Manifesto from Dagstuhl Perspectives Workshop 22282

Current and Future Challenges in Knowledge Representation and Reasoning

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

1 **Simon Fraser University, CA.** jim@cs.sfu.ca
2 **Ulm University, DE.** birte.glimm@uni-ulm.de
3 **University of Cape Town, ZA.** tmeyer@cair.org.za
4 **University of Kentucky, US.** mirektruszczynski@gmail.com
5 **University of Liverpool, GB.** wolter@liverpool.ac.uk

Abstract

Knowledge Representation and Reasoning is a central, longstanding, and active area of Artificial Intelligence. Over the years it has evolved significantly; more recently it has been challenged and complemented by research in areas such as machine learning and reasoning under uncertainty. In July 2022,sser a Dagstuhl Perspectives workshop was held on Knowledge Representation and Reasoning. The goal of the workshop was to describe the state of the art in the field, including its relation with other areas, its shortcomings and strengths, together with recommendations for future progress. We developed this *manifesto* based on the presentations, panels, working groups, and discussions that took place at the Dagstuhl Workshop. It is a declaration of our views on Knowledge Representation: its origins, goals, milestones, and current foci; its relation to other disciplines, especially to Artificial Intelligence; and on its challenges, along with key priorities for the next decade.

Perspectives Workshop July 11–15, 2022 – www.dagstuhl.de/22282

2012 ACM Subject Classification Theory of computation → Semantics and reasoning; Theory of computation → Logic; Theory of computation → Complexity theory and logic; Information systems → Information integration; Computing methodologies → Artificial intelligence; Computing methodologies → Knowledge representation and reasoning

Keywords and phrases Knowledge representation and reasoning, Applications of logics, Declarative representations, Formal logic

Digital Object Identifier 10.4230/DagMan.10.1.1

Edited in cooperation with Vaishak Belle, Diego Calvanese, Jens Claßen, Giuseppe de Giacomo, Wolfgang Dvořák, Anthony Hunter, Gerhard Lakemeyer, Marco Montali, Ana Ozaki, Francesco Ricca, Steven Schokaert, Francesca Toni, Johannes Wallner, and Stefan Woltran

Acknowledgements We thank the Dagstuhl staff for their outstanding administrative and operational support in the organisation and running of the workshop. We thank the participants of the workshop, without whose involvement this manifesto, of course, would not be possible. We also thank those individuals who helped us with parts of this manifesto: Vaishak Belle (Reasoning under Uncertainty), Diego Calvanese (Information Systems), Jens Claßen (Reasoning about Action), Giuseppe de Giacomo (Reasoning about Action), Wolfgang Dvořák (Argumentation), Anthony Hunter (Argumentation), Gerhard Lakemeyer (KR and Robotics), Marco Montali (Information Systems), Ana Ozaki (KR and ML), Francesco Ricca (ASP), Steven Schokaert (KR and ML), Francesca Toni (Argumentation), Johannes Wallner (Argumentation), Stefan Woltran (Argumentation).

* Editor / Organizer

Executive Summary

Knowledge Representation and Reasoning (KR) is a field of Artificial Intelligence (AI) that deals with explicit, symbolic, declarative representations of knowledge along with inference procedures for deriving further, implicit information from these representations. Even though KR is one of the oldest and best-established areas of AI, it continues to grow and thrive. 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. Other areas, such as answer set programming and argumentation, arose 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 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 today 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 crucial for the overall development of AI and for its already wide and growing range of applications. 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 are, for instance, the recent interest in *Explainable AI*, which requires a reference to declarative structures and reasoning over such structures, as well as the need for enhancing current AI systems with commonsense reasoning capabilities. Thus, in common with the majority opinion in AI, cognitive science, and philosophy, we espouse the position that 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 is timely. The Dagstuhl Perspectives Workshop 22282 "Current and Future Challenges in Knowledge Representation and Reasoning" had this as its objective. During the workshop, the participants assessed the current state of KR along with future trends and developments. Further, the workshop offered ideas for developing an innovative agenda for the next phase of KR research. Key proposals targeted supporting a synergistic relationship with other subareas of the rapidly-changing field of AI, and of computer science as a whole. The workshop 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.

This manifesto, while representing the views of the authors and several other researchers who offered input and comments, to a large degree is a reflection of discussions at the Dagstuhl workshop. Its first section outlines the field of Knowledge Representation, briefly characterising the area and explaining why a *knowledge-based* approach is important to Artificial Intelligence, and arguably *necessary* for the development of general intelligent agents. The following section reviews areas seen to be in KR, or related to it. Some of these areas fall squarely within KR. Some others have a strong overlap with KR, while for yet others there is a strongly felt need for their overlap with KR to grow. For each of these areas we give a snapshot of their state (as it relates to KR), along with their research issues and challenges. Where appropriate we indicate how KR may contribute to these areas, and how they may in turn contribute to KR. The third section considers research challenges to the field of KR as a whole: what these challenges are and how they may be met. The fourth section discusses ways of promoting KR, including expanding the field and enhancing its visibility. Section 5 provides a brief conclusion. While the history of KR is interesting and instructive, we feel that a detailed history would overly lengthen the manifesto; we have however included a brief history in the appendix.

Table of Contents

1 Introduction to Knowledge Representation and Reasoning

1.1 What is Knowledge Representation and Reasoning?

Knowledge Representation and Reasoning (KR) is a field of Artificial Intelligence (AI) that deals with explicit, symbolic, declarative representations of information along with inference procedures for deriving further, implicit information from these representations. Typically, these symbolic representations encode information on some domain of application. They may also encode other forms of information. For instance, when describing knowledge of an agent, they may encode that agent's goals and preferences, beliefs about other agents, etc. This collection of symbolic expressions is called a *knowledge base (KB)*. The key idea is that the information in a KB describes the domain, without specifying how information is to be used. Accompanying a KB is a *knowledge base manager* consisting of a collection of procedures that perform inferences on the knowledge base, possibly interacting with the external world via sensors and actuators, and possibly modifying the knowledge base. Together, the KB and the KB manager form a *knowledge-based system*. This general paradigm of building systems based on knowledge is also referred to as the *declarative* approach to system building, in contrast with a *procedural* paradigm, in which a system consists of a collection of specific procedures or commands that determine its scope of functionality (as in an operating system, for example).

1.2 Why Knowledge Representation and Reasoning?

The case for declarative representations was made early on in the history of AI, motivated by analogies with how humans address a wide range of problems they encounter in everyday life. A similar case was also made in other areas, including databases, information systems, and logic programming. The arguments put forward in those areas provide a strong support for declarative approaches to knowledge representation. We review some of them below.

First, a symbolic framework allows for a principled, formal account to both representation and reasoning. The meaning of sentences in a KB can be defined via a semantic account, while specialised inference procedures can be developed for the given semantics. Typically, one seeks procedures that are *sound* and *complete*, the characteristics that establish their formal adequacy for reasoning.[1] Other important considerations are their termination and efficiency.

Second, in the declarative paradigm the information in a KB represents *what* an agent knows about a domain but does not constrain *how* to reason with this information: what types of reasoning tasks one might want to perform and how to execute them. For example, in a simple KB system in the form of a relational database, we have a declarative specification of the knowledge given by the set of ground atoms in the language determined by the relational schema. This representation leaves open what reasoning tasks (queries) the user might want to execute, and how they are actually executed once posed to the system.

An important point about this separation of information from reasoning and its implementation is that a KB can easily be updated (new information added, information no longer relevant removed, and erroneous information corrected), without requiring any changes to inference procedures. For example, in a planning domain, information about a particular planning problem can be modified, independently of how a plan is constructed.

[1] A *sound* procedure is one that will derive true consequences from true premises, while a *complete* procedure is one that can derive all true consequences of the premises.

Next, a knowledge-based system is, in principle, able to explain and justify its behaviour, based on its KB and the inference steps it took to arrive at a conclusion. For example a KB may indicate that a specific drug was not prescribed for a patient, even though the drug is effective against the patient's ailment, because the patient is allergic to the class of drugs to which the proposed drug belongs.

Another important consideration concerns the type of information one may want to have in a knowledge-based system. This will clearly include the factual information pertaining to a particular application domain. However, a KB may also contain *any* additional pertinent information such as the agent's (user's) knowledge, goals, obligations, and preferences, or that of other agents. Further, it can include hypothetical or counterfactual information, and assertions regarding the past or future, and also information about actions available to the agent (or agents).

At the same time, KB systems pose numerous challenges (in other words, there is no free lunch). First, both the representation and reasoning problems in a KB may be extremely challenging. A notable example is in commonsense reasoning, where simple assertions such as "birds normally fly" but "penguins normally do not fly" have been the impetus for much research, but where arguably we still do not know how to fully represent and reason with such assertions.

Second, reasoning required by knowledge-based approaches is often inherently computationally hard. This stands in contrast with procedural representations which, when available, are generally more efficient (for example, to find a shortest path, an agent is better off using a specialised shortest-path algorithm rather than inference based on the declarative specification of the concept of a shortest path).

Next, the information that humans have about the world is imperfect. The same is true about the information stored in knowledge-based systems. Often this information is incomplete, imprecise, inconsistent, inaccurate, or otherwise incorrect. Thus, a major research direction is to design inference procedures to handle such challenges. But in many cases, an agreement has not yet emerged on the appropriate choice of such inference mechanisms. And even in cases where there is agreement, reasoning tasks may be undecidable and, when decidable, often are intractable.

Given this, the desiderata of soundness and completeness for reasoning procedures can be difficult or impossible to meet in practice. Hence in the interest of efficient reasoning, a practical inference procedure might be unsound or incomplete in general, or a representation language may be limited in expressibility (for example, in the case of first-order logic languages, to Horn clauses), in return for guaranteed bounds on inference.

A separate issue concerns the matter of acquiring relevant knowledge, and effectively structuring it in a KB. In some cases (for example, description logics) substantial progress has been made, but in others the complexity of the tasks remain a hindrance to the development of practical applications.

The challenge then is to create practical knowledge-based systems for managing complex, real-world tasks. The fundamental problems of KR may be considered in two categories: first, the development of sufficiently precise notations with which to represent knowledge and, second, the development of effective procedures for deriving further knowledge. These have been referred to as, respectively, the *epistemological* and *heuristic* adequacy [265] of a knowledge-based system.

