algorithm. We have implemented this planner and applied it to several existing domains. The experimental results show that the PbR approach provides significant savings in planning effort while generating high-quality plans.

54 Seiten.

Ambite, J. L. and Knoblock, C. A. (2001). Planning by rewriting. *Journal of Artificial Intelligence Research*, 15:207–261

Thema 9: Query Rewriting and Optimization for Ontological Databases

Ontological queries are evaluated against a knowledge base consisting of an extensional database and an ontology (i.e., a set of logical assertions and constraints that derive new intensional knowledge from the extensional database), rather than directly on the extensional database. The evaluation and optimization of such queries is an intriguing new problem for database research. In this article, we discuss two important aspects of this problem: query rewriting and query optimization. Query rewriting consists of the compilation of an ontological query into an equivalent first-order query against the underlying extensional database. We present a novel query rewriting algorithm for rather general types of ontological constraints that is well suited for practical implementations. In particular, we show how a conjunctive query against a knowledge base, expressed using linear and sticky existential rules, that is, members of the recently introduced *datalog* family of ontology languages, can be compiled into a union of conjunctive queries (UCQ) against the underlying database. Ontological query optimization, in this context, attempts to improve this rewriting process soas to produce possibly small and cost-effective UCQ rewritings for an input query.

45 Seiten.

Gottlob, G., Orsi, G., and Pieris, A. (2014b). Query rewriting and optimization for ontological databases. *ACM Transactions on Database Systems (TODS)*, 39(3):1–46

Thema 10: Query Rewriting for DL-Lite with n-ary Concrete Domains

We investigate ontology-based query answering (OBQA) in a setting where both the ontology and the query can refer to concrete values such as numbers and strings. In contrast to previous work on this topic, the built-in predicates used to compare values are not restricted to being unary. We introduce restrictions on these predicates and on the ontology language that allow us to reduce OBQA to query answering in databases using the so-called combined rewriting approach. Though at first sight our restrictions are different from the ones used in previous work, we show that our results strictly subsume some of the existing first-order rewritability results for unary predicates.

7 Seiten.

Baader, F., Borgwardt, S., and Lippmann, M. (2017). Query rewriting for dl-lite with n-ary concrete domains. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 786–792

Thema 11: First-Order Rewritability of Frontier-Guarded Ontology-Mediated Queries

We focus on ontology-mediated queries (OMQs) based on (frontier-)guarded existential rules and (unions of) conjunctive queries, and we investigate the problem of FO-rewritability, i.e., whether an OMQ can be rewritten as a first-order query. We adopt two different approaches. The first approach employs standard two-way alternating parity tree automata. Although it does not lead to a tight complexity bound, it provides a transparent solution based on widely known tools. The second approach relies on a sophisticated automata model, known as cost automata. This allows us to show that our problem is 2EXPTIME-complete. In both approaches, we provide semantic characterizations of FO-rewritability that are of independent interest.

7 Seiten.

Barceló, P., Berger, G., Lutz, C., and Pieris, A. (2018). First-order rewritability of frontier-guarded ontology-mediated queries. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 1707–1713. International Joint Conferences on Artificial Intelligence Organization

Thema 12: Gradual Semantics Accounting for Similarity between Arguments

Argumentation is a reasoning model based on the justification of claims by arguments. Often, arguments to be considered are not completely independent, two arguments can be related for different reasons, they may overlap, or given by two persons that make similar statements during a debate, but express them differently, etc. This paper studies for the first time the impact of similarity (i.e., when pairs of arguments are related) in the context of gradual evaluation in abstract argumentation. We present principles that a semantics accounting for similarities should satisfy, and show how to extend gradual semantics for this purpose. We propose three original methods to do so, and study their properties. In particular, the new semantics are evaluated with respect to the new principles, and others from the literature.

10 Seiten.

Amgoud, L., Bonzon, E., Delobelle, J., Doder, D., Konieczny, S., and Maudet, N. (2018). Gradual semantics accounting for similarity between arguments. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*

Thema 13: Measuring Similarity between Logical Arguments

Argumentation is a prominent approach for reasoning with (inconsistent) propositional information. It is based on the justification of formulas by arguments, which are minimal and consistent logical proofs of the formulas. The aim of this paper is to evaluate to what extent two such arguments are similar. For that purpose, we introduce a notion of similarity measure and a set of principles that such a measure should satisfy. We propose some intuitive extensions of measures from the literature, and show that they fail to satisfy some of the principles. Then, we come up with a more discriminating measure which satisfies them all.

