Report from Dagstuhl Seminar 22282

Current and Future Challenges in Knowledge Representation and Reasoning

James P. Delgrande*[1], Birte Glimm*[2], Thomas Meyer*[3], Miroslaw Truszczynski*[4], Milene Santos Teixeira†[5], and Frank Wolter*[6]

1 **Simon Fraser University – Burnaby, CA.** jim@cs.sfu.ca
2 **Universität Ulm, DE.** birte.glimm@uni-ulm.de
3 **University of Cape Town, ZA.** tmeyer@cs.uct.ac.za
4 **University of Kentucky – Lexington, US.** mirek@cs.uky.edu
5 **Universität Ulm, DE.** milene.santos-teixeira@uni-ulm.de
6 **University of Liverpool, GB.** wolter@liverpool.ac.uk

Abstract

The area of Knowledge Representation and Reasoning (KR) is a central area in Artificial Intelligence that deals with the explicit, declarative representation of knowledge along with inference procedures for deriving further, implicit information from this knowledge. The goal of this Perspectives Seminar was to assess the area of KR, including its history, current state, and future prospects, and from this assessment to provide suggestions and recommendations for advancing the field, increasing participation in the area, and furthering links with related areas. Over the course of 5 days, 25 participants from a cross-section of subareas in KR and areas adjacent to KR met to discuss these topics. The workshop was composed of a number of invited talks and panels for reviewing the history and state of the art of KR, along with several working groups and general open discussions. In common with other Perspectives Workshops, a Manifesto will be produced; as well, recommendations contained in the manifesto will be also forwarded to the steering committee of the Principles of Knowledge Representation and Reasoning conference series for their consideration.

Seminar July 10–15, 2022 – http://www.dagstuhl.de/22282
2012 ACM Subject Classification Theory of computation → Complexity theory and logic; Computing methodologies → Knowledge representation and reasoning
Keywords and phrases applications of logics, declarative representations, formal logic, knowledge representation and reasoning
Digital Object Identifier 10.4230/DagRep.12.7.62

* Editor / Organizer
† Editorial Assistant / Collector

Current and Future Challenges in Knowledge Representation and Reasoning, *Dagstuhl Reports*, Vol. 12, Issue 7, pp. 62–79
Editors: James P. Delgrande, Birte Glimm, Thomas Meyer, Miroslaw Truszczynski, Milene S. Teixeira, and Frank Wolter

Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Executive Summary

James P. Delgrande (Simon Fraser University – Burnaby, CA)
Birte Glimm (Universität Ulm, DE)
Thomas Meyer (University of Cape Town, ZA)
Mirek Truszczynski (University of Kentucky – Lexington, US)
Milene Santos Teixeira (Universität Ulm, DE)
Frank Wolter (University of Liverpool, GB)

Knowledge Representation and Reasoning (KR) is the field of Artificial Intelligence (AI) that deals with explicit, declarative representations of knowledge along with inference procedures for deriving further, implicit information from these symbolic representations. Research in KR as a mature area of AI is commonly taken as being marked by an Artificial Intelligence Journal Special Issue on Nonmonotonic Reasoning in 1980. In 1989 the Principles of Knowledge Representation and Reasoning Conference was founded, providing a dedicated, specialised forum for research in the area. While KR is one of the oldest and best-established areas of AI, it has continued to grow and thrive over the years. Most of the original research areas have evolved significantly, and have matured from the discovery and exploration of foundations, to the development and analysis of systems for emerging or established applications. Yet other areas, such as argumentation, arose much more recently, and are now thriving areas of KR.

While progress in KR has been steady and often impressive, it has not kept pace with the recent significant successes in AI in the use of statistical techniques and machine learning (ML). As a result, much of the work in AI, and much of the public perception of AI, centres on machine learning and on statistical applications. Nonetheless, we take it as given that KR is a vital, essential area of AI, and that research and development in KR remains necessary. Indeed, despite the unquestionable successes in machine learning and statistical techniques, limitations of these approaches are now emerging that, we believe, can only be overcome with advances in KR. Indicative of this is the recent interest in "Explainable AI", which requires a reference to declarative structures and reasoning over such structures. Furthermore, and in common with the majority opinion in AI, cognitive science, and philosophy, we take it as given that symbolic, declarative representations of knowledge are essential for any ultimate, general theory of intelligence.

For all of these reasons, a reassessment of the area of Knowledge Representation was a very timely undertaking of the Dagstuhl Perspectives Workshop 22282 "Current and Future Challenges in Knowledge Representation and Reasoning". During the seminar, the participants assessed the current state of KR along with future trends and developments. A questionnaire, which had been earlier distributed to the participants, helped in this assessment. Alltogether, the seminar served as a basis for developing an innovative agenda for the next 10-20 years of KR research. Key findings are measures to support a synergistic relationship with other subareas of the rapidly-changing field of AI and of computer science as a whole, e.g, through tutorials at the major KR conference, through new conference tracks and updated reviewing guidelines. The seminar further identified research areas for emphasis, assessed prospects for practical application of techniques, and considered how KR may address limitations of statistical techniques and machine learning.

The program comprised invited talks, panel discussions, working groups, and general dicussions. While the invited talks were agreed upon beforehand, the topics of the working groups (apart from Day 1) were decided interactively with all participants to allow for flexibility and reacting to the talks and the triggered discussions. Day 1 started with a short welcome and participant introduction session, followed by an assessment of *the past and present of KR* in the form of two invited talks by Anthony Cohn and Thomas Eiter. The remainder of Day 1 was dedicated to presenting the questionnaire results, which also prepared for the first working group on rethinking the call for papers (CfP) for the main KR conference, which not only served as rethinking the CfP, but also steered the working groups into thinking about the definition of KR as an area. The day closed with a report from the four working groups and indeed identified changes for the CfP, but also for the track structure and the recruitment and instructions for reviewers.

Day 2 focussed on the relationships of KR with four neighboring areas. For each sub-areas we began with a short invited talk (20 min) followed by a commentary (5 min), also invited, and a short general discussion (5 min). The function of the commentator was to look at the area from a different angle or give another perspective to avoid a too personal or narrow a perspective. The four talks addressed ""KR and AI" (Ian Horrocks, commentator: Sébastien Konieczny), "KR and ML" (Francesca Toni, commentator: Ana Ozaki), "KR and Information Systems" (Diego Calvanese, commentator: Meghyn Bienvenu), and "KR and Robotics" (Gerhard Lakemeyer, commentator: Michael Beetz). Working groups on research challenges for these subareas concluded the day.

The third day began with a short talk on "Handling Uncertainty" (Jean Christoph Jung), for initiating a panel discussion on this topic. The morning concluded with a continuation of the working groups on sub-areas of KR from the previous day. The afternoon was dedicated to hiking and biking in smaller groups.

Day 4 started with short invited talks on "Applications of KR" (Esra Erdem, Thorsten Schaub, Michael Tielscher). The remainder of the day was dedicated to working groups on assessing the state of the art in sub-areas of KR and to expanding KR. For this latter group, we discussed the fact that geographically KR is stronger in Europe than in other parts of the world. As well, we considered how to attract new talent and how to reach out to disadvantaged groups, along with thinking of new forms of events such as hybrid conferences or virtual seminar series.

The final day of the seminar looked at strengthening the interaction between sub-areas of KR and wrapped up with statements of the participants regarding their personal impressions and "take-home" messages. This has, for example, already led to the creation of a novel KR discussion channel (on a Discord server). Key findings include that KR applications are very important to make the field visible and that applications are to be made more visible, e.g., through a journal special issue. Another outcome includes measures to reach out to other areas of AI, in particular machine learning and statistical techniques, where symbolic approaches can make contributions, e.g., for general intelligent agents. A separate Manifesto will provide an assessment of the area, and will give a set of recommendations regarding the future of KR and its promotion.

2 Table of Contents

3 Overview of Talks

3.1 An (Abbreviated, Partial) History of Knowledge Representation: A Personal Perspective

Anthony Cohn (University of Leeds, GB)

In this invited talk I give a brief, and selective history of KR, starting with Aristotle and focusing mostly on work between 1950 and 2000. I talk about Newell's Knowledge Level, the Physical Symbol Hypothesis and early theorem provers, and key advances such as the resolution rule of inference which then led to Prolog and also Datalog. I talk about Frames, and the "scruffy vs neat" debate in the 1970s. I mention the rise of description logics, the semantic web, constraint reasoning and the origins of non monotonic logic. I discuss the representation of particular kinds of knowledge, including taxonomic knowledge and qualitative representation and reasoning methods focusing on spatial knowledge. Finally I note the founding of KR Inc in 1993 and the series of very successful KR conferences it has run since 1989.

3.2 History of Knowledge Representation: A Personal View

Thomas Eiter (TU Wien, AT)

In this talk, I give a brief account of the history of KR, which however is done in a selective manner and in a subjective perspective in which some developments and results are cherry-picked. The presentation will go by decades, covering the time from the beginnings of AI in the 1950s up to the late decade, even if this view is not well-suited as there are overlaps and interesting developments may start or end a bit earlier, and similarly done or start a bit later. Special emphasis is given to the KR conference and the co-development of other venues and communities such as for logic programming, constraint satisfaction, and planning. The talk ends with a status-quo assessment and an outlook on future challenges.

3.3 Knowledge Representation and Artificial Intelligence

Ian Horrocks (University of Oxford, GB)

Main reference Yavor Nenov, Robert Piro, Boris Motik, Ian Horrocks, Zhe Wu, Jay Banerjee: "RDFox: A Highly-Scalable RDF Store", in Proc. of the The Semantic Web – ISWC 2015 – 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part II, Lecture Notes in Computer Science, Vol. 9367, pp. 3–20, Springer, 2015.
URL http://dx.doi.org/10.1007/978-3-319-25010-6_1

KR was central to early work on AI, e.g., at the famous Dartmouth Conference, where several of the well-known protagonists made important contributions to KR. Expert systems, a subsequent high-profile development in AI, is also closely linked to KR. Although there have

been ups and downs in the intervening years, we now have practical KR systems that are used in important applications, particularly in bio-health. An interesting recent (ish) phenomenon is the development of large-scale knowledge graphs (KGs) such as the Google KG. Such KGs are now pervasive in search, e-commerce, and personal assistants (Alexa, Siri, etc.). This success has encouraged the development of general-purpose KG systems and these are now being successfully applied in areas such as industrial design and configuration. So why are KR systems "suddenly" so successful? At least in part due to (i) increasing processing power; (ii) advances in theory and algorithms; (iii) availability of data and digitisation. Of course, many challenges remain, not least knowledge creation and curation – these are still hard problems!

3.4 Knowledge Representation and Machine Learning

Francesca Toni (Imperial College London, GB)

Machine Learning (ML) has grown massively in the last 10 years or so, predominantly due to increased processing power availability, big data, and powerful statistical and probabilistic models. ML is predominantly data-driven these days, with the ambition to automate reasoning as a vector, rather than a manipulator. In this talk, I have expressed the role that KR may have in ML, as well as (to a lower extent) the role that ML may have to KR. I have identified the need for ML to be verified and explained, so as to identify any artifacts and biases that may be present in ML models. KR can (and does) play an important role to support ML. Also, KR can contribute to 'hybrid ML models, integrating reasoning components with statistical/neural ML, as in inductive logic programming. KR-based verification/explanation methods and hybrid systems are all examples of how KR can support ML. On the other hand, ML can help KR with knowledge elicitation (e.g. for argumentation or knowledge graphs). Overall, ML is an important area of AI research and the KR community can gain lots from joining forces with the ML community.

3.5 Knowledge Representation and Information Systems

Diego Calvanese (Free University of Bozen-Bolzano, IT)

URL http://www.inf.unibz.it/ calvanese/presentations/2022-KR-IS-calvanese.pdf

An Information System (IS) is an integration of components for collecting, storing, and processing data, where the data is used to provide information and contribute to knowledge and digital products that facilitate decision making. This definition illustrates that the key elements that an IS has to deal with are data, knowledge, and processes operating over them. These are exactly the elements Knowledge Representation and Reasoning (KRR) has been concerned with, by studying semantic and computational aspects, developing techniques, and building tools. In this presentation we start by providing a brief overview of the main formalisms, techniques, and tools that have been proposed in KRR (notably, based on lightweight Description Logics) to address static and structural aspects of data

and knowledge as they are encountered in ISs. We then move to dynamic and temporal aspects, which are related to the evolution and change over time of data and knowledge due to processes that operate over them. We discuss proposals and results on the integrated modeling and reasoning over data, knowledge, and processes, which is a major concern in ISs. We conclude by presenting some challenges that this poses for KRR.

3.6 Knowledge Representation and Robotics

Gerhard Lakemeyer (RWTH Aachen, DE)

While KR was central to building autonomous robots initially, as demonstrated by the robot Shakey in the late sixties, the field diverged, and KR techniques have started to play a role starting in the late nineties, with good examples being the museum tour guides Rhino and Minerva. One of the most addressed KR frameworks for robots is the KnowRob system developed by Tenorth and Beetz, combining rich ontologies with specialized reasoning. Planning techniques also play a major role, starting with domain-dependent planners like IxTeT and TAL, and later, domain-independent planners like FF and TFD, followed by a combination of action programming languages like Golog and PDDL planners, task and motion planners as well as conditional planners. Task planning needs to be complemented by action execution and maintenance. For the diagnosis of failures model-based techniques with strong KR foundations have been developed. After touching on these issues, my talk ended by proposing the Robocup Logistic League, where a team of robots needs to dynamically assemble products with the help of machines, as a rich benchmark for research in KR and robotics.

3.7 Knowledge Representation and Uncertainty

Jean Christoph Jung (Universität Hildesheim, DE)

Uncertainty arises in many applications: it may come from incomplete information in games like Poker (or other adversarial scenarios), in the presence of random events (like dice rolls), and is a result of extracting information from text, images, audio, video. KR has early on recognized the need of dealing with uncertainty and proposed a range of solutions to various problems. In my talk, I survey the most important concepts that have been identified during the last 40 years (of course the choice is highly subjective): possible world semantics, probabilities, independence, graphical models, updates, and probabilistic and epistemic logic.

3.8 Applications of Knowledge Representation and Reasoning

Esra Erdem (Sabanci University – Istanbul, TR)

As the definition of AI changes towards building rational agents that are provably beneficial for humans, KR&R plays an important role in addressing the user-oriented challenges in applications that come along during this shift, such as generality, flexibility, provability, hybridity, bi-directional interactions, and explainability. In this talk, I present three applications underlining how KR&R addresses these challenges: to solve combinatorial search problems in cladistics, to address knowledge-intensive problems in bioinformatics, and to solve hybrid reasoning problems in robotics. I also discuss how evaluations of human-centric KR&R applications could be extended to include subjective quantitative/qualitative measures.

3.9 Knowledge-driven Artificial Intelligence

Torsten Schaub (Universität Potsdam, DE)

Knowledge plays a vital role in modern companies and organizations, be it in production, storage, or workforce management. Most crucial is the knowledge needed to accomplish creative processes, such as designing a product, an assembly line, a shift schedule or timetable, or planning a trip, routing vehicles, or diagnosing remote systems. All of these tasks involve taking knowledgeable decisions while respecting constraints and preferences. Answer Set Programming (ASP) has become a popular approach for modeling and solving such knowledge-intense combinatorial (optimization) problems. What makes ASP attractive is its combination of a declarative modeling language with highly effective solving engines. This allows us to concentrate on specifying – rather than programming the algorithm for solving – a problem at hand. The talk highlighted some current research topics, and concluded with an outlook on the ASP's potential impact as a knowledge-driven AI tool.

3.10 Application of Knowledge Representation and Reasoning

Michael Thielscher (UNSW – Sydney, AU)

Three systems with a clear KR&R component were presented: (1) A two-armed Blocksworld problem-solving robot with a high-level, symbolic (re-)planning component connected to a low-level robotic controller. (2) An interactive artwork with a BDI-based high-level agent programming component connected to a unity-based controller for virtual characters interacting with human users. (3) A general game-playing system that understands logic-based rule descriptions of new games and uses a propositional logic interface for reasoning connected to a neural network that, with the help of Monte Carlo tree search, learns to play any new game without human intervention.

4 Working groups

4.1 Rethinking the KR Call For Papers

James P. Delgrande (Simon Fraser University – Burnaby, CA), Birte Glimm (Universität Ulm, DE), Thomas Meyer (University of Cape Town, ZA), and Frank Wolter (University of Liverpool, GB)

Participants were divided in four different groups to discuss possible updates of the call for papers of KRR. Each group had 45 minutes to discuss between themselves and their conclusions were shared afterwards. The main points elicited were:

(i) Title of the conference: participants discussed the possibility of updating the conference title (e.g., replacing it by the more general "The international conference on knowledge representation and reasoning"). However, changing the name also has disadvantages since information might be lost (on google search engine, for example). An alternative is to keep the current official name and add the "new title" as a subtitle (as done by other conferences).

(ii) Definition of what KRR is: participants agreed that the main text introducing KRR needs updates. Some updates discussed at the Workshop have already been implemented in the Call for Papers for KR 2023.

(iii) Topics of interest: there are too many topics listed. An alternative is to elicit about 10 topics and add examples of subtopics (e.g. KRR and cognition (e.g. cognitive systems, cognitive reasoning...)). Just hiding all subtopics, however, might not be a good idea since less knowledgeable authors might not understand that their work does not fit the conference.

4.2 Research Challenges: Knowledge Representation and Robotics

Michael Beetz (Universität Bremen, DE), Anthony Cohn (University of Leeds, GB), Esra Erdem (Sabanci University – Istanbul, TR), Andreas Herzig (Paul Sabatier University – Toulouse, FR), Gerhard Lakemeyer (RWTH Aachen, DE), and Michael Thielscher (UNSW – Sydney, AU)

The working group emphasized that there were few submissions to KR and robotics 2022, there being only one paper accepted. The discussion on how it can be improved elicited a few alternatives: (i) besides a special track, introduce a workshop or summer school on the topic; (ii) inclusion of short papers; (iii) linking to a recently created European conference that integrates KR and robotics. Overall, the three main challenges identified by the participants were the following: (1) continuously learning about the world that is changing, i.e. the need for the identification of what KR issues are raised for a robot that works for long periods of time (not just setting a table, but a whole day task, for example); (ii) cognition: learn from demonstration, exchange information and knowledge given by human or available in ontologies; (iii) multi-agent path finding: a complete representation of everything going on robot to understand why they fail is necessary; there is still a lack of good solutions for this problem.

4.3 Research Challenges: Foundations of Knowledge Representation

James P. Delgrande (Simon Fraser University – Burnaby, CA), Marc Denecker (KU Leuven, BE), Sebastien Konieczny (University of Artois/CNRS – Lens, FR), and Thomas Meyer (University of Cape Town, ZA)

The participants of the work group agreed that KR is based on a sound foundational footing, but that many foundational questions remain open. An example that was mentioned is the case of a particular form of non-monotonic reasoning known as "rational closure" (or system Z). The question was raised whether this form of reasoning can really be described as rational. One possibility that was mooted is that of collaboration with psychologists to help with analyzing these theories. The participants also discussed the issue that, although there are many success stories of foundational KR being turned into applications, more ought to be done when it comes to being sufficiently successful with solving real-world problems. In summary, three main challenges were identified: (i) obtaining more precise models of different types of knowledge; (ii) the well-known bottleneck on knowledge acquisition; (iii) differentiating quantitative approaches to reasoning from qualitative approaches.

4.4 Research Challenges: Knowledge Representation and Machine Learning

Birte Glimm (Universität Ulm, DE), Ian Horrocks (University of Oxford, GB), Jean Christoph Jung (Universität Hildesheim, DE), Ana Ozaki (University of Bergen, NO), Steven Schockaert (Cardiff University, GB), and Francesca Toni (Imperial College London, GB)

The working group discussed the challenges regarding the integration of KR and ML. The participants highlighted that KR can contribute to ML in different ways; examples are through the manipulation of rewards or the injection of knowledge to speed up the learning process. However, it was brought to attention that it is still necessary to represent uncertainty on reasoning processes (probabilistic reasoning and conclusions) and insisting on "crisp" knowledge limits collaborations with other fields like ML. Participants agreed that, as a community, we need to be more open to these intersections and even promote them (e.g. special tracks, tutorials, summer schools, promoting a competition within KR&R). Finally, the three main challenges identified by the group were: (1) neuro-symbolic integration: dealing with uncertainty that comes from learning; (2) defining a benchmark/resource for a competition; (3) knowledge compilation: energy consumption, interpretability, identify which rule formalisms are feasible for each model.

4.5 Research Challenges: Knowledge Representation and Information Systems

Magdalena Ortiz (TU Wien, AT), Meghyn Bienvenu (University of Bordeaux, FR), Piero Andrea Bonatti (University of Naples, IT), Diego Calvanese (Free University of Bozen-Bolzano, IT), and Frank Wolter (University of Liverpool, GB)

The working group identified three concrete challenges within KR and information systems: (1) high-level descriptions and abstractions of dynamic data-centric systems and processes, and reasoning about them; (2) support for more comprehensive data management tasks in knowledge-enriched systems, such as analytical queries (numeric values, aggregation), updates, security and privacy, customization, and efficient system design; (3) better collaboration with communities like KGs, the semantic Web, and graph DBs. In particular, it is necessary to be aware and understand the standards of these communities so that we can leverage them to (a) make our techniques readily usable in practice, (b) identify and address their research challenges, and (c) increase the confidence in our solutions in different contexts. Regarding success stories in the field, we can highlight data integration (although several challenges still remain) and ontologies in biomedical domains. Medical ontologies are purpose-specific and hard to reuse, making it challenging their use in practice. One typically needs to extract some "relevant" knowledge, combine ontologies, match terms from different sources, etc., and the tools for such tasks are still underdeveloped. We lack easy-to-use tools that would allow developers to quickly use the existing knowledge. When it comes to processes with data, there are some theoretical results on temporal verification of dynamic data-centric systems, but still far from what is needed in practice. An example is the existence of a model-checker based on SMT, which still plays with toy examples, far from being deployed into real-world business processes. There are also some open challenges regarding privacy and security of information systems, namely (i) confidentiality: keeping the information secure and private (existing approaches are not secure and have many vulnerabilities); (ii) integrity: KBs should be correct (who inserts it? Is it trusted? Did I manage my ontology?); and (iii) availability: the KB should be reliable and not easy to crash.

4.6 Current Trends and Challenges in Knowledge Acquisition

Birte Glimm (Universität Ulm, DE), Steven Schockaert (Cardiff University, GB), and Francesca Toni (Imperial College London, GB)

Obtaining formalized knowledge is seen as a challenge that was discussed in this working group. The participants identified three main challenges regarding knowledge acquisition: (i) learning: it is possible to learn axioms, for example, but it still presents high uncertainty, (ii) maintenance: there are methods to develop ontologies, but not to maintain them, (iii) reasoning on (medical) guidelines, extracting and reasoning on rules is still challenging. Regarding success from last 5 years, the participants highlight: (i) a long running competition, (ii) learned graphs, (iii) availability of huge knowledge bases (fact bases, including certain levels of taxonomy), (iv) wide use of Protege, which offers a plugin architecture with successful built-in reasoners. Protege is not restricted to only the last 5 years, but it is still strongly relevant for the community.

4.7 Current Trends and Challenges in Description Logics

Magdalena Ortiz (TU Wien, AT), Meghyn Bienvenu (University of Bordeaux, FR), Piero Andrea Bonatti (University of Naples, IT), Diego Calvanese (Free University of Bozen-Bolzano, IT), Jean Christoph Jung (Universität Hildesheim, DE), and Frank Wolter (University of Liverpool, GB)

The participants discussed major trends and challenges in description logic research. For instance,

(i) Ontology-based data access has been a major research topic in description logic over the past 15 years. The development of novel and efficient query answering algorithms has been an important achievement. There are now powerful implemented tools for ontology-based data access that are used in various applications. Meaningful benchmarking for query answering remains an important challenge.

(ii) A comprehensive study of the complexity of query answering is another highlight. This includes novel alternative approaches to complexity such as non-uniform and parametrized complexity.

(iii) Making description logics nonmonotonic has been a major research topic with contributions ranging from the development of defeasible description logics to applying circumscription to description logics. Developing tool support remains a major challenge.

(iv) The development of support for ontology engineering has been another major area of research. Principled approaches to explanation, forgetting, repairing, modularity, and versioning have been suggested and investigated. Tools have been developed for some of these approaches.

(v) Explanation of entailments, abduction, and provenance have also been investigated in depth, with significant development of tool support in particular for explanation.

(vi) The close link between description logics and datalog/existential rules has triggered fruitful interaction with the database community; with results and techniques being transferred in both directions.

(vii) Temporal description logics have been investigated in depth over the past 20 years. Here implemented systems are still missing.

(viii) Inconsistency handling in ontology-based data access has been a major research area, again with significant interaction with the database community.

(xii) Adding features for handling uncertain knowledge to description logics has a long history and many approaches have been proposed and investigated. Tool support is still limited.

(xi) Areas that have recently attracted attention but in which challenging problems remain include description logic and learning, description logic and knowledge graphs, security and privacy of description logic ontologies.

4.8 Current Trends and Challenges in Reasoning about Action

Michael Beetz (Universität Bremen, DE), Anthony Cohn (University of Leeds, GB), Andreas Herzig (Paul Sabatier University – Toulouse, FR), Gerhard Lakemeyer (RWTH Aachen, DE), and Michael Thielscher (UNSW – Sydney, AU)

Participants discussed current trends and challenges regarding reasoning about action. Several trends were identified, including:
(i) epistemic planning
(ii) generalized planning
(iii) controller synthesis (e.g. using LTL, LTLf , timed automata, situation calculus)
(iv) expressive action logics with uncertainty, including regression and progression reasoning, and verification of belief programs
(v) ontologies for robots
(vi) causality
(vii) goal reasoning
(viii) learning action representations, and
(ix) KR techniques for informed reinforcement learning.

As well, various challenges were identified, including
(i) solvers for epistemic planners (the problem is undecidable in general)
(ii) learning representations of actions, affordances, and game rules from data
(iii) rational reconstruction of implemented KR systems for robots, like KnowRob, and
(iv) connections between KR work on causality and machine learning.

4.9 Expanding Knowledge Representation: Attracting People

Anthony Cohn (University of Leeds, GB), Thomas Eiter (TU Wien, AT), Gerhard Lakemeyer (RWTH Aachen, DE), and Michael Thielscher (UNSW – Sydney, AU)

Participants discussed alternatives to attract new talents to the community. Launching tangible practical challenges (e.g. angry birds, robotics), as done by other communities, is an alternative that would catch the attention of students. However, it was also highlighted that the process to attract new students must start early, with undergraduate students. It was identified that students might be interested in the field, but they do not have the necessary background. This way, it is important that universities keep elective courses on symbolic reasoning, for example. In alternative, providing high quality online material would support this lack of background. Finally, the participants discussed the creation of very short videos, possibly with testimonials, to motivate new students.

4.10 Expanding Knowledge Representation: Geographics

Esra Erdem (Sabanci University – Istanbul, TR), Ian Horrocks (University of Oxford, GB), Michael Thielscher (UNSW – Sydney, AU), and Frank Wolter (University of Liverpool, GB)

The group identified that the KR community is underrepresented in 3 important regions: North America, China, and southern Asia. Organising workshops addressing relevant KR topics (e.g. knowledge graphs) in North America could contribute to the promotion of the field in this area. Regarding China, participants identified that there are not many submissions from this area as KR is not ranked as a top conference there. The possibility of contacting acquaintances in China to discuss this issue and try an application for a re-ranking was considered. Finally, the participants agreed that to promote KR in southern Asia members of the KR community could offer to give tutorials and guest lectures in the area. Such tutorials and lectures would also be an opportunity to establish new collaborations in KR with researchers in the region.

4.11 Expanding Knowledge Representation: Other Event Types

Andreas Herzig (Paul Sabatier University – Toulouse, FR), James P. Delgrande (Simon Fraser University – Burnaby, CA), Birte Glimm (Universität Ulm, DE), Sebastien Konieczny (University of Artois/CNRS – Lens, FR), Torsten Schaub (Universität Potsdam, DE), and Steven Schockaert (Cardiff University, GB)

The group discussed several kinds of events differing from the standard format of the KR conference. Four recommendations are proposed.

First, we propose to reconsider holding KR every second year as an online event (either fully online or hybrid). Instead of simply copying the on-site format and having a one-week online event, we propose a series of sessions. This is inspired from successful recent experiences with thematic groups (e.g. the Online Social Choice and Welfare Seminar Series, `https://www.sites.google.com/view/2021onlinescwseminars`), projects (e.g. the EU TAILOR project), local groups advertising their online seminars (e.g. the LUCI Lunch Seminar Series, `https://luci.unimi.it/events/`), and national seminars (e.g. the French KR seminar, `https://www.gdria.fr/seminaire/`). Among the pros of this suggestion are: (i) no travel costs; (ii) good for the planet; (iii) it favors participation from developing countries, (iv) it is easier to attract people from adjacent fields such as roboticists (similar to colocation with another conference), and (v) talks can be recorded easily and cheaply. Among the cons, we list: (i) people from the KR field favor of on-site conferences; (ii) choosing the slot(s) will be problematic: there are several solutions but none is optimal, the merits of each of them should be weighed carefully; (iii) it may become apparent that there is too much heterogeneity in our field, e.g. only argumentation people attend argumentation sessions.

Second, we should try to establish benchmark and system competitions in areas where this makes sense. We might take inspiration from the NSF-funded StarExec platform (`https://www.starexec.org/starexec/public/about.jsp`) for first-order SAT solvers or the ICLP Prolog Programming Contests (`https://people.cs.kuleuven.be/~bart.demoen/PrologProgrammingContests/Contest99.html`).

Third, we propose to get more involved in summer schools such as: (i) ESSLLI (https://2022.esslli.eu/), (ii) NASSLLI (https://ml-la.github.io/nasslli2022/), (iii) the EurAI ACAI Summer School (eurai.org/acai), and (iv) the Reasoning Web Summer School (reasoningweb.org). Beyond KR people individually submitting courses, KR Inc. may propose to sponsor 2 or 3 courses.

Finally, KR Inc. might support tutorials and other small events such as schools in underrepresented countries in order to attract locals; local attendance can be expected to go beyond the KR field.

4.12 Expanding Knowledge Representation: Underrepresented Groups

Renata Wassermann (University of Sao Paulo, BR), Meghyn Bienvenu (University of Bordeaux, FR), Diego Calvanese (Free University of Bozen-Bolzano, IT), Jean Christoph Jung (Universität Hildesheim, DE), Thomas Meyer (University of Cape Town, ZA), Magdalena Ortiz (TU Wien, AT), and Ana Ozaki (University of Bergen, NO)

The group started the discussion with the questioning on whether a session to promote the role of women within KR should be held every year. Such a session results in lots of work for very specific participants and might not reach the expected goals. Alternatives for this session are presenting statistics (or a documentary) during the introduction session or displaying short videos during "food sessions". Next, participants also agreed that funding is necessary for inclusion, but it gets expensive for a single student. Therefore, when looking for sponsors for KR, a specific request for D&I funding can be launched. Finally, participants discussed the creation of a mentoring pool, where people from our community that are interested in supporting underrepresented students can act. Related activities include meeting online with students, discussing career prospects, what to work on, and where to publish.

5 Panel discussions

5.1 Current and Future Challenges in Knowledge Representation and Reasoning: Questionnaire Results

Birte Glimm (Universität Ulm, DE)

The panel session started with a presentation of results from a survey with 24 participants to identify key challenges, areas of increasing and decreasing importance in knowledge representation and reasoning and topics to which participants plan to make contributions. As key successes the participants identified existing industry-strength applications, in particular in the areas of answer set programming, description logics and ontologies, and knowledge graphs. In the following discussion it became clear that applications use KR techniques, but this use is often not very visible. In order to make KR applications more visible, a journal special issue on applications was proposed as well as special tracks with instructions

to reviewers as to how applications papers are to be reviewed were identified. A topic that is seen as very important is the area of hybrid AI, i.e., combining knowledge and learning-based methods. While progress in this direction is seen, this is an area that needs significant attention in the future to which at least some of the seminar participants want to contribute. To support this kind of interdisciplinary work, more tutorials and invited talks from the machine learning community can be incorporated into conferences and (summer) schools. Progress in this area was identified as crucial in order for KR to stay relevant as an area, while it was also clear that KR techniques have clear potential to advance intelligent systems, e.g., in terms of explainability or interpretability and when it comes to integrate existing knowledge. Finally, reaching out to neighboring areas is seen as very important as well as to actively showcase the KR successes and work on the still open challenges that so far hinder a wider adoption of KR techniques such as handling uncertainty or inconsistencies.

5.2 Interaction between Subareas

James P. Delgrande (Simon Fraser University – Burnaby, CA), Birte Glimm (Universität Ulm, DE), Thomas Meyer (University of Cape Town, ZA), and Frank Wolter (University of Liverpool, GB)

The final session of the seminar started with a discussion about the possibility of collocating KR with other conferences. For this, it would be interesting to look for papers that bridge gaps between fields (for whomever we collocate with). Participants also discussed the idea of having a session where anyone can "informally" present something in 3 minutes or, instead, researchers could submit an abstract and present it in 5 minutes. These works would be lightly refereed and would not be available in the KR proceedings, since their aim is exposure, e.g. to promote integrations between fields or propose new ideas. At the end of the session, the idea of creating a forum to instigate discussions and works in KR was launched. As a result, a Discord channel to announce talks and activities in the field was created by Ana Osaki and Renata Wassermann. The session concluded with individual feedback from each participant.

Participants

- Michael Beetz
Universität Bremen, DE
- Meghyn Bienvenu
University of Bordeaux, FR
- Piero Andrea Bonatti
University of Naples, IT
- Diego Calvanese
Free University of Bozen-Bolzano, IT
- Anthony Cohn
University of Leeds, GB
- James P. Delgrande
Simon Fraser University – Burnaby, CA
- Marc Denecker
KU Leuven, BE
- Thomas Eiter
TU Wien, AT
- Esra Erdem
Sabanci University – Istanbul, TR
- Birte Glimm
Universität Ulm, DE
- Andreas Herzig
Paul Sabatier University – Toulouse, FR
- Ian Horrocks
University of Oxford, GB
- Jean Christoph Jung
Universität Hildesheim, DE
- Sebastien Konieczny
University of Artois/CNRS – Lens, FR
- Gerhard Lakemeyer
RWTH Aachen, DE
- Thomas Meyer
University of Cape Town, ZA
- Magdalena Ortiz
TU Wien, AT
- Ana Ozaki
University of Bergen, NO
- Milene Santos Teixeira
Universität Ulm, DE
- Torsten Schaub
Universität Potsdam, DE
- Steven Schockaert
Cardiff University, GB
- Michael Thielscher
UNSW – Sydney, AU
- Francesca Toni
Imperial College London, GB
- Renata Wassermann
University of Sao Paulo, BR
- Frank Wolter
University of Liverpool, GB