We conclude this section by looking at what is often taken as the ultimate goal of AI, that is, developing a general artificially-intelligent agent. To this end, there is a stance, known as the *KR Hypothesis*, that asserts that a knowledge-based approach is *essential* for the construction of any generally intelligent agent. This was articulated by Brian Smith [339] as follows:

> "Any mechanically embodied intelligent process will be comprised of structural ingredients that
> 1. we as external observers naturally take to represent a propositional account of the knowledge that the overall process exhibits, and
> 2. independent of such external semantical attribution, play a formal but causal and essential role in engendering the behaviour that manifests that knowledge."

That is, the KR Hypothesis claims that any intelligent agent will contain symbolic structures (the knowledge base) where the symbols in these structures *mean* or *represent* something. Moreover, the manipulation of these symbols will play an essential role in determining the behaviour of the system. This *stance* (or *thesis*) directly points to the key role of knowledge representation and reasoning for achieving general intelligence. It is worthwhile to note, though, that some researchers, notably in the connectionist community, disagree. Thus, at present, whether general intelligence is necessarily knowledge-based remains an open question.

2 Areas In and Related to Knowledge Representation and Reasoning

KR is a broad area aiming to establish foundations and general principles for representing knowledge in ways that make that knowledge usable by computers, and to develop effective automated reasoning mechanisms. It can be broken down into several subareas, some of a broad foundational nature, some concerned with specific formalisms and tools, and some focused on applications. As one of the main areas in AI, KR is also related to several important subareas of AI. It also has overlap with other areas in computer science such as database and information systems. Finally, because of its significant use of logic and its focus on the concept of knowledge, KR is also related to mathematical and philosophical logic and, more generally, to philosophy.

In this section we offer an overview of KR and its main research efforts, both past and current. These areas are presented in two rough groups. First are those that fall squarely within KR, such as nonmonotonic reasoning. Second are those areas that are distinct from KR but have strong overlap; these subsections have titles of the form "KR and . . . ".

2.1 Non-monotonic Reasoning

Non-monotonic reasoning (NMR) [64] is one of the original areas of KR and a driving force behind its early development. It remains a major area of KR today. Unlike other areas of KR, such as reasoning about action (Section 2.4), description logics and ontologies (Section 2.5), or argumentation (Section 2.7), NMR is not focused on any particular knowledge representation language or any particular application domain. Rather, it concerns a major aspect of *commonsense reasoning*, that of reasoning (and acting) based on incomplete information by "filling in the gaps" of this missing information in some fashion. In such cases, when additional information becomes available, it is natural that some conclusions that have been reached earlier must be withdrawn. Thus in the standard example (some would say hackneyed) of non-monotonic reasoning, if one learned that a certain animal was a bird, it would be inferred that that animal can fly, and hence for example one would build a cage with a cover on it. If one learned in addition that the animal was a penguin they the default conclusion regarding flying would be withdrawn.

This type of reasoning cannot be modelled in concise, natural ways in classical inference systems such as first-order logic.[2] Classical inference systems are designed to derive new information only when the reasoning that leads to them is iron-clad and cannot be invalidated when new facts become known. More formally, classical reasoning is *monotone*, that is, a conclusion derived from a set of facts remains a conclusion when this set of facts is enlarged. *Defeasibility* of human commonsense reasoning when complete information is not available means that new information may render invalid some of the conclusions obtained in its absence. Therefore, commonsense reasoning requires a new type of logic as its natural formal expression, a logic whose inference mechanism lacks the monotonicity property that classical inference systems have. Several such logics, commonly referred to as *non-monotonic logics*, have been proposed and studied for their capacity to account for ways in which one may draw *non-monotonic* (or defeasible) conclusions from incomplete or imprecise information.

An early non-monotonic mechanism is the *closed world assumption* [320], which expresses the notion that a ground atomic fact can be assumed to be false if it can not be demonstrated to be true. (This, of course, is how negation is handled in standard relational databases.) *Default logic* [321] generalises the intuition behind the closed world assumption. It augmentes classical logic theories (*base* theories) by a set of rules, called *defaults* whose applicability depends on a consistency condition. *Answer set programming* (Section 2.2) in its most basic form can be seen as a fragment of default logic, with its semantics directly traceable to that of default *extensions* [41, 261, 42]. *Circumscription* [267], another major attempt at formalising reasoning with incomplete information, expresses the notion that when faced with such situations, humans make the extension of incompletely specified predicates as small as possible. For example, a circumscriptive interpretation of "birds fly" would have the set of non-flying birds be as small as consistently possible. Yet another line of research of non-monotonic inferences was set in the language of modal logics. Two prominent examples of such logics are the modal non-monotonic logics S4F [334] and KD45 [282]. The former has been shown to be equivalent to default logic. The latter, broadly known as *autoepistemic logic* [283], formalises reasoning of an introspective agent based on what she knew and what she did not know. Over the years, a major research effort was extended to establish fundamental properties of non-monotonic logics and their interrelations [262], and to understand the computational complexity of reasoning with such logics [175, 132].

A different attempt to provide a direction for non-monotonic formalisms is to propose axioms against which any non-monotonic consequence relation could be tested. A prime example of this is the so called *KLM approach* to non-monotonic reasoning [229]. Seeking semantic accounts of non-monotonic inference relations satisfying selected sets of *KLM postulates*, the authors proposed the so called *preferential semantics* and developed fundamental characterisation results. Although that work focused on non-monotonic consequence for classical propositional logic, it was subsequently extended to study non-monotonic consequence over conditional propositional logics [240]. This extension has given rise to a flurry of activity in designing non-monotonic reasoning systems for versions of conditional logics going beyond propositional logic, especially description logics [170, 50, 67] and other fragments of first-order logic. While this work shows some progress, the transfer of results from propositional to first-order logic or its decidable fragments remains a challenge. For example, the question of how to properly account for the additional expressive power associated with quantifiers is not yet resolved [306].

[2] However, by expanding the underlying formal system, for instance to allow time-stamping of each formula or moving to a second-order logic, we can capture some non-monotonic behaviour within classical systems.

Over the years NMR has been taken to include, or apply to, commonsense reasoning, answer-set programming, reasoning about action, query-answering in database systems, description logics, argumentation, and belief change. Indeed following the early development of domain-independent approaches, most work on non-monotonic reasoning has focused on addressing non-monotonic reasoning aspects arising in the areas just mentioned. Current research in NMR may be seen as further mapping out the phenomenon of NMR, investigating its applications and relations to other areas of AI and information systems, and developing further applications.

While research on non-monotonic forms of reasoning has contributed substantially to progress in commonsense reasoning, the key question of how to ensure that the correct inferences are drawn still remains a challenge [119]. At the centre of this challenge is the question of the link between NMR and cognitive science with its two facets: potential influence of NMR research to topics in cognition, and conversely, the extent to which work in cognitive science can benefit NMR, whether by providing benchmarks or phenomena for modelling [340, 317].

Another major area where NMR can make significant contributions is machine learning (Sections 2.9 and 3.3). It is broadly recognised that a major limitation of machine learning in applications is that they often require, but currently lack, common sense and an ability to adjust to unanticipated or novel situations. As noted above, non-monotonic reasoning seems to be a key tool to formalise and implement such abilities. However, this cuts both ways. A successful KR system will require large amounts of real-world information, including defeasible inference rules such as defaults. There is a growing evidence that machine learning may be a primary tool for acquiring these rules.

For further information of current issues in NMR we refer to the materials of a Dagstuhl Perspectives workshop in 2019 dedicated to non-monotonic reasoning [202].

2.2 Answer Set Programming

Answer Set Programming is a declarative formalism for modeling and solving search and optimization problems specified in terms of constraints [65]. The modeling language is that of logic programming, with the meaning of programs determined by the *stable model semantics*.

The stable model semantics that provided the theoretical underpinning to Answer Set Programming (ASP) was introduced in the late 1980s [156]. The view of programs as models of search problems was articulated about ten years later in the papers by Marek and Truszczynski [263] and Niemelä [294]. Supported by the first computationally-promising program-processing tools *lparse* and *smodels* [295, 348], it quickly attracted the interest of researchers dealing with computational solutions to problems in knowledge representation and constraint satisfaction. In the twenty years since its inception, ASP has grown into one of the most vibrant areas of research in KR. A good introduction and an overview of the field can be found in the survey paper by Brewka et al. [65] and in the issue of the AI Magazine dedicated to ASP [66].

ASP continues to attract researchers. On the one hand, the theoretical foundations of ASP abound in challenging problems motivated by its intended use as a tool for modelling knowledge. On the other hand, the demonstrated power of ASP in addressing problems from a broad spectrum of application domains draws those interested in tool development as well as those who are interested in applying the tools. Below we list some of the current research directions that we find particularly interesting or important.

A thorough understanding of why ASP has proved to be so effective still eludes us. Recent work by Hecher, comparing ASP restricted to normal logic programs with propositional logic, provides some insights [188]. That work and several other related papers bring the concept of *tree width* to the forefront of studies of the expressive power, computational complexity, and algorithms for computing answer sets, promising a more nuanced understanding of the inherent computational properties of the stable model semantics. A subsequent paper by Fandinno and Hecher is a good representative of this line of research [139].

Another important line of research concerns extensions of the language of ASP with aggregates. Such extensions are necessary to facilitate both knowledge modelling and program processing. But these extensions invariably bring with them the challenge of defining the "right" semantics. A good overview of this work can be found in the survey paper by Alviano et al. [7]. Another perspective on the matter was proposed by Vanbesien et al. [361]. Finally, extensions of the language by more general classes of constraints by integrating ASP and constraint programming are also being actively researched [249].

An important emerging theme in ASP is to develop methods for counting answer sets of programs. Recent papers by Fichte et al. [142], Kabir et al. [213], and Pajunen and Janhunen [300] are good examples of this line of work. The work is motivated by possible applications of ASP in settings requiring probabilistic reasoning, which is increasingly attracting attention given the current emphasis in AI on probabilistic and neuro-symbolic approaches [238, 377].