10 Seiten.

Amgoud, L. and David, V. (2018). Measuring similarity between logical arguments. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*

Thema 14: On the Progression of Situation – Calculus Universal Theories with Constants

The progression of action theories is an important problem in knowledge representation. Progression is second-order definable and known to be first-order definable and effectively computable for restricted classes of theories. Motivated by the fact that universal theories with constants (UTCs) are expressive and natural theories whose satisfiability is decidable, in this paper we provide a thorough study of the progression of situation calculus UTCs. First, we prove that progression of a (possibly infinite) UTC is always first-order definable and results in a UTC. Though first-order definable, we show that the progression of a UTC may be infeasible, that is, it may result in an infinite UTC that is not equivalent to any finite set of first-order sentences. We then show that deciding whether or not there is a feasible progression of a UTC is undecidable. Moreover, we show that deciding whether or not a sentence (in an expressive fragment of first-order logic) is in the progression of a UTC is CONEXPTIME-complete, and that there exists a family of UTCs for which the size of every feasible progression grows exponentially. Finally, we discuss resolution-based approaches to compute the progression of a UTC. This comprehensive analysis contributes to a better understanding of progression in action theories, both in terms of feasibility and difficulty.

10 Seiten.

Arenas, M., Baier, J. A., Navarro, J. S., and Sardina, S. (2018). On the progression of situation calculus universal theories with constants. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*

Thema 15: TensorLog: A Probabilistic Database Implemented Using Deep-Learning Infrastructure

We present an implementation of a probabilistic first-order logic called TensorLog, in which classes of logical queries are compiled into differentiable functions in a neural-network infrastructure such as Tensorflow or Theano. This leads to a close integration of probabilistic logical reasoning with deep-learning infrastructure: in particular, it enables high-performance deep learning frameworks to be used for tuning the parameters of a probabilistic logic. The integration with these frameworks enables use of GPU-based parallel processors for inference and learning, making TensorLog the first highly parallellizable probabilistic logic. Experimental results show that TensorLog scales to problems involving hundreds of thousands of knowledge-base triples and tens of thousands of examples.

41 Seiten.

Cohen, W., Yang, F., and Mazaitis, K. R. (2020). Tensorlog: A probabilistic database implemented using deep-learning infrastructure. ***Journal of Artificial Intelligence Research*****, 67:285–325**

Thema 16: Verifying Procedural Programs via Constrained Rewriting Induction

This article aims to develop a verification method for procedural programs via a transformation into logically constrained term rewriting systems (LCTRSs). To this end, we extend transformation methods based on integer term rewriting systems to handle arbitrary data types, global variables, function calls, and arrays, and to encode safety checks. Then we adapt existing rewriting induction methods to LCTRSs and propose a simple yet effective method to generalize equations. We show that we can automatically verify memory safety and prove correctness of realistic functions. Our approach proves equivalence between two implementations; thus, in contrast to other works, we do not require an explicit specification in a separate specification language.

50 Seiten.

Fuhs, C., Kop, C., and Nishida, N. (2017). Verifying procedural programs via constrained rewriting induction. ***ACM Transactions on Computational Logic (TOCL)*****, 18(2):1–50**

Liste der Titel

Hier nochmal in Kurzform als Liste nur mit den Titeln.

Thema 1:	Learning Plannable Representations with Causal InfoGAN
Thema 2:	The price of query rewriting in ontology-based data access
Thema 3:	Flexible Representative Democracy: An Introduction with Binary Issues
Thema 4:	Exact Query Reformulation over Databases with First-order and Description Logics Ontologies
Thema 5:	Fusing uncertain knowledge and evidence for maritime situational awareness via Markov Logic Networks
Thema 6:	BLOG: Probabilistic Models with Unknown Objects
Thema 7:	A Contribution to the Critique of Liquid Democracy
Thema 8:	Planning by Rewriting
Thema 9:	Query Rewriting and Optimization for Ontological Databases
Thema 10:	Query Rewriting for DL-Lite with n-ary Concrete Domains
Thema 11:	First-Order Rewritability of Frontier-Guarded Ontology-Mediated Queries
Thema 12:	Gradual Semantics Accounting for Similarity between Arguments
Thema 13:	Measuring Similarity between Logical Arguments
Thema 14:	On the Progression of Situation Calculus Universal Theories with Constants
Thema 15:	TensorLog: A Probabilistic Database Implemented Using Deep-Learning Infrastructure
Thema 16:	Verifying Procedural Programs via Constrained Rewriting Induction