

MIT Open Access Articles

Semantic Technologies for Situation Awareness

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

As Published: <https://doi.org/10.1007/s13218-020-00694-3>

Publisher: Springer Berlin Heidelberg

Persistent URL: <https://hdl.handle.net/1721.1/131831>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution





Semantic Technologies for Situation Awareness

Report on a Project Within the CRC “Highly Adaptive Energy-Efficient Computing”

Franz Baader¹ · Stefan Borgwardt¹ · Patrick Koopmann¹ · Veronika Thost² · Anni-Yasmin Turhan¹

Received: 20 September 2020 / Accepted: 15 October 2020
© The Author(s) 2020

Abstract

The project “Semantic Technologies for Situation Awareness” was concerned with detecting certain critical situations from data obtained by observing a complex hard- and software system, in order to trigger actions that allow this system to save energy. The general idea was to formalize situations as ontology-mediated queries, but in order to express the relevant situations, both the employed ontology language and the query language had to be extended. In this paper we sketch the general approach and then concentrate on reporting the formal results obtained for reasoning in these extensions, but do not describe the application that triggered these extensions in detail.

Keywords Situation awareness · Description logics · Reasoning · Complexity · Ontology-based data access

1 Introduction

The project “Semantic Technologies for Situation Awareness” was funded by the German Research Foundation as part of the Collaborative Research Centre “Highly adaptive energy-efficient computing” (HAEC) at TU Dresden [39, 40], which was a joint effort of the faculties of Electrical and

Computer Engineering, Computer Science, and the Department of Mathematics, encompassing more than 20 projects. This project was led by Franz Baader in the first phase of the CRC (2011–2015) and in the second (2015–2019) by him together with Anni-Yasmin Turhan.

The task of our project was to develop formally well-founded methods for achieving context-awareness under incomplete information, which are able to recognize complex situations by means of logical reasoning. The application motivation in this project was to equip a complex hard- and software system with context awareness in order to achieve adaptivity that enables increasing the energy efficiency of the overall system. For example, such a system should be able to adapt its configuration to optimize its energy consumption based on knowledge about user preferences and intentions, quality of service requirements, platforms, location, time, CPU and network load, etc.

The main challenges addressed in this project were thus how to represent context information relevant for enhancing energy efficiency in a formally well-founded way, how to integrate context information obtained from different sources into a coherent semantic view, and how to reason about this view in order to decide on an appropriate action, like moving running processes from an underutilized compute node to another one with remaining capacity, to allow shutting down the former one. To address these challenges, we employed ontologies expressed in appropriate decidable,

The work reported on here was supported by the German Research Foundation (DFG) within the Collaborative Research Center (CRC) 912 – HAEC.

✉ Anni-Yasmin Turhan
Anni-Yasmin.Turhan@tu-dresden.de

Franz Baader
Franz.Baader@tu-dresden.de

Stefan Borgwardt
Stefan.Borgwardt@tu-dresden.de

Patrick Koopmann
Patrick.Koopmann@tu-dresden.de

Veronika Thost
veronika.thost@ibm.com

¹ Institute of Theoretical Computer Science, Technische Universität Dresden, Nöthnitzer Straße 46, 01187 Dresden, Germany

² MIT-IBM Watson AI Lab, IBM Research, Cambridge, MA, USA

logic-based languages. Since such logic-based formalisms are usually interpreted using an open-world semantics, they are well-suited in a setting where the information on the observed system may not be complete.

In order to apply such an ontology-based approach to context recognition [3, 38], one needs to transform the raw data (e.g., numeric sensor information) into a corresponding symbolic context representation expressed in terms of an appropriate application ontology. We adopted *ontology-based data access* (OBDA) [62] as general framework for realizing context recognition since it allows for such a transformation. In OBDA, a database is used together with an ontology, which contains definitions and additional constraints for the concepts relevant in the application domain. First, the raw data, such as subsymbolic sensor information, information on the hardware provided by the operating system, or information from other sources, are preprocessed, aggregated and cleaned (by the database system) and then used to populate the fact base, which contains information expressed in terms of the ontology. Technically this pre-processing is realized by database queries. This framework enables us to derive complex context information from basic information provided by sensors and to reason about this information by abstracting from the actual relations in the database. The logical reasoning is performed over this abstract information using the concept definitions and constraints from the ontology together with the fact base. More precisely, in this ODBA framework situations are represented by conjunctive queries (CQs), which are a well-known class of database queries. However, the data are not assumed to be complete (no closed-world assumption) and the queries may contain (unary and binary) predicates that are defined in an ontology. By using these predicates, the data can be enriched.

The complex contexts recognized this way can then be used to lower the energy consumption of the system by appropriately adapting it to the context. The descriptions of the situations¹ to be recognized are given as queries formulated over the vocabulary of the ontology. These queries are answered over the fact base, using the definitions and constraints contained in the ontology. The existing reasoning methods for OBDA are mostly rewriting-based approaches that transform the initial query such that it incorporates the relevant information from the ontology. After rewriting, the fact base is viewed as a data base, and the rewritten query can be answered over this database using standard database techniques.

Our main extensions to this OBDA-based framework in phase I of the CRC HAEC were to add temporal operators to CQs and to allow the use of fuzzy concepts. Fuzzy concepts are needed to model notions without crisp boundaries (such

as “high processor load”) in an appropriate way. Regarding the temporal extension, information relevant for recognizing critical situations is then stored in time-stamped fact bases. Raw data provided by sensors at a certain point in time are aggregated and abstracted into symbolic information, and stored in the fact base with an appropriate time stamp. When answering temporal queries, not only the time-stamped facts, but also the background knowledge in the ontology is taken into account. This ontology is global (i.e., supposed to hold at every point in time) and expressed in an appropriate Description Logic (DL) using so-called general concept inclusions (GCIs) [11]. The complexity of answering temporal CQs depends on the expressiveness of the DL and on which temporal logic provides us with the temporal operators in the CQs.

In phase I of the CRC we had assumed that information from sensors and other sources is always correct. We dropped this assumption in phase II and thus had to cope with the problem that information stored in the fact base may be incorrect. One possible way of dealing with this problem is to assume that facts are true only with a certain probability. For this reason, we have investigated in phase II how to combine ontological with probabilistic reasoning. Alternatively, one can allow for possibly faulty information in the fact base, and deal with the inconsistencies that this may cause by employing inconsistency-tolerant reasoning and nonmonotonic reasoning.

In the following, we survey the results of phase I in Sect. 2 and the ones obtained during phase II in Sect. 3. We concentrate on describing our own work within the project. For work by others we refer the reader to the descriptions of related work in the cited papers.

2 Extending DL Reasoning for Basic Context Recognition

Our original ideas for how to transform subsymbolic information, such as numerical sensor values, into a logical representation were based on the approach described in [3], where so-called pre-processors were used to clean and aggregate the raw data and populate the fact base. In parallel to our own work in this direction, several other groups considered such a scenario, with the additional restriction that the pre-processors were assumed to be database queries, which transform the raw data into a representation using the concepts of an ontology [26, 62]. Since quite a number of research groups started to investigate this approach, called ontology-based data access [62], we decided to adopt OBDA as the general framework for our situation awareness approach. Since the systems we wanted to monitor showed a dynamic behavior, it was, however, clear from the outset that OBDA had to be extended to the temporal case. Later

¹ We use the notions “context” and “situation” interchangeably.

on, also other research groups adopted temporalized OBDA as a framework for achieving context-awareness [44, 58]. In a keynote address [2] at the Vienna Summer of Logic in 2014 (which was one of the biggest scientific events dedicated to logic to date), F. Baader presented the more general framework of ontology-based monitoring of dynamic systems (OBMDS), which encompasses the framework developed by our project for HAEC, but also has other interesting instances, e.g., ones where the behavior of the system is not only observable at each point in time, but its dynamic behavior (i.e., how a state is transformed into its successor state) is formally specified using an appropriate action description language [14]. Our work on DL-based action description languages was partially funded by DFG within the Research Unit “Hybrid Reasoning for Intelligent Systems” [48].

2.1 Temporalizing OBDA-Based Context Recognition

Our initial studies [30, 42, 43] on situation recognition in the HAEC scenario using standard ontology languages (i.e., the profiles of the web ontology language OWL 2) have underlined that temporal information is fundamental for situation recognition and that the representation of fuzzy information is important as well. Situation descriptions often need to refer to different time points. For instance, in order to express that an application has just terminated, we could say that it was running at the previous time point, but is no longer running now. Thus, instead of having a single fact base that describes the actual state, a snapshot of the relevant system properties is taken several times a minute, and preprocessed into a time-stamped fact base as described above. Overall, this yields a sequence of fact bases. Contexts can then be described using a temporalized query language, which not only considers the actual fact base, but also refers to previous ones.

Since the use of LTL operators within a DL ontology easily renders reasoning undecidable [41, 53], we limited the use of temporal operators to the queries. More precisely, in our temporalized version of OBDA one has a global² ontology written in a certain DL, a sequence of fact bases formulated using the vocabulary of the ontology, and a temporal query, which is an LTL formula in which conjunctive queries can occur in place of the propositional variables of propositional LTL. Based on this overall setting, we have investigated several combinations of DLs and LTL. Basically, one can vary the used DL and decide whether rigid symbols (which do not change their interpretation over time) are allowed or not. As expressive DL, we first considered the DL \mathcal{ALC} and investigated the effect that rigid symbols

have on the computational complexity of answering temporal queries [4]. As usual in research on OBDA, we distinguish data complexity (measured only in the size of the fact bases) from combined complexity (measured in the overall size of fact bases, ontology, and query). Without any rigid symbols, query answering for temporal queries has the same complexity as in the atemporal case for combined complexity. Regarding data complexity, the same holds without rigid symbols and if only concepts can be rigid. If all symbols (i.e., also binary roles) may be rigid, the complexity of the query answering method for the temporal case introduced in [4] is higher than in the atemporal case, but in [63] we could show that this increase can be avoided. The results obtained in [4] for \mathcal{ALC} were extended in [5, 6] to the considerably more expressive DL \mathcal{SHQ} .

As light-weight DL we considered a member of the DL-Lite family of inexpressive DLs, called DL-Lite_{core}, together with a negation-free variant of the LTL-based temporal conjunctive queries described above. Many members of the DL-Lite family (including DL-Lite_{core}) are first-order (FOL) rewritable, which ensures that (atemporal) query answering can be realized via a rewriting of the query into SQL and answering this SQL query over the fact base (without ontology). In [18], we were able to prove that this rewriting approach also works in the temporal case. This paper also addresses the problem that in temporalized OBDA one needs to keep all the fact bases obtained during the run of the system, which results in a huge number of fact bases if the system runs for a long time. For the case where all the temporal queries to be asked are known beforehand, we have devised a method that allows to compile the necessary information about a sequence of fact bases into a single fact base, whose size does not depend on the length of the sequence, but only on the ontology and the predefined queries. In [19] we extended the work on temporal query rewriting in [18] to other DLs by introducing generic approaches for how to transfer rewritability from the atemporal to the temporal case.

For the case of LTL-based temporal conjunctive queries with negation, the complexity of query answering w.r.t. ontologies formulated in \mathcal{EL} or different members of the DL-Lite family was investigated in [22, 23].

2.2 Admitting Vagueness for Ontology-Based Context Recognition

The necessity for vague information in representations of contexts is motivated by our use of symbolic information obtained from numerical sensor values. In fact, concepts like “high CPU load” are not well represented by a crisp representation language where e.g. 90% load is high, but 89.9% load is no longer high. In a fuzzy representation, both values

² in the sense that it is required to hold at all points in time.

would belong to the concept “high CPU load” with slightly different membership degrees.

For the fuzzy representations of contexts, we intended to start from existing results on the decidability and complexity of reasoning in fuzzy Description Logics. However, on closer inspection it turned out that the approach used for showing these results (tableau-based algorithm with a naive blocking mechanism to ensure termination) actually does not work in cases where the fuzzy ontology contains so-called general concept inclusion axioms (GCIs), which are available in all of the OWL 2 profiles that we used for modeling contexts. Thus, instead of being able to employ existing reasoners and results for fuzzy DLs, and just extend them by temporal logics and to query answering, we had to re-investigate the decidability status of fuzzy DLs with GCIs. Our initial investigations within this project produced undecidability results for several fuzzy DLs with GCIs and showed that the decidable cases are actually not truly fuzzy since decidability can be obtained by a reduction to the crisp case [8, 13, 17, 21]. Since within our project in HAEC we did not have enough time for a thorough analysis of the decidability status of fuzzy DLs, these initial results were the starting point for a separate research project funded by DFG and dedicated to investigating fuzzy DLs [9]. This project achieved an almost complete classification of the border between decidability and undecidability for fuzzy DLs with GCIs.

Within our HAEC project we tried to find a practical approach for expressing fuzzy contexts, which trades off expressiveness for decidability. One possible way to proceed was to restrict the fuzzy logic such that it offers only a finite set of fuzzy values. In [20, 54], we combined such a fuzzy logic with very expressive DLs, and showed that answers to conjunctive queries w.r.t. an ontology expressed in the resulting fuzzy DL can still effectively be computed. A second approach developed for query answering in fuzzy DL-Lite performs reasoning by using query answering for the crisp DL as a black box procedure to obtain a rewriting into SQL, where numerical SQL predicates are added in a separate rewriting step to retrieve and compute the fuzzy values [55]. This simple approach yields correct results only for the Gödel semantics, which is a simple but frequently used semantics for fuzzy logics. The approach was implemented [56] and extended to fuzzy temporal query answering [64].

Instead of fuzzy logics, one can also use other means for expressing vague information. One such approach considered in our project are rough logics, which we adapted to the DL setting. In rough DLs, the vagueness information is captured by an indiscernibility relation ρ , which is an equivalence relation over the domain. In [59] we extended the well-known combined approach for query answering in \mathcal{EL} [52] to the rough DL $\mathcal{EL}_{\perp, \rho}$ and showed that the complexity stays the same as for classical \mathcal{EL} .

In the experimental evaluation of our situation recognition framework it turned out that it is useful to return not just answers that exactly match the query (i.e., completely satisfy the description of the situation), but also return answers that “almost” satisfy the query. In fact, situations similar to the ones described are usually also good candidates for an adaptation. This new kind of reasoning problem was defined and investigated in our project for instance queries in (variants of) the DL \mathcal{EL} [36, 37]. The approach uses concept similarity measures (CSMs), which are functions that map pairs of concepts to a similarity value in the interval $[0, 1]$. The new reasoning service retrieves for a given query concept Q , a CSM \sim , and a threshold value $t \in [0, 1)$ all those individuals that are instances of those concepts C that are similar to Q with degree greater than t , i.e., for which $C \sim Q > t$ holds. This form of retrieval supports top- k queries by decreasing thresholds. The approach heavily relies on properties of CSMs, which we investigated in [31, 32, 49] and on the computation of generalizations of individuals, which were studied in [65]. Similarity measures can also be used to extend DLs by constructors that enable the approximate definition of concepts [10].

2.3 Incorporating Concrete Data Values

Besides temporal and vague information, we have considered means to represent concrete data values in our ontology language directly. Concrete domains enable the ontology language to refer to such concrete values, as for instance numbers, and compare them via built-in predicates. They can be used to represent and reason about numerical values, instead of abstracting them away by using preprocessors. Adding concrete domains to DLs can easily make them undecidable in the presence of GCIs. To retain decidability, the concrete domain must satisfy strong restrictions. For almost a decade, the restrictions formulated in [51] were the only ones known to preserve decidability. In [57] we extend these conditions to fuzzy concrete domains. An important result of our project is that we were able to establish new restrictions that guarantee decidability [27], which are orthogonal to the previously known ones, and yield decidability for some interesting concrete domains over the integers that were not covered by the restrictions in [51]. Furthermore, we have investigated CQ answering over ontologies that use concrete domains in [7], which is the foundation for reasoning with probabilistic concrete domains, studied in phase II (see below) and for a practical reasoning procedure for query answering in DL-Lite augmented with a concrete domain over the real numbers [1].

3 Addressing Advanced Challenges for Context Recognition

In phase II of our project, we dropped the simplifying assumption that the data populating the fact base is always free of errors. Thus, the main challenges addressed in phase II were to handle data that is true only with a certain probability and to develop methods that make ontology reasoning resilient against inconsistencies. In cooperation with the HAEC project headed by Christel Baier, we also explored how the design-time probabilistic model checking (PMC) techniques employed in that project can be combined with our run-time monitoring approach.

3.1 Incorporating Probabilities

In practice, we cannot assume that the information in our fact base is always correct, since the raw data may already be erroneous. If information on the reliability of input sources is available (e.g., error probabilities for sensors), then this can be modeled using probabilistic fact bases.

First, we considered a setting where symbolic facts are assigned probabilities, and thus answers to queries are true only with a certain probability. If the probability exceeds a given threshold, then the situation described by the query is recognized. OBDA in this setup had first been investigated in [45] for light-weight DLs, and we extended this work to more expressive DLs in [16]. To support explaining why a situation is recognized, we developed methods for finding most probable explanations for query answers [29].

If sensor values from a continuous domain are directly represented in the fact base, giving a probability to a single value does not make sense as the actual value follows a probability distribution. Thus it is more appropriate to use continuous probability distributions to deal with values obtained from uncertain sensor measurements. This setting was investigated in [12]. Since probabilities may be non-rational real numbers, we employed methods from the complexity theory of real functions, and defined probabilistic query entailment w.r.t. a given precision parameter.

Finally, our framework for probabilistic query answering was integrated with our temporal extension, yielding a new formal framework for specifying temporal probabilistic knowledge bases, and a corresponding query language [46]. Probabilistic temporal query answering had been considered earlier for Datalog [35], but with a query language that does not have temporal and probabilistic operators. Such operators are expedient to describe situations occurring in the HAEC scenario and are available in the query language introduced in [46].

3.2 Resilience Against Inconsistency

If no error probabilities are available, the approaches described in the last subsection cannot be used. Since errors in the raw data may then be propagated to the fact base, it may become inconsistent w.r.t. the ontology. In classical logic, inconsistencies render a fact base useless since everything follows from it. In this case, inconsistency-tolerant reasoning still allows to deduce consequences that are not affected by the error. We have investigated two approaches for achieving such resilience against errors in the data.

First, together with the distinguished female post-doc researcher of the CRC, Camille Bourgaux, we have extended atemporal CQ answering under the so-called repair semantics [15, 50] to the temporal setting. To be more precise, we have investigated the complexity of reasoning in light-weight DLs [24, 25] under different types of repair semantics, also taking rigidity of concepts and roles into account, and were able to show that the complexity often does not increase compared to the case of classical temporal reasoning.

The second approach to deal with erroneous data considered in our project is to use non-monotonic variants of DLs, such as defeasible DLs. In addition to (strict) GCIs these logics use defeasible GCIs that “typically” hold, but which may be defeated by other information. Semantics and reasoning methods for defeasible DLs introduced in previous work mostly ignored defeasible information for anonymous objects introduced by quantifiers. This severe limitation was overcome by our novel reasoning methods for defeasible subsumption and defeasible instance checking in the light-weight defeasible DL \mathcal{EL}_\perp [60, 61].

3.3 Linking Context Recognition and Probabilistic Model Checking

In Christel Baier’s project “Formal Methods for Quantitative Analysis and Optimization of Energy Models”, probabilistic model checking was employed at design-time to verify certain temporal properties of all possible runs of the HAEC system. Our context recognition system observes the HAEC system at run-time. Both approaches use the notion of a critical situation, which is formalized in logic. In fact, also the model checker identifies situations to be avoided or in which an adaptation is required.

Together with members of the group investigating probabilistic model checking, we have developed an approach for integrating situation recognition and PMC: situations and constraints on the system behavior are modeled in our ontology, from which we generate formulas to be used in PMC. This enables the PMC-based analysis to be consistent with our situation descriptions, and allows us to determine off-line, based on the PMC analysis, which situations are most relevant for adaptations towards energy-savings. This

framework, described in [33], was recently investigated in regard to different semantics that can handle inconsistencies between the two components in [34].

This approach also facilitates the use of temporal system properties verified by PMC as additional probabilistic and temporal background knowledge in our ontology. To this end, the language used by PMC (i.e., temporal formulas describing properties of system states) needs to be translated into the ontology and the query language. To some extent this is possible using the query language developed for probabilistic temporal query answering [46] mentioned earlier.

3.4 Other Cooperation

The stimulating environment within the CRC also lead to cooperation with other projects that was not directly motivated by situation recognition. For example, together with the operating systems group of TU Dresden, we conducted an empirical study [47] of the energy consumption of standard ontology reasoners using a standard corpus of ontologies and state-of-the-art DL reasoners. Besides providing a detailed picture of the energy consumption of DL reasoning, our study also explored the relationship between computation power of the CPU, reasoning time, and energy consumption. Together with the group of Markus Krötzsch, we developed a technique for rewriting ontologies formulated in Horn-SRIQ into Datalog programs, which resulted in a much better performance [28].

4 Conclusions

In this report, we have presented the results obtained in the project “Semantic Technologies for Situation Awareness”, which was part of the DFG-funded Collaborative Research Center “Highly Adaptive Energy-efficient Computing” from 2011–2019. Though motivated by the need to save energy in large compute and data servers of the future, our ontology-based framework for situation recognition can, of course, also be employed in other applications. In addition to developing the general framework and prototypical implementations for parts of it, the project has also produced a host of fundamental results on the decidability and complexity of various extensions of DLs, e.g. with fuzzy logic, temporal logic, non-monotonic logic, and probabilistic logic. It has also triggered a growing interest of other groups in temporalized OBDA and its use for complex situation recognition.

Acknowledgements We would like to thank our colleagues Christian Alrabbaa, Camille Bourgaux, Claudia Carapelle, Ismail Ceylan, Felix Distel, Marcel Lippmann, Theofilos Mailis, and Erik Zenker for their

contributions to this project, and all the other members of the CRC HAEC for their cooperation.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alrabbaa C, Koopmann P, Turhan AY (2019) Practical query rewriting for DL-Lite with numerical predicates. In: GCAI 2019. Proceedings of the 5th global conference on artificial intelligence, EPiC series in computing, vol 65. pp 15–27
2. Baader F (2014) Ontology-based monitoring of dynamic systems. In: Proceedings of 14th international conference on principles of knowledge representation and reasoning. Invited contribution
3. Baader F, Bauer A, Baumgartner P, Cregan A, Gabaldon A, Ji K, Lee K, Rajaratnam D, Schwitler R (2009) A novel architecture for situation awareness systems. In: Proceedings of 18th international conference on automation reasoning with analytic tableaux and related methods
4. Baader F, Borgwardt S, Lippmann M (2013) Temporalizing ontology-based data access. In: Proceedings of 24th international conference on automated deduction
5. Baader F, Borgwardt S, Lippmann M (2015) Temporal conjunctive queries in expressive description logics with transitive roles. In: AI 2015: Advances in artificial intelligence, lecture notes in computer science, vol 9457. Springer, pp 21–33
6. Baader F, Borgwardt S, Lippmann M (2015) Temporal query entailment in the description logic *SHQ*. J Web Semant 33:71–93
7. Baader F, Borgwardt S, Lippmann M (2017) Query rewriting for DL-Lite with n -ary concrete domains. In: Proceedings of the international joint conference on artificial intelligence (IJCAI)
8. Baader F, Borgwardt S, Peñaloza R (2014) On the decidability status of fuzzy *ALC* with general concept inclusions. J Philos Log 44:117–146
9. Baader F, Borgwardt S, Peñaloza R (2017) Decidability and complexity of fuzzy description logics. Künstliche Intell 31(1):85–90
10. Baader F, Fernández Gil O (2016) Extending the description logic $\tau EL(deg)$ with acyclic TBoxes. In: Proceedings of the European conference on artificial intelligence (ECAI)
11. Baader F, Horrocks I, Lutz C, Sattler U (2017) An introduction to description logic. Cambridge University Press, Cambridge
12. Baader F, Koopmann P, Turhan, A-Y (2017) Using ontologies to query probabilistic numerical data. In: Proceedings of the international symposium on frontiers of combining systems (FroCoS)
13. Baader F, Peñaloza R (2011) On the undecidability of fuzzy description logics with GCIs and product t -norm. In: Proceedings of international symposium on frontiers of combining systems

14. Baader F, Zarri   B (2013) Verification of golog programs over description logic actions. In: Proceedings of 9th international symposium on frontiers of combining systems
15. Bienvenu M, Bourgaux C, Goasdou   F (2014) Querying inconsistent description logic knowledge bases under preferred repair semantics. In: Proceedings of the AAAI conference on artificial intelligence (AAAI)
16. Borgwardt S, Ceylan II, Lukasiewicz T (2017) Ontology-mediated queries for probabilistic databases. In: Proceedings of the AAAI conference on artificial intelligence (AAAI)
17. Borgwardt S, Distel F, Pe  aloza R (2014) Decidable G  del description logics without the finitely-valued model property. In: Proceedings of 14th international conference on principles of knowledge representation and reasoning
18. Borgwardt S, Lippmann M, Thost V (2013) Temporal query answering in the description logic *DL-Lite*. In: Proceedings of 9th international symposium on frontiers of combining systems
19. Borgwardt S, Lippmann M, Thost V (2015) Temporalizing rewritable query languages over knowledge bases. *J Web Seman* 33:50–70
20. Borgwardt S, Mailis T, Pe  aloza R, Turhan A-Y (2016) Answering fuzzy conjunctive queries over finitely valued fuzzy ontologies. *J Data Seman* 5(2):55–75
21. Borgwardt S, Pe  aloza R (2012) Undecidability of fuzzy description logics. In: Proceedings of 13th international conference on principles of knowledge representation and reasoning
22. Borgwardt S, Thost V (2015) Temporal query answering in *DL-Lite* with negation. In: Proceedings of the global conference on artificial intelligence (GCAI)
23. Borgwardt S, Thost V (2015) Temporal query answering in the description logic *EL*. In: Proceedings of the international joint conference on artificial intelligence (IJCAI)
24. Bourgaux C, Koopmann P, Turhan A-Y (2018) Ontology-mediated query answering over temporal and inconsistent data. *Semant Web J Special Issue on Stream Reasoning*: 1–46
25. Bourgaux C, Turhan A-Y (2017) Temporal query answering in DL-Lite over inconsistent data. In: 16th international semantic web conference (ISWC 2017), LNCS, vol 10587. Springer, pp 121–137
26. Calvanese D, De Giacomo G, Lembo D, Lenzerini M, Poggi A, Rodriguez-Muro M, Rosati R, Ruzzi M, Savo DF (2011) The mastro system for ontology-based data access. *Seman Web* 2:43–53
27. Carapelle C, Turhan A-Y (2016) Description logics reasoning w.r.t. general TBoxes is decidable for concrete domains with the EHD-property. In: Proceedings of the European conference on artificial intelligence (ECAI)
28. Carral D, Gonzalez L, Koopmann P (2019) From Horn-*SRIQ* to Datalog: A data-independent transformation that preserves assertion entailment. In: Proceedings of the AAAI conference on artificial intelligence (AAAI)
29. Ceylan II, Borgwardt S, Lukasiewicz T (2017) Most probable explanations for probabilistic database queries. In: Proceedings of the international joint conference on artificial intelligence (IJCAI)
30. Dargie W, Eldora Mendez J, M  bius C, Rybina K, Thost V, Turhan A-Y (2013) Situation recognition for service management systems using OWL 2 reasoners. In: Proceedings of 10th IEEE workshop on context modeling and reasoning
31. Distel F, Atif J, Bloch I (2014) Concept dissimilarity with triangle inequality. In: Proceedings of 14th international conference on principles of knowledge representation and reasoning
32. Distel F, Jamal A, Bloch I (2014) Concept dissimilarity based on tree edit distances and morphological dilations. In: Proceedings of 21st international conference on artificial intelligence
33. Dubsloff C, Koopmann P, Turhan A-Y (2019) Ontology-mediated probabilistic model checking. In: Proceedings of the 15th international conference on integrated formal methods (iFM'19), LNCS, vol 11918. Springer, pp 194–211
34. Dubsloff C, Koopmann P, Turhan A-Y (2020) Enhancing probabilistic model checking with ontologies. *Form Asp Comput (to appear)*
35. Dylla M, Miliaraki I, Theobald M (2013) A temporal-probabilistic database model for information extraction. In: Proceedings of the VLDB endowment
36. Ecke A, Pe  aloza R, Turhan A-Y (2014) Answering instance queries relaxed by concept similarity. In: Proceedings of 14th international conference on principles of knowledge representation and reasoning
37. Ecke A, Pe  aloza R, Turhan A-Y (2015) Similarity-based relaxed instance queries. *J Appl Log* 13(4):480–508
38. Endsley MR (1995) Toward a theory of situation awareness in dynamic systems. *Hum Factors* 37(1):32–64
39. Fettweis GP, Dorpinghaus M, Castrill  n J, Kumar A, Baier C, Bock K, Ellinger F, Fery A, Fitzek FHP, H  rtig H, Jamshidi K, Kissinger T, Lehner W, Mertig M, Nagel WE, Nguyen GT, Plette-meier D, Schr  ter M, Strufe T (2019) Architecture and advanced electronics pathways toward highly adaptive energy-efficient computing. *Proc IEEE* 107(1):204–231
40. Fettweis GP, Nagel WE, Lehner W (2012) Pathways to servers of the future. In: 2012 design, automation and test in Europe conference and exhibition, DATE 2012. IEEE, pp 1161–1166
41. Gabbay DM, Kurucz A, Wolter F, Zakharyashev M (2003) Many-dimensional modal logics: theory and applications, vol 148. North Holland
42. G  tz S, Mendez J, Thost V, Turhan A-Y (2013) OWL 2 reasoning to detect energy-efficient software variants from context. In: Proceedings of 10th OWL: experiences and directions workshop
43. H  hnel M, Mendez J, Thost V, Turhan A-Y (2014) Bridging the application knowledge gap: using ontology-based situation recognition to support energy-aware resource scheduling. In: Workshop on adaptive and reflective middleware'14
44. Horrocks I, Hubauer T, Jim  nez-Ruiz E, Kharlamov E, Koubarakis M, M  ller R, Bereta K, Neuenstadt C,   z  ep   , Roshchin M, Smeros P, Zheleznyakov D (2013) Addressing streaming and historical data in OBDA systems: Optique's approach. In: Proceedings of 2nd international WS on knowledge discovery and data mining meets linked open data
45. Jung JC, Lutz C (2012) Ontology-based access to probabilistic data with OWL QL. In: The semantic web-ISWC 2012, lecture notes in computer science, vol 7649. Springer, pp 182–197
46. Koopmann P (2019) Ontology-based query answering for probabilistic temporal data. In: Proceedings of the AAAI conference on artificial intelligence (AAAI)
47. Koopmann P, H  hnel M, Turhan A-Y (2017) Energy-efficiency of OWL reasoners: frequency matters. In: Proceedings of the joint international semantic technology conference (JIST)
48. Lakemeyer G (2017) DFG research unit (Forscherguppe) FOR 1513 hybrid reasoning for intelligent systems. *K  nstliche Intell* 31(1):91–92
49. Lehmann K, Turhan A-Y (2012) A framework for semantic-based similarity measures for *ELH*-concepts. In: Proceedings of 13th European conference on logics in artificial intelligence
50. Lembo D, Lenzerini M, Rosati R, Ruzzi M, Savo DF (2010) Inconsistency-tolerant semantics for description logics. In: Web reasoning and rule systems, lecture notes in computer science, vol 6333. Springer, pp 103–117
51. Lutz C, Milicic M (2007) A tableau algorithm for description logics with concrete domains and general TBoxes. *J Autom Reason* 38:227–259
52. Lutz C, Toman D, Wolter F (2009) Conjunctive query answering in the description logic *EL* using a relational database system.

- In: Proceedings of 21st international joint conference on artificial intelligence
53. Lutz C, Wolter F, Zakharyashev M (2008) Temporal description logics: a survey. In: 15th international symposium on temporal representation and reasoning
 54. Mailis T, Peñaloza R, Turhan A-Y (2014) Conjunctive query answering in finitely-valued fuzzy description logics. In: Proceedings of 8th international conference on web reasoning and rule systems
 55. Mailis T, Turhan A-Y (2014) Employing DL-Lite_R-reasoners for fuzzy query answering. In: Proceedings of the joint international semantic technology conference (JIST)
 56. Mailis T, Turhan A-Y, Zenker E (2015) A pragmatic approach to answering CQs over fuzzy DL-Lite-ontologies: introducing FLite. In: Proceedings of the international workshop on description logics (DL)
 57. Merz D, Peñaloza R, Turhan A-Y (2014) Reasoning in ALC with fuzzy concrete domains. In: Proceedings of 37th edition of the German conference on artificial intelligence
 58. Özçep Ö, Möller R, Neuenstadt C (2014) A stream-temporal query language for ontology based data access. In: Proceedings of 27th international workshop on description logics
 59. Peñaloza R, Thost V, Turhan A-Y (2018) Query answering for rough EL ontologies. In: Proceedings of international conference on principles of knowledge representation and reasoning (KR)
 60. PenseL M, Turhan A-Y (2017) Including quantification in defeasible reasoning for the description logic EL_{\perp} . In: Proceedings of the 14th international conference on logic programming and non-monotonic reasoning: LPNMR. Springer, pp 78–84
 61. PenseL M, Turhan A-Y (2018) Reasoning in the defeasible description logic EL_{\perp} : computing standard inferences under rational and relevant semantics. Int J Approx Reason 103:28–70
 62. Poggi A, Lembo D, Calvanese D, De Giacomo G, Lenzerini M, Rosati R (2008) Linking data to ontologies. J Data Semant 10:133–173
 63. Thost V (2017) News on temporal conjunctive queries. In: Joint proceedings of the web stream processing workshop (WSP 2017) and the 2nd international workshop on ontology modularity, contextuality, and evolution (WOMoCoE 2017), CEUR workshop proceedings, vol 1936. pp 1–16
 64. Turhan A-Y, Zenker E (2015) Towards temporal fuzzy query answering on stream-based data. In: Proceedings of the workshop on high-level declarative stream processing (HiDest), vol 1447. CEUR, pp 56–69
 65. Zarriß B, Turhan A-Y (2013) Most specific generalizations w.r.t. general EL-TBoxes. In: Proceedings of 23rd international joint conference on artificial intelligence