The success of ASP stems to a large degree from the availability of excellent computational tools. Most notable here is the suite of tools developed under the *Potassco* project [310]. This project is ongoing with the effort focused on improvements to its flagship programs *gringo*, *clasp* and *clingo*. However, increasingly, the developers address requirements of emerging extensions of ASP: *plingo* for probabilistic reasoning, *aspirin* for reasoning with qualitative preferences, *telingo* for temporal reasoning, and *plasp* for planning. The main challenge of this line of work lies in developing fast programs for model generation, or ASP solvers. Successful research efforts, besides the *Potassco* project, include *wasp* [367] and *dlv* [126].

Computing the models of programs in ASP most often consists of executing two main steps: *grounding*, which creates a propositional program with the same stable models as the original one, and *solving*, where models of the grounded propositional program are computed. In most situations, it is the second step that dominates the computational effort (the problem involved is NP-hard). However, in some applications, the grounding step is the bottleneck – while the problem admits polynomial-time algorithms, the order of the polynomial may be high. Recent efforts to improve grounding tools include the development of grounders that aim to minimise the size of the ground program [77], and lazy grounders [369] that delay grounding of some parts of the program until those parts are needed for solving. Related efforts are aimed at establishing formal foundations of grounding [182, 215].

A crucial aspect of making ASP easier to use and more effective is to develop tools for program development and verification. The tutorial paper by Kaminski et al. [214] offers an overview of the contributions of the *Potassco* project in this area. Other recent notable work on tools for ASP include the development of IDEs [140, 73], debuggers [127] and visualisation software [177]. These and other related projects are discussed in a survey paper by Lierler et al. [250]. A major role in the recent advances in ASP program development and processing tools has been played by the language standard ASP-Core-2 adopted by the community [76].

Thanks primarily to the availability of high quality grounders and solvers, ASP has proved to be an effective technology in addressing problems arising in a wide range of application domains. This continues to be a vibrant area of research. Papers by Erdem et al. [134] and Falkner et al. [138] provide good overviews of applications of ASP in solving problems of practical interest. Recent interesting examples of applications of ASP concern train scheduling [3], course scheduling [27], the stable roommates problem [133], and robotics [133, 325].

2.3 Belief Revision

Belief revision addresses the problem of how an agent should change its beliefs in light of new information. It began as an area of philosophy with the seminal work of Alchourrón, Gärdenfors, and Makinson [6, 307] called, after its founders, the *AGM approach.* There are two facets to this approach, somewhat analogous to the proof theory and semantics of a classical logic. On the one hand, there is a set of belief change *postulates* that arguably any intelligent agent should follow in changing its set of beliefs. On the other hand, there are various formal *constructions* for specifying a belief change operator; these include orderings over formulas in a language, called an *epistemic entrenchment ordering*, or orderings over possible worlds. The belief change postulates and belief operator constructions are related via a *representation theorem*, showing that an operator that satisfies the postulates can be specified via a formal construction, and vice versa. These approaches don't specify a unique operator, but rather they guide the design of specific operators for implementing belief change. The primary operators are belief *revision*, in which an agent consistently incorporates a new belief into its belief corpus, and belief *contraction* in which an agent loses belief in a formula without necessarily believing its negation. Given that belief revision deals broadly with how an agent may change its beliefs in the face of new information, it is not surprising that belief revision attracted attention in the KR community. Today, it is the KR community that is responsible for most research in the area, including a comprehensive exploration of the original AGM framework, and proposed modifications and extensions to other forms of belief change such as belief *merging* [224, 223].

In a related development, researchers have investigated connections between belief revision and non-monotonic reasoning. Firstly, Gärdenfors and Makinson [151] showed that belief revision and the KLM style of non-monotonic reasoning [229] are in a certain formal way so similar that they can be viewed as two sides of the same coin. More recently, Casini et al. [87] proposed a framework for belief revision and belief contraction within a non-monotonic formalism. That the latter even makes sense was somewhat surprising, since one of the basic tenets in belief change used to be that it needs to be based on a monotonic logic.

The original AGM approach had nothing to say about how an agent's *epistemic state* changes following a belief change operation, and it considered the result of a single change operation only. This issue was addressed by Darwiche and Pearl [111] who considered the problem of *iterated* revision by developing a qualitative counterpart to Spohn's [341] *ordinal conditional functions.* In their approach, an agent's epistemic state was represented by a total preorder over possible worlds (with the agent's beliefs characterised by the least set of worlds in the preorder). Then, belief revision could be expressed in terms of the modification of the total preorder. Jin and Thielscher [210], Booth and Meyer [53], as well as Nayak et al. [289] subsequently expanded this line of research. In an orthogonal direction, the AGM approach to revision has also been shown to be applicable in any logic [121], including those weaker than classical propositional logic [54, 122, 120]. As a result, the original AGM approach has been substantially broadened in recent years.

Nonetheless, the area suffers from a number of limitations. Currently, the dominant framework by far is the AGM approach, so much so that on occasion approaches that violate one or another of the AGM principles are seen as being problematic or not *really* about belief change. This is in contrast with non-monotonic reasoning (NMR), where a wide range of approaches, based on differing underlying intuitions or motivation, have been developed. While there are alternative approaches, such as based on distance between models or on syntactic considerations, these approaches have attracted little attention, and it remains an open question whether there are alternative, compelling accounts of belief change to complement the AGM approach.

Perhaps the most significant problem with the area, at least from a computer science perspective, is that research has been almost entirely theoretical in nature. The area largely lacks compelling motivating examples, such as those that drove early work in NMR, and it lacks benchmark problems or domains. While there have been some prototypical implementations, these implementations are usually limited, do not scale well, and again lack a compelling application. Thus, arguably, the major challenge facing the area is to develop useful applications and implementations. Such applications would of course be valuable in their own right. Further, they would provide a driver toward the development of new formalisms, or the elaboration of existing ones. To this end, work on belief change in description logics [276, 239, 385], or in reasoning about action [335, 96] show promise in embedding belief change in potentially practical areas. In a similar vein, it is worth pointing out that the subareas in description logic research known as *axiom pinpointing* [305] and *ontology repair* [24] are closely related to belief change. Roughly speaking, axiom pinpointing is concerned with identifying the statements needed to draw a conclusion, and can be used as the basis for performing belief contraction. Similarly, ontology repair has as its aim the removal, or weakening of description logic axioms in order to regain consistency, and can therefore be viewed as a form of belief revision.

Lastly, belief change can be seen as a qualitative means of dealing with uncertainty and the acquisition of new knowledge; consequently it would be very interesting to consider belief change in the context of current work on uncertainty or in the use of machine learning techniques toward the learning of contingent, qualitative information.

2.4 Reasoning about Action and Planning

Reasoning about action and change is one of the original areas in AI and KR, and it remains an active and growing area of research within KR. Research in the area began with the *situation calculus* [268], a formalism in which actions take one state of the world to another; it was subsequently developed into a full theory by Reiter [322]. Other approaches include the *event calculus* [227], and answer set-based *action languages* [157]. In contrast to reasoning about action, work in *planning* is much more engineering-based. Research in this area goes back to the early AI work using the STRIPS system [144]. Planning in general, and planning to address the needs of autonomous agents, robots remains one of the core areas of AI. Many issues in planning involve representing knowledge about the world, and reasoning about the world as it changes. Nevertheless, today planning and reasoning about action in KR are two largely separate areas of AI. Consequently, a major challenge remains that of (depending on how one looks at things) generalising planners to incorporate the generality of KR theories of reasoning about action or, alternatively, coming up with computational models of reasoning about action that are the same order of efficiency as current planning systems.

Within KR, an interesting trend is the convergence of reasoning about action and the foundations of planning in AI on the one hand, and model checking, games on graphs, and synthesis in formal methods, on the other. In particular, the rise of linear temporal logic on finite traces [168, 169], which is particularly well-behaved computationally and applicable in practice, has shed a new light on the interrelations among logic, automata, and games, well-known in formal methods. It also provided a fertile ground for a new kind of advanced research on reasoning about actions and strategic reasoning for autonomous agents, including planning with temporally extended goals [117], self-programming agents [116], and strategy logic and strategic reasoning [143].

Another trend, as exhibited by a corresponding series of workshops,[3] is research on *generalisation in planning*, where one is interested in general solutions for classes of problems, instead of just solving single instances. In particular, instead of relying on plans in the form of sequences of actions, solutions to generalised planning problems can be programs (policies, controllers) that include control structures such as branches and loops. Two representative examples of this research effort are contributions by De Giacomo et al. [160], who explore controller synthesis in manufacturing scenarios based on first-order representations in the situation calculus, and by Cui et al. [105], who present a framework for generalised planning based on abstraction.

The trend to combine symbolic and machine learning approaches (see Section 3.3) also arises in the area of planning and reasoning about action. One issue is that approaches to planning and reasoning about action rely on representations of action models, which are traditionally crafted by hand. Instead, one wants to be able to learn actions, affordances, and game rules from data or observations. Asai and Muise [15] present an architecture for end-to-end learning of STRIPS representations from images. Geffner [154] argues more generally that symbolic target languages are to be preferred over inductive biases for learning to act and plan.

One form of machine learning that lends itself particularly well to integration with reasoning about action is that of *reinforcement learning*, where the goal is to train an agent to act in an environment based on rewards administered after the execution of each action. Here, symbolic representations can be used for representing the learned policies (for example, in methods based on *policy gradient search*), or for the purpose of encoding reward functions more succinctly, which can help to learn an optimal policy in a more sample-efficient manner. Icarte et al. [205] explore the effectiveness of this approach based on a finite state machine representation, while De Giacomo et al. [161] study “restraining bolts” expressed through temporal logics over finite traces.

Below, we briefly discuss other research directions in reasoning about action and planning.

Goal Reasoning. Whereas classical planning is concerned with finding a sequence of actions to achieve a given, fixed goal, dynamic multi-agent environments require that agents are also able to reason about *which* goals to pursue. Hofmann et al. [195] elaborate on goal reasoning in the context of the Robocup Logistics League (RCLL), where teams of autonomous robots have to coordinate their actions in a simulated factory environment. Roberts et al. [326] present ActorSIM, a general software framework for studying multi-agent goal reasoning in simulated environments. Goal reasoning has also been addressed in the situation calculus [163].

Beliefs. Uncertainty is often connected to some notion of belief. It is helpful to explicitly encode what an agent believes and does not believe, and consider how an agent’s beliefs are affected by its acting and sensing. In this instance, action formalisms are extended

[3] `https://www.genplan.ai/resources/`

by epistemic modalities, one prominent example being Dynamic Epistemic Logic (DEL). For example, Bolander et al. [47] consider the application of DEL for epistemic planning in the context of human-robot collaboration. While these formalisms traditionally rely on qualitative notions of belief and uncertainty, probabilistic variants have been studied as well. For example, Liu and Lakemeyer [254] study belief and meta-belief for a probabilistic action logic based on a modal variant of the situation calculus.

Hybrid Reasoning. While traditionally most approaches make use of qualitative representations, many real-world applications require the ability to handle quantitative information as well. Notable quantitative aspects include time as studied, for example, by Cabalar et al. [74], who present an approach for incorporating metric time into Answer Set Programming for the purpose of planning and scheduling. Another aspect is probabilistic uncertainty, where Zarrieß [379], for instance, presents results on the projection of stochastic actions in a probabilistic description logic.

Infinite-State Systems and the Situation Calculus. While most research in planning is based on finite-state systems, important research has also been conducted on infinite-state systems, most notably, at least in recent years, in the Situation Calculus [322]. This includes decidable reasoning about actions through a bounded (FOL) state assumption [164], a non-Markovian version of the situation calculus [165], research on abstraction in the situation calculus [28], and controller/program synthesis in the situation calculus [160].

Ethics and Norms. Where embodied agents (such as mobile robots or autonomous cars) act in physical environments that they share with human beings, it becomes increasingly important that their actions are governed by social norms and moral guidelines. The subarea of *normative (multi-)agent systems* [93] is concerned with the study of such agents. Many approaches are based on integrating reasoning about actions with some form of *deontic logic* to represent notions such as obligation, permission, and prohibition. For instance, Horty [199] studies obligations in the context of an epistemic variant of *stit* ("see-to-it-that") semantics. Further, Lindner et al. [253] present an approach for incorporating ethical reasoning into a planning system, distinguishing act-based from goal-based deontological principles.

2.5 Description Logics and Ontologies

Description logics (DLs) [20, 22] were first introduced in the 1980s as an attempt to create a formal semantics for frames and semantic networks. Initially, DLs had limited expressive power and were designed for polynomial time reasoning. However, today they have evolved into a large family of languages that typically permit unary and binary relation symbols and some form of universal or existential quantification. While some DLs also permit second-order features like transitive closure or more general fixpoints, most DLs are fragments of first-order logic and inherit its classical open world semantics. The research community generally agrees that DLs should be decidable in order to enable effective and robust reasoning support.

The DL community meets annually at the Description Logic Workshop,[4] which had its 35^{th} jubilee in 2022. It closely interacts with many other research communities, including the ASP and general non-monotonic reasoning community in efforts to combine the open world semantics of DLs with closed world features required in many applications. As well, it

[4] https://dl.kr.org/

interacts with the datalog and rule-based reasoning community to develop ontology-mediated data access. Because of its focus on decidable reasoning, the DL community also contributes to ongoing efforts of mapping out the boundary between decidable and undecidable fragments of first-order logic.

Currently, the main view of DLs is as an underpinning language for ontologies, regarded in this document as logical theories that define the relevant classes, attributes, and relations for a domain of interest and that specify the relationships between them by means of logical axioms. Indeed, ontologies can be regarded as knowledge bases with a focus on terminological, schema, or conceptual knowledge. The design of ontology languages together with efficient reasoning engines has been a major research direction in DL. This research has led to the World Wide Web Consortium (W3C) OWL ontology language [275] and its successor OWL 2 [364] specifying three sublanguages, called *profiles*, of OWL [284]: OWL 2 EL, a fragment with polynomial time reasoning complexity used for large scale ontologies and based on the description logic $\mathcal{EL}$ [19], OWL 2 QL, designed to support efficient access to data using ontologies and based on the description logic DL-Lite [78], and the rule-based language OWL 2 RL.

A basic research challenge for any ontology language is to understand the trade-off between expressivity and complexity of reasoning, and to develop, implement, and analyse efficient reasoning systems. In what follows, we distinguish between two rather different types of reasoning: *terminological reasoning* that aims at extracting knowledge from an ontology by computing, for example, the induced concept hierarchy, and *ontology-mediated querying of data* that supports access to data modulo an ontology. In the former case, ontologies are typically very large and often it is not the full logical theory underpinning the ontology that is needed in applications, but only the induced concept hierarchy or variants thereof. In contrast, in the latter case, ontologies are typically much smaller and their main purpose is to provide a schema for accessing, managing, and interpreting various types of data.

Regarding terminological reasoning, significant progress has been made over the past 15 years. There are now extremely powerful reasoning engines that compute the concept hierarchy induced by very large ontologies formulated in OWL 2 EL (for example, ELK [217]), in Horn extensions of OWL 2 EL, and even in very expressive description logics (for example, HermiT [171], Pellet [338], or Konclude [342]). A typical application of this type of reasoning is modelling support for SNOMED CT.[5] Reasoners for OWL 2 are, by now, so highly optimised that it is very challenging to implement novel optimisations into an existing reasoner. Equally challenging is the development of novel competitive reasoners based on new approaches. One possible direction is to investigate a reasoner that translates to SAT, so that it can be computed by a state of the art SAT solver.

It remains a major research challenge to provide principled logic-based support for the development and maintenance of large-scale ontologies. Active research areas include the computation and presentation of *explanations* of subsumptions (or, even more challenging, non-subsumptions) between classes [197, 172], automated support for *modularisation and module extraction* [344], *versioning* and collaborative development of ontologies [209, 57], and *schema manipulation* such as forgetting [222, 384]. All mainstream ontology languages in this area are fragments of first-order logic. Extensions of these languages to deal with defeasible knowledge and exceptions [281, 51], time [258], and uncertainty [257] remain important research questions.

[5] SNOMED CT (`https://www.snomed.org/`) is a comprehensive, multilingual clinical healthcare terminology with more that 350,000 concepts and used in more than 80 countries. It enables the consistent representation of clinical content in electronic health records.

If ontologies are to be used to access data, the main reasoning problems often have a rather different flavour. In this case, reasoning becomes an extension of query answering over data, a problem originally addressed in database research. As data stored on the web and in data warehouses is often heterogeneous, distributed, and only partially structured, it is often incomplete and even logically inconsistent. Dealing with such data requires complex and expressive data representation models and dedicated reasoning services to support data access, management, and interpretation. Hence the distinction between data and knowledge bases has become blurred and the insights gained in KR research for dealing with issues such as incomplete information and logical inconsistency have become directly relevant for mainstream database applications.

A basic application of ontology-mediated query answering is to enrich an incomplete data source with background knowledge in order to obtain a more complete set of answers to a query. Another application is *data integration*, where an ontology is used to provide a uniform view on multiple data sources abstracting from specialised schemata and implementation details. Starting with DL-Lite [78], great progress has been made in understanding the complexity of answering queries under DL ontologies [79, 45, 44] and under existential rules [75]. Various types of systems such as Ontop and Vadalog have been developed that support ontology-mediated query answering [373, 37, 374].

Many research challenges related to ontology-mediated data access remain to be solved, however. Active research areas include dealing with inconsistency in ontology-mediated querying-answering [43], temporal ontology-based data access [14, 365], and reverse engineering of queries from data examples [12, 212]. Moreover, research into crucial issues such as privacy and data provenance under ontologies has only just begun.

2.6 Knowledge Graphs

In 2012 Google introduced the term *knowledge graph* to refer to a large collection of facts accompanied by a basic ontology providing semantics. Since then, the term knowledge graph has been used to denote a wide variety of data models, which have in common that information is stored in a graph-like fashion, making use of labelled nodes and edges to represent objects and entities, facts about them, and relationships between them (see the paper by Hogan et al. [196] for a discussion of the basic concepts). These data models have become an important paradigm in both industry and research, with applications, for instance, in (web) search, as a backbone of Wikipedia through the structured knowledge in Wikidata [363], and in personal assistants such as Alexa or Siri. Knowledge graphs are often of impressive size. For example, Google's knowledge graph started with public sources such as Freebase, Wikipedia and the CIA World Factbook, but was quickly extended to contain more than 500 million nodes and 3.5 billion facts and relationships. In 2020, Google reported a size of 500 billion facts about five billion entities.[6] Microsoft, in its Bing search engine, also uses a knowledge graph.[7] In 2015, that graph contained over one billion entities with more than 21 billion facts associated to them and over 5 billion relationships between them. Apart from numerous commercial knowledge graphs, Wikidata is an example of a large publicly maintained knowledge graph.

[6] https://blog.google/products/search/about-knowledge-graph-and-knowledge-panels/

[7] https://blogs.bing.com/search/2015/08/20/bing-announces-availability-of-the-knowledge-and-action-graph-api-for-developers/

Knowledge graphs give rise to difficult but exciting research challenges, both for the database and the knowledge representation community. From a database perspective, the design of new query languages that can navigate graphs in ways that go beyond first-order queries is ongoing [9]. From a knowledge representation perspective, a principled approach to adding semantics to knowledge graphs without affecting the efficiency of query answering is an important but very difficult problem. For instance, to enable the addition of logical rules to knowledge graphs, highly optimised rule systems such as RDFox [291] and VLog [85] have been developed. Further open research problems in this direction include dealing with meta annotations and with restrictions on relations. For example, in many knowledge graphs, relations are time-dependent; they are true for a certain amount of time, but not before or after. Modelling this in standard description logics requires expensive workarounds such as reification [183]. First steps towards new modelling languages dealing with such issues are *attributed description logics* presented by Krötzsch et al. [231]. Another fundamental problem is support for automated updating and revising knowledge graphs, as discussed by Chaudhri et al. [90].

Finally, one of the main practical problems in the area is the need for robust automated support for constructing and completing knowledge graphs. Standard ML approaches use embeddings of entities and relations into continuous vector spaces; see, for instance, the survey by Wang et al. [366]. These approaches, however, neither use nor respect the semantics that knowledge representation adds to knowledge graphs. Recent approaches to deal with this issue exploit textual knowledge within knowledge graphs in the form of literals [4] or include background knowledge in the embedding [290].

2.7 Argumentation

Argumentation [30] is concerned with how arguments are supported or undermined by other arguments, and how these arguments and their components (supporting or undermining claims) interact. It addresses the resolving of conflicts among arguments along with measures to evaluate an argument's plausibility.

An argument may be stated as a logical formula, and studying argumentation in such a setting is ongoing. However, arguably the most consequential development in the theory of argumentation in KR has been the emergence of abstract argumentation [130]. This approach does away with the structure of arguments and language-specific inference, and focuses instead on the basic relation between arguments, in particular when one argument *attacks* another. These relations can be represented as a directed graph, or *(abstract) argumentation framework*, where arguments are given by vertices and the attacks relation is a directed edge. The goal is to determine a set of accepted arguments, or *extension*, S, where informally S has no internal attacks relation instance and if an argument in S is attacked by another argument then that argument in turn is *defended* by an argument in S. How this is to be done exactly is not clear, and has led to the development of a large number of different semantics for specifying extensions [31].

Nonetheless, this approach has attracted attention of researchers in various fields of AI, including multi-agent systems, natural language understanding, machine learning and, more recently, explainable AI. Below, we outline some of the major, current research themes in argumentation.

Computational aspects of argumentation. A major theme in argumentation research concerns computational aspects, both regarding complexity of reasoning, and the design and implementation of argumentation systems. Initial research was limited to the domain of abstract argumentation and its plethora of semantics. A comprehensive account of the complexity landscape is presented by Charwat et al. [89] and Cerutti et al. [88]. However, it is also clear that arguments in practical settings have structure, including their claims and support not captured by abstract argumentation. Therefore, formalisms have been developed in which this structure is represented explicitly. These formalisms include assumption-based argumentation [52], ASPIC$^+$ [279, 280], and deductive argumentation [39, 40], among others. Recent papers include work by Lehtonen et al. [241], which provides a treatment of complexity and algorithms for assumption-based argumentation, and Dvořák et al. [131], which deals with the instantiation-based approach [174].

Argumentation for explainable AI. Decisions, choices and recommendations made by humans or by decision-support systems must be *explainable*. AI software systems, whether based on machine learning or logical rules and constraints, are often opaque and their outputs do not come with human-understandable explanations. *Explainable* AI (XAI) aims to address this shortcoming. One of the promising directions in XAI is based on the concept of *argumentative* explanations, and focusses on the human perspective [10]. It turns out that many forms of argumentative explanations can be cast in terms of current argumentation frameworks [99, 318, 311]. This provides a setting for stating explanations and computational tools for establishing them. The approach applies both when explaining models in an argumentation formalism and when explaining machine learning models [107]. Argumentation-based XAI is in its early stages; one of its key challenges is to develop human-computer interaction tools that will effectively construct argumentative explanations in a human-understandable form.

Argument mining. This line of research studies techniques to extract, from text, arguments and their components such as premises and claims. It studies methods to arrange this data into abstract argumentation frameworks and applies inferential techniques to discover fallacies and inconsistencies in the original text. The area can be traced back to the work of Teufel et al. [350], as well as Mochales and Moenas [278]. The survey by Villata and Cabrio [362] provides a good overview of the state of the art and a roadmap for the future. Recent work by Goffredo et al. [173] showcases the potential of argument mining by identifying fallacious arguments from the US presidential campaigns.

Applications. From its inception, research in argumentation has been driven by practical applications. Currently, developing and implementing applications of argumentation is one of the most actively pursued research directions by the community. For instance, in medicine, argumentation is applied as a patient management tool [106] and for persuasion to bring about changes in behaviour [201]; in shared governance, argumentation is used in collaborative decision support systems [333]. Law is another area that shows particularly strong connections with argumentation, both motivating research and benefiting from it [16]. At present no argumentation-based systems aimed at legal reasoning have been deployed in practice, making law a particularly urgent and promising target area. Another emerging area of applications is *judgmental forecasting* [381], a decision-making approach to situations when statistical methods are not applicable. A recently proposed variant of an argumentation formalism, a *forecasting argumentation framework* [206] is guided by forecasting research and aims to support argumentation-based forecasting.

2.8 Reasoning under Uncertainty

The unification of logic and probability has always been a major concern in AI. It is, therefore, not surprising that John McCarthy, who was first to suggest the use of logic for representing the knowledge of AI agents and the calculability of that knowledge, was also concerned about the role of probability. However, he was not quite convinced that there is any easy approach to adding probabilities to knowledge bases [271]. These concerns notwithstanding, there is a great body of work on integrating knowledge representation and uncertainty, as years of engineering efforts in knowledge representation have shown us that there is pervasive uncertainty in almost every domain of interest. The upshot is that the "rigidity" (sentences always evaluate to *true* or *false*), "brittleness" (sentences in the knowledge base must be true in all possible worlds), and "discreteness" of classical logic have forced scientists to look at formalisms such as fuzzy and probabilistic logics. Below we enumerate some of the key developments and current research trends in this area.

Expressive probabilistic logics. Starting with Nilsson's probabilistic logic [297], and Halpern's and Bacchus' investigations on first-order logics of probability [179], we now have a range of expressive modal propositional and first-order logics that allow us to reason about meta-knowledge, actions, plans and programs [228, 299, 176, 32]. In recent years, there has been interest in epistemic planning [29, 35], which can involve the modelling of mental states of multiple agents in domains with noisy effectors and sensors.

Tractable probabilistic logics. Expressive probabilistic logics tend to be hopelessly undecidable [1, 299], and so given the success of probabilistic graphical models [304], there has been a happy marriage of finite-domain relational logic and Bayesian and Markov networks [218, 329, 68], most recently further integrated with neural networks [259]. In particular, to go beyond essentially propositional models, there has been considerable recent work on probabilistic description logics [221], as well as probabilistic logic programs [316]. It is also worth noting that from the database community, a related formalism called probabilistic databases has emerged [346].

Learning of logical axioms. There is a considerable body of work on learning propositional and relational formulas [220, 36], and in the context of probabilistic information, learning weighted formulas [315], probabilistic automata [255], and grammars [236].

Neural reasoning. Given the increasing popularity of neural networks for low-level perception tasks, an emerging concern in both the neural and logical communities is how the two areas can be bridged. The resulting area of neuro-symbolic AI has now come to include approaches such as fuzzy logic [25], logic-based regularisation [150, 192], and differential statistical relational learning [259] and inductive logic programming [135].

For earlier work on reasoning under uncertainty, see the work by Pearl [303], while the Russell and Norvig text [328] provides an extensive overview. For a more recent overview and further details, see the work by Belle [33, 34].

2.9 KR and Machine Learning

The breakthroughs in *machine learning* (ML), in particular the emergence of deep neural networks, have brought about a new area of research on combining machine learning with symbolic or knowledge-based approaches. On the one hand, researchers expect KR methods to help tackle ML problems. On the other hand, they expect ML to help address some of the challenges facing KR systems. Below we discuss how the two fields interact and list some of the key challenges that arise. A more through discussion can be found in the paper by Benedikt et al. [38].

KR for ML. Despite the undeniable success of deep learning, the ability of neural networks to reason and to generalise in systematic ways has been called into question [155, 203, 382]. This view has led to the development of neuro-symbolic methods, which integrate deep learning architectures with explicit symbolic reasoning processes [191]. For instance, DeepProbLog [260] uses probabilistic logic programs to reason about the predictions of a deep learning model. Similar combinations of neural networks with Answer Set Programming [377] and Markov Logic Networks [264] have also been proposed. More broadly, the usefulness of symbolic reasoning for machine learning has been studied in the context of Inductive Logic Programming [286], and more recently, under the umbrella of Statistical Relational Learning [158]. Machine learning models can benefit from reasoning components in at least two different ways. First, the use of (rule-based) reasoning helps models to generalise in predictable and systematic ways. Second, training neuro-symbolic methods in an end-to-end fashion can help bridge the gap between the training data that is available, and the data that would be needed for training the individual components of the system in isolation. To illustrate, in DeepProbLog one can train a neural network to recognise hand-written digits when provided only with training examples that specify the sum of two hand-written digits [260].

Beyond reasoning, machine learning models can also benefit from symbolic knowledge in other ways. For instance, the so-called *knowledge injection* methods use rules and constraints to regularise neural network models [123, 247, 375]. The underlying idea is essentially to discourage the model from making predictions that are incompatible with a given set of (soft) rules. Rules are also sometimes used as a form of weak supervision, to deal with a scarcity of labelled training examples [17]. While rules and constraints encode specific semantic dependencies, simply knowing which concepts are relevant for a given domain can also be important. For instance, Concept Bottleneck Models [219] use such knowledge to ensure that the representation spaces that are learned by neural networks are semantically meaningful, and to some extent interpretable. Other approaches aim to discover semantically meaningful concepts, without prior knowledge, by designing models that learn vectors that can be interpreted as prototypes of concepts [246]. These prototypes can then be used for explanations.

Explainability and interpretability have become important topics within machine learning. Explanations can take many forms, including sets of input features [343], linear combinations of input features [324], or generated natural language sentences [371]. Given that transparency is one of the key strengths of KR systems, it should come as no surprise that ideas from KR often play a central role in this context. For instance, one problem with generating explanations using language models is that such explanations are often not faithful [82]. One possible solution is to develop models that infer the answer by incrementally constructing the analogue of a proof tree [109]. Somewhat related, large language models are capable of generating step-by-step derivations for answers that require reasoning (known as *chains of thought* in this context), by generalising from a few examples of such derivations [368]. However, the proof strategies that are implicitly employed by these models are rather primitive [330], which suggests that a hybridisation with KR methods would improve their capabilities. Specific KR frameworks can also be used more directly. For instance, techniques from argumentation have long been used for modelling explanations [108]. As another example, when explaining entire models, rather than individual predictions, default rules often allow for significantly more compact explanations than decision trees or traditional rules [234]. Another potential application of KR methods is to formally verify properties of ML models, such as their robustness against adversarial attacks [288, 287, 208].

ML for KR. The knowledge acquisition bottleneck is one of the core challenges in KR (see also Section 3.2). KR systems need access to structured knowledge, encoded in some formalism, but such encodings are not readily available for most domains, and tend to be expensive to obtain. There are several ways in which ML techniques can be used to alleviate this issue. ML models can be used, for instance, to identify and exploit statistical regularities in existing knowledge bases. This has resulted in strategies for automatically extending knowledge graphs with plausible triples [55, 376, 356, 347, 26] and methods for identifying plausible missing subsumptions in ontologies [245, 91]. A large number of methods have been developed for converting knowledge expressed in text into a structured format [84]. While most work has focused on knowledge graph triples, learning more expressive knowledge has also been considered [309]. More recently, the focus has shifted to extracting knowledge directly from (large) language models [308, 180, 100]. The success of language models as knowledge acquisition tools, however, extends beyond knowledge graphs. For instance, Hwang et al. [204] use language models to capture social commonsense knowledge, while Jin et al. [211] and Zellers et al. [380] induce script knowledge.

ML methods have also been proposed for improving reasoners, or even as an alternative to the use of classical reasoners. One strand of work is focused on approximating the entire reasoning process using neural networks, for instance to allow for efficient and inconsistency tolerant reasoning with ontologies [2], or to heuristically search for solutions to computationally intractable problems [312]. Another possibility is to use ML methods within traditional reasoners, for instance, for guiding the proof process, e.g., through learning suitable heuristics, (premise) selection strategies, or other techniques to better control the search space [178, 248, 5, 94].

2.10 KR and Robotics

A robot is a physical agent that carries out tasks by interacting with the world via its sensors and effectors. Prominent examples are vacuum robots or robots operating in warehouses, and also include autonomous vehicles in urban environments. Given that robots need to be given or to learn or acquire domain and general knowledge and to reason with this knowledge to solve problems such as building plans, the field of robotics would appear to be a prime application of KR. Indeed, in one of the first robotics projects in the late 1960s, the robot *Shakey* made use of explicit representations of knowledge and was able to plan a course of actions to move objects between rooms. While the reasoning techniques developed for Shakey have long been superseded by more advanced techniques, it is remarkable that the planning language STRIPS [144], which resulted from the Shakey project, has survived until today. However, after Shakey, KR played almost no role in robotics for a long time, mainly because more immediate problems like safe navigation in unknown environments needed to be solved first. A breakthrough in this regard came about in the mid 1990s with the development of probabilistic robotics [353]. It resulted in impressive artefacts like the museum tour guide robots Rhino and Minerva [72, 352], and in the re-emergence of KR techniques like the action language Golog [244] controlling the high-level behaviour of Rhino.

Around the same time, Ray Reiter and colleagues established a brand of *Cognitive Robotics* where KR plays an essential role. To quote from a recent Dagstuhl report on the area [189]:

> *Cognitive Robotics is concerned with endowing robots or software agents with higher level cognitive functions that involve reasoning, for example, about goals, perception, actions, the mental states of other agents, collaborative task execution, etc. This research agenda has historically been pursued by describing, in a language suitable for automated reasoning, enough of the properties of the robot, its abilities, and its environment, to permit it to make high-level decisions about how to act.*

A concise summary of some of the early work on Cognitive Robotics based on the situation calculus is given by Levesque and Lakemeyer [243]. Despite the advances of Cognitive Robotics, it is fair to say that this vision of placing KR at the heart of robotics has not yet been fully realised.[8]

Nevertheless, KR plays an important role in controlling the high-level behaviour of many of today's robotic systems. One of the best known KR frameworks for robots is the *KnowRob* system [349], combining rich ontologies with specialised reasoners like temporal and qualitative spatial reasoners. Planning techniques (see Section 2.4) have also played a major role. While early work focused on domain-dependent planners like IxTeT [159] and TAL [235], domain-independent PDDL planners like FF [193], TFD [136], and ROSPlan [86] are now used routinely in robotic systems. More recently, Golog has been integrated with PDDL planners [98]. Besides task-level planning, goal reasoning has also found its way into the world of robotics. For example, Hofmann et al. [195] consider production logistics scenarios, where robots need to concurrently entertain multiple goals (requests for products), which may need to be prioritised and sometimes abandoned due to timing constraints or failures. An interesting direction specific to robotics is the combination of both task and motion planners (see [152] for a survey), which can help produce more robust plans, especially in scenarios involving object manipulation. Extending this idea to conditional planning [298] allows for handling incomplete information and partial observability.

When a robot has decided on a course of actions, their execution needs to be constantly monitored. An early KR approach to execution monitoring in the context of Golog programs was proposed by De Giacomo et al. [167]. More recently *continual planning* [63, 194] not only tightly integrates planning and execution monitoring, but also allows for postponing plan refinement until enough information is gathered at runtime. Regarding the diagnosis of failures, model-based techniques with strong KR foundations have also been developed [378, 237]. While execution failures cannot be avoided in general, efforts are being made to make robots safer, more robust, and trustworthy. In this regard, verification plays an important role with many connections to KR. For a recent survey on verifying robotic systems, we refer interested readers to the work of Luckcuck et al. [256].

Needless, many challenges remain when it comes to realising the vision of KR-based cognitive robots. A list of many of these challenges can be found in the Dagstuhl report on cognitive robotics [189].

2.11 KR and Information Systems

Ontologies and knowledge graphs (see Sections 2.5 and 2.6) are two key KR contributions to information systems research and practice. In Section 2.5, we also considered ontology-mediated data access, which addresses fundamental questions originating in database research

[8] Indeed, the robotics community has developed other versions of Cognitive Robotics with much less emphasis on KR techniques as in the work of Cangelosi and Asada [83].

by utilising advanced KR techniques. Variants of datalog, description logics, and existential rules are its main underpinning logical formalisms. Closely related to ontology-mediated data access in terms of techniques and goals are declarative approaches to data exchange [137, 11] (transform data structured under a source schema into a target schema using logical rules) and data integration [242, 162] (combine data residing in different sources and provide users with a unified view of them). Description logics have also been proposed as a logical underpinning of standard database design formalisms such as ER and UML diagrams [147, 313]. Highly optimised description logic reasoners discussed earlier are then used to check their consistency and logical consequences. Recently, KR methods developed in non-monotonic reasoning and description logics have been used to further develop the W3C standard SHACL (Shapes Constraint Language) [301] for validating graph-based data against sets of constraints to enhance semantic and technical interoperability [8].

These applications of KR methods to problems in information systems and databases have the common feature that they are concerned with structural and static aspects of information. In contrast, KR methods have only recently been applied when the dynamic behavior of information systems is taken into account. This is despite the fact that the integration of structural and behavioural aspects to capture how information systems dynamically operate over data is recognised as one of the main challenges in business process management (BPM) and, more generally, information systems engineering. Three aspects of BPM are particularly relevant for KR and should be considered when KR methods are applied:

1. Business processes are modelled, configured, executed, and continuously improved based on the so-called business process lifecycle, where every phase calls for reasoning support, ranging from model-driven verification to the combined analysis of models and event data tracing the actual process execution (known as process mining [358]).
2. Business processes vary depending on their complexity, predictability, and repetitiveness. The modelling paradigms used in BPM reflect these differences, ranging for instance from procedural to declarative approaches.
3. Business processes are specified in BPM using concrete modelling patterns. These should guide the use of modelling restrictions needed to achieve decidability or tractability of reasoning tasks.

We briefly discuss the state of the art of KR applications in this field. First, the integrated modelling of processes and data yields infinite-state relational transition systems where each state comes with a first-order interpretation. The analysis of standard properties (such as reachability and safety), or more sophisticated properties expressed in variants of first-order temporal logics have been studied both under complete [80] and incomplete [97, 181] information over relational states, and also in the presence of read-only data [125]. Relevant KR techniques include reasoning about action [81] and planning [56].

Second, KR methods are suitable for the analysis of knowledge-intensive processes, where flexibility is a key requirement [319]. Here methods developed in the context of decidable first-order temporal logic [149, 21] and temporal conceptual modelling [13] are relevant. Applications of KR approaches in declarative process modelling, management, and mining are discussed by Di Ciccio and Montali [95], who make a strong case that KR methods should be utilised to tackle new reasoning problems emerging from process mining. This should facilitate the formal definition of tasks and the exploration of their decidability and complexity status; and it should then also be possible to utilise automated reasoning techniques instead of ad-hoc algorithmic approaches. Notable examples of logic-based techniques employed already are planning [166], SAT [48], and SMT [141].

2.12 KR and Logic/Philosophy

Classical first-order logic was originally developed in an attempt to formalise the foundations of mathematics, with the monumental work of Russel and Whitehead perhaps being the primary example [370]. A good selection of early papers has been collected by Van Heijenoort [360], and an entertaining history of logic has been provided by Doxiadis and Papadimitriou [128]. Subsequent work in philosophy focussed on notions including strict vs. material implication, necessity vs. contingent truth [200], issues concerning naming and reference [230], and the like. So the use of logics – whether propositional, first-order, modal, relevance, or others – to represent and reason about real-world domains can be regarded as a radical shift in the application of logic in general. Indeed, it can be noted that the KR Hypothesis [339], which asserts that a knowledge based approach is *essential* for the construction of any generally intelligent agent, inextricably links logic and KR since it posits semantically meaningful declarative structures in any generally-intelligent agent.

The relation between KR on the one hand, and philosophy and logic on the other hand, has been a mutually beneficial one. The example of belief revision originating as part of philosophy (Section 2.3) is but one of many. Very broadly, KR (and indeed AI as a whole) has been a source of new problems and issues for logic, and through real-world applications has served to sharpen such issues. On the other hand, formal logic has provided the tools and direction for addressing such problems and issues. Indeed, much formal work in KR involves the development and elaboration of new and existing logics.

Here, we merely provide a few illustrative examples of such benefits and synergies. Interesting reasoning problems, such as the frame, qualification, and ramification problems have been identified and addressed [271, 265, 351, 322]. Areas originating in philosophy, such as belief revision (Section 2.3) and deontic reasoning [148], see much (if not most) of their progress coming from the KR community. Much current work in modal logic is driven by research in KR and computer science more broadly. Work in non-monotonic reasoning has contributed significantly to "extended" logical reasoning, in particular fixed-point semantics [124], and the preferential structures advocated for by Shoham [336] and subsequently extended by others. Work in description logics has provided a broad and deep analysis of useful fragments of first-order logic and attendant complexity properties [18]. In the other direction, the realization that many description logics are syntactic variants of modal logics [332] enabled the description logic community to draw on many of the theoretical and technical results that have been developed for modal logics over the years.

2.13 Summary

The areas described in the preceding subsections is to be understood as a starting point for a further, individual literature research; and, certainly, much further works exists within the described areas, between them, and also beyond them. For example, dealing with uncertainty (Section 2.8) is an important issue in essentially all other above described areas, while there are further areas in or linked to KR such as in spatial and qualitative reasoning (see, for example, the survey by Chen et al. [92]), epistemic logic (see, for example, the introductory text by van Benthem [357]), or commonsense reasoning (see, for example, the work by Davis and Marcus [114] or the recent book by Brachman and Levesque [60]).

3 Major Research Challenges for KR

In the previous section, we presented several areas within or related to KR. We included in that discussion research problems and challenges specific to individual areas. In this section, we discuss major overarching challenges for KR that cut across several areas. In doing so, it is useful to keep in mind the following questions that have always been fundamental to KR.

- What knowledge does a system need to have in advance, as opposed to what can be acquired by observations?
- What is the language for representing and reasoning with the background knowledge and observations?
- What kind of semantics governs the updating of knowledge, given new and possibly conflicting observations?
- What impact does continual learning have on the entailments of a knowledge base?
- What is the mechanism for requesting "shallow" versus "deep" reasoning (for example, quickly reacting to a tiger versus reflecting on the nature of the universe)?
- How does a system generalise from low-level observations to high-level structured knowledge? What effect might this have on the computational tractability of the overall system?

As is clear from the remainder of this section, much work remains ahead of us, and these are deeply challenging and exciting questions, where knowledge representation and symbolic approaches more generally will have a significant role to play.

3.1 Commonsense Reasoning

Commonsense reasoning can be defined as reasoning in which an agent has "the ability to make effective use of ordinary, everyday, experiential knowledge in achieving ordinary, everyday, practical goals" [60]. Broadly speaking, this might be thought of as "what a typical seven year old knows about the world, including fundamental categories like time and space, and specific domains such as physical objects and substances; . . . humans, their psychology, and their interactions; and society at large" [113].[9] Clearly it is something that humans possess and use effectively. It is broadly relevant in many areas of AI, including natural language understanding, constructing visual interpretations, planning, and interacting with the real world. In KR, different aspects of commonsense reasoning have been addressed by means of concepts and techniques developed for non-monotonic reasoning, temporal and spatial reasoning, approaches to action and change, qualitative reasoning, and others.

However, no successful commonsense reasoning system of substantial breadth has yet emerged from KR research, efforts by projects such as CYC[10] notwithstanding, and major challenges remain. First, many domains are not well understood, and it is not clear how they may be formalised, both with respect to representing commonsense knowledge and in defining effective inference mechanisms. While substantial progress has been made in areas such as temporal reasoning, and in reasoning about action and change, much remains to be done with regards to physical processes, knowledge and mental attitudes, social mores and attitudes, and the like.

[9] Davis [113] gives a comprehensive survey of issues in commonsense reasoning from a KR perspective. Several points below are drawn from this paper.
[10] `https://cyc.com`

Second, while some knowledge domains have a well-founded theoretical basis and come with effective reasoning tools, combining them into a single coherent knowledge system remains a challenge. It is so, for instance, when an application relies on *different* aspects or instances of commonsense reasoning, in particular, when some of them concern qualitative and some other quantitative knowledge and reasoning.

Third, in many formal settings such as mathematics, databases, information systems, and most forms of declarative programming, reasoning is deductive, and is based on classical logic, or a fragment thereof, as in description logics. However, in commonsense reasoning, conclusions are rarely deductive, but rather are tentative and based on the best information (and perhaps lack of information) available. So reasoning for the most part is *plausible*, and conclusions may be retracted as other information is gleaned. Again, much progress has been made, notably in non-monotonic reasoning on the one hand, and reasoning under uncertainty on the other. But there is no comprehensive account of plausible reasoning in general, nor any substantial understanding of how different approaches relate to each other.

Further, commonsense domains are subject to the so-called *long tail phenomenon.* While general rules may account for a large proportion of domain instances, most often there is also a large (arguably unbounded) number of exceptional instances. This crops up in, for example, the notion of exceptional individuals in default reasoning, and in the *qualification problem* in reasoning about action, wherein it is often impossible to list all preconditions to an action. As Brachman and Levesque [60] note, even when an individual exceptional condition occurs very rarely, there will usually be enough of them that *some* exceptional circumstance will arise with reasonable likelihood. The challenge for a commonsense reasoner is how to handle exceptional circumstances that are impossible to foresee.

Last, it is a natural question to ask whether generative systems such as ChatGPT, trained on vast amounts of human knowledge, possess common sense and to what degree. Thus, an interesting question arises as to how systems such as ChatGPT can be analysed for their commonsense capabilities: where they display comprehensive and convincing commonsense capabilities, and where they are lacking. Related to this is the question of how KR can be used to enhance the commonsense capabilities of such systems. This emphasises another major challenge for KR of developing frameworks to support integration of diverse systems, both knowledge-based and non-declarative.

3.2 Knowledge Acquisition and Maintenance

One of the main challenges for knowledge representation and reasoning was and is the task of knowledge acquisition and maintenance, that is, the formulation of suitable domain knowledge by knowledge engineers and domain experts. While this challenge is not a new one, for instance, Studer et al. discussed the challenges already in 1998 [345], it is still a challenge today. In the area of ontologies, the presence of tools such as the Protégé ontology editor[11] certainly helped to promote the area. In other KR-related areas, for example, planning, the development of methods and tools for modelling support are currently starting to emerge (see, for example, the work of Lin et al. [252] for plan modelling support). The presence of tools for modelling support alone is, however, not sufficient to solve the knowledge acquisition and maintenance problem.

[11] https://protege.stanford.edu/

Some large knowledge graphs are successfully built and maintained by a community, for example, Wikidata [363], or are based on semi-automatic or automatic extraction (see, for example, the work of Ji et al. [207] for a recent survey on knowledge graph acquisition and use) and there are even some patents for knowledge acquisition in knowledge graphs.[12] These examples, however, typically use rather lightweight schema languages and stay away from treating evolving knowledge as a belief revision task.

There are also approaches to learn schema axioms/rules in more expressive languages [372] and first-order rule learning has been extensively studied in the area of Inductive Logic Programming [285, 104]. Such learned knowledge, however, may be inconsistent or it may come with a degree of uncertainty. In order to deal with this, a knowledge engineer could check and disambiguate the learned knowledge, which is difficult, at least for some domains. Another option is better support for reasoning with uncertainty (Section 2.8). Depending on the domain, one might also need support for commonsense reasoning, which we discussed in Section 3.1. Overall, better support for acquiring and maintaining large, complex, but still *consistent* knowledge bases or better capabilities for reasoning with *inconsistent* knowledge is a challenging but necessary task for demonstrating the use of KR techniques within applications [49, 43].

Work on these topics generally requires empirical evaluations, rather than establishing theoretical foundations. The KR community should support such efforts more actively.

3.3 KR and Hybrid Systems

The paradigms that have been developed in KR and ML clearly have complementary strengths. Integrating approaches from both fields into hybrid systems, thus, is intuitively appealing and, as described in Section 2.9, several works already go in this direction and many events now provide special tracks, tutorials, or invited talks in this area.[13] There are, however, still a number of important and non-trivial challenges that need to be overcome. One issue is that KR methods are typically designed to work with clearly defined and carefully formulated knowledge. The knowledge generated by ML methods, on the other hand, is inevitably noisy. Beyond the possibility of errors in the extracted knowledge, there are also issues related to the vagueness of natural language concepts. KR approaches normally assume that predicates have well-defined and precise meanings, even if these meanings can be application-specific and sometimes somewhat arbitrary. When knowledge is obtained from text, we have no guarantees that the concepts or predicates involved have the intended meaning. Similar problems arise in the context of ontology alignment, but the informal nature of extracted knowledge tends to make these problems more severe. In terms of reasoning, neuro-symbolic methods essentially aim to reason about the output of a neural network model. This means that the learning and reasoning components may be only loosely coupled. It would be beneficial to have a tighter integration, where the steps of knowledge extraction and reasoning are intertwined. There are also various challenges of a more practical nature. Work in this area would clearly benefit from more easy-to-use implementations of reasoners, that are both scalable and that can deal with uncertainty.

[12] For example, patent US20180144252A1 “Method and apparatus for completing a knowledge graph” by Fujitsu Ltd.

[13] For example, the track “Reasoning, Learning, and Decision Making” and two of the invited talks at the KR conference in 2024 go in this direction.

In light of recent developments, a very important aspect to consider here is the case of Generative Models and their relationship to KR. Software tools such as ChatGPT offer a mapping from the vast amounts of information used to train them to well-structured, understandable natural language responses to prompts (queries) also presented to it in natural language. The scope of such systems is quite broad with no topic being off-limits. They provide responses to prompts concerning all aspect of human daily life, respond to questions concerning science and solve technical problems. These responses are often quite accurate and detailed. They also demonstrate a good level of reasoning ability. While it is not hard to find prompts that result in incorrect or confusing answers, the question is whether these systems are (or ought to be viewed as) knowledge-based systems. If they are, what are their weaknesses, and how can they be improved? That is, how can they be analysed from a knowledge-based perspective? And if they aren't, why not? Such questions are now critically important as they could well determine the future direction of KR.

Finally, there are a number of unfortunate barriers between the KR and ML communities, going from differences in terminology to differences in expectations for publications (for example, in terms of the balance between theoretical development and empirical validation).

3.4 Explainable AI (XAI)

It has been argued that KR can address the central problem of explainable AI (XAI), which is to develop tools and methods to interpret predictions or recommendations made by models developed using machine learning and to present them in a human-understandable form [383, 354, 103]. Indeed, "interpreters" of machine learning models implemented as declarative, rule-based knowledge representation systems would arguably allow us to build explanations of decisions or recommendation, by employing the abductive reasoning capabilities of such systems. Recent work in KR, for instance in the area of argumentation (see Section 2.7), inductive logic programming [383], and ontologies and knowledge graphs [354, 103], supports the view that KR is potentially a viable approach to XAI. XAI is, however, still in its infancy and providing compelling evidence of applicability of KR for XAI remains an open problem. One point of concern is whether one can build KR-based interpreters to machine learning models that are accurate enough to provide a basis for building explanations. If one could, they could be used instead of machine learning models in the first place.

Even for the more modest goal of explaining predictions and recommendations made using symbolic KR methods, many challenges remain to be solved. In fact, recently there has been a revival of interest in "explainable KR", witnessed for instance by the workshop series Explainable Logic-Based Knowledge Representation (XLoKR), which started in 2020.[14] In principle, if standard symbolic KR methods are used, predictions and decisions should be easy to explain. For example, if knowledge is represented in (some fragment of) first-order logic, and a decision is made based on the result of a first-order reasoning process, then one can use a formal proof to explain a positive reasoning result, and a counter-model to explain a negative result. Despite many years of research, however, in practice things are not so easy. For example, proofs and counter-models may be very large or very hard to comprehend for non-specialist users. To come up with good explanations, even in the symbolic case, many challenges remain to be solved.

[14] https://lat.inf.tu-dresden.de/XLoKR20/

4 Promoting KR

Perhaps the single most important ingredient necessary for success of any field of science is to have an exciting, relevant and potentially high-impact research program that promises to advance broad societal and scientific needs. We have argued in previous sections that KR is an active and important field, both as a subarea of AI, and as a contributor or partner to other areas. Moreover we have suggested that KR is fundamental for the development of any form of general artificial intelligence. Thus, the field of KR as outlined in this manifesto is clearly in a position to continue its vibrant development.

Nevertheless, KR faces significant challenges. It has often been viewed by researchers in AI and, more broadly, CS, as being highly theoretical, perhaps esoteric, and lacking in practical applications. More recently, significant recent successes of machine learning methods, showcased by compelling applications, have overshadowed work in KR. This has resulted in a decrease in open academic positions in KR, as well as in the number of students and researchers attracted to the field. It might also be the reason for a geographic imbalance in KR research activity, with a stronger research contingent to be found in Europe than in other parts of the world. In this section, we suggest a number of measures to promote KR, so that it remains an attractive, relevant, and thriving research area.

4.1 Broadening the Scope of the KR Conference Series

The International Conference on Principles of Knowledge Representation and Reasoning, that is, the KR conference series, has been one of the most significant vehicles driving research in the area of knowledge representation. Consequently, it provides a major means and opportunity for promoting KR research and development in the future. Here we make some suggestions as to how this may be accomplished.

While previous conferences in the series have welcomed and encouraged papers on applications, the conference should nonetheless strive to make more room for research focusing on practical applications relevant to current societal needs and goals. It should become a forum to discuss and showcase not just theoretical advances, but also successful deployments of KR systems. As part of this effort we recommend that program committees make a distinction between actual applications deployed and in use from work that is mostly theoretical (although of course theoretical work should also be motivated by possible future application). Further, as a step towards practical applications, the conference should promote experimental studies by seeking submissions presenting data sets, and collections of benchmark problems, and should sponsor KR system competitions as a mechanism for achieving performance improvements necessary for moving proof-of-concept solutions to practical and effective implementations.

As we argued, it is important to see KR in the broader context of other areas of AI and beyond AI. To promote these synergies and strengthen research cutting across different fields, the KR conference should seek productive collocations with other conferences such as AAAI, ECAI, IJCAI, ISWC, ICAPS, TARK, and ICLP, as well as with prominent conferences in machine learning. Further, the KR conference should continue to be open to participate in federated conferences such as FLoC. Starting a federated conference series centred around knowledge representation and with the KR conference as its centrepiece should also be explored. The program of KR conferences should reflect the importance of promoting crossarea interactions by the strategic selections of invited speakers and tutorials from bordering areas, and by holding panels and special topic sessions covering areas most promising for successful integration with KR.

A related topic concerns the location of KR conferences. We suggest that KR conferences should try to strike a balance between being located where they may benefit and have an impact on the local KR community, and being easily reachable by train for many participants to reduce travel emissions. While esoteric locations may make for interesting travel destinations, they potentially cause significant travel emissions and are often expensive, thus making it difficult for researchers, particularly students, to attend. As well, it is essential that any conference location allows for low-cost accommodation. It should be a goal of the conference to keep registration fees affordable; in selecting a conference venue, it may be worthwhile considering less expensive settings, such as those that may be provided by a university, in place of resorts or high-end conference facilities.

From its inception, the KR conference was distinguished by its emphasis on more mature and in-depth research than typically found in major AI conferences. Consequently the page limit for submissions has generally been 9-10 pages, as opposed to the more usual 5-7 pages (double column, AAAI/IJCAI style). This has arguably resulted in a higher portion of archival-quality papers. However, this higher threshold comes at a price. Authors eager to present the most recent results of their research may find the technical depth expectation hurdle hard to overcome, particularly in the presence of deadlines and with other highly competitive venues with less stringent requirements for a submission. As well, papers that are accepted for a KR conference may be already developed to the degree that makes extensions required by a subsequent journal publication much harder. This has been addressed to some extent by allowing short (4 page) submissions, where these short papers are otherwise subject to the same reviewing criteria. We recommend that the KR steering committee review the current full paper submission requirements, with an eye to increasing submissions while not compromising technical quality. Here are a few suggestions to this end: First, the KR submission requirements could be brought into line with other conferences such as AAAI, ECAI, or IJCAI; however this might be seen as erasing part of KR's unique character. Second, the call for papers could be rephrased to emphasise that regular submissions need not be of maximum length, stating something like "contributions are welcome for both regular papers (in the range of 6 - 9 pages) and short papers (up to 4 pages). Regular papers that are under 9 pages will not be penalised in the reviewing process". (It is ironic that when a conference talks about a page limit of n, while this is technically an upper limit, in practice it is treated as a lower limit by authors and reviewers alike. This suggested change tries to emphasis that shorter regular papers are welcome and encouraged.)

A third, intriguing possibility follows from the observation that accepted KR papers are already at, or close to, the requirements for a journal submission. Consequently, the steering committee should investigate whether there may be a simple process for having regular accepted papers appear as journal articles with minimal additional effort. This is not a new suggestions, and other conferences such as ICLP, SIGGRAPH, and VLBD have been doing this for years. However, given KR's emphasis on longer, mature submissions, KR papers would be uniquely suited for such treatment. *Which* journal may be suitable for this is open to debate, but it may be that JAIR or some other forum would be open to having a special "Transactions in KR", or some such stream.

4.2 Broadening the KR Community

Another essential goal for the KR community is to ensure fresh cohorts of researchers who consider KR an attractive, inspiring and exciting area. It is important to educate students in KR related topics, and to engage them early on in research efforts. It is similarly important to support young researchers, promote collaborations, and recognise research accomplishments. Finally, it is critical to develop programs bringing awareness of KR and its central role in AI research, especially outside of the geographic areas of current KR strength, now primarily Europe and, to a lesser degree, North America. We believe, KR Inc. through its steering committee is in a position to develop initiatives aiming to address these goals.

Beyond some notable exceptions, such as answer set programming and description logics,[15] a general issue concerns a lack of introductory or overview material for areas of KR: the main KR textbook [59] is nearly 20 years old, and the more technical Handbook of KR [359] is over 15 years old. A new textbook would be valuable as would an updated handbook, although this would require substantial effort and it is not clear where such effort would come from. One option would be to develop a repository of KR materials as has been done for example for description logics.[16] Regardless, instructional materials should be developed introducing KR to undergraduate students and more advanced materials aimed at graduate students. The KR steering committee should promote undergraduate research projects in KR and showcase the best of them. It should encourage the KR community to develop tutorials and summer schools on KR topics that would bring the area of KR closer to students and young researchers outside of the regular academic curriculum. KR-themed tutorials at general AI conferences would be particularly well suited for reaching a broader audience. When these programs are offered in person, KR Inc. should offer travel support to students and young researchers, especially from places with limited resources.

An in-person conference is unquestionably valuable, allowing one to meet and interact with researchers, students, and other practitioners, and as an incubator of new research and collaborations. However, recent years have shown that virtual events can be effective in supporting research interactions and the exchange of ideas. These can be especially important in bringing KR to research communities facing economic barriers, and in reaching researchers from underrepresented groups and researchers with family conditions limiting their ability to travel. Tutorials, summer schools, panels, and seminar series are examples of events that lend themselves to virtual offerings. Further, even the KR conference might benefit by incorporating some online events into its program. Most simply, conference presentations can be recorded and made generally available. More radically the conference could switch to a virtual event every two or three years. This is something that can be investigated by the KR community through the KR steering committee.

5 Conclusions

This manifesto is based on the presentations, panels, working groups, and discussions at the Dagstuhl Perspectives workshop 22282 “Current and Future Challenges in Knowledge Representation and Reasoning”. To avoid an overly lengthy introduction, we deliberately

[15] For instance, textbooks and tutorials for the former [153, 251, 214], and the DL handbook [20], the DL textbook [22], or the shorter DL primer [232], for the latter.

[16] https://dl.kr.org/

kept our introduction in Section 1 short. Interested readers may, however, continue with the appendix, which covers the history of KR in more depth. The previous sections highlight that KR is a central, longstanding, and active area of Artificial Intelligence. While some sub-areas of KR have a long tradition, others, such as argumentation or hybrid systems, evolved more recently as core KR research areas. In Section 2, we reviewed these different areas within and related to KR, presented an account of the state of the art as well as challenges within the sub-areas. Section 3 then took a higher-level perspective and introduced general and major research challenges for KR as a whole. Finally, in Section 4, we concluded with recommendations for promoting the area.

6 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 University, 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

References

1 M. Abadi and J.Y. Halpern. Decidability and expressiveness for first-order logics of probability. *Information and Computation*, 112(1):1–36, 1994. doi:10.1006/inco.1994.1049.

2 Ralph Abboud, İsmail İlkan Ceylan, and Thomas Lukasiewicz. Learning to reason: Leveraging neural networks for approximate DNF counting. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 3097–3104. AAAI Press, 2020. URL: https://ojs.aaai.org/index.php/AAAI/article/view/5705, doi:10.1609/AAAI.V34I04.5705.

3 Dirk Abels, Julian Jordi, Max Ostrowski, Torsten Schaub, Ambra Toletti, and Philipp Wanko. Train scheduling with hybrid answer set programming. *Theory and Practice of Logic Programming*, 21(3):317–347, 2021. doi:10.1017/S1471068420000046.

4 Mirza Mohtashim Alam, Md. Rashad Al Hasan Rony, Mojtaba Nayyeri, Karishma Mohiuddin, M. S. T. Mahfuja Akter, Sahar Vahdati, and Jens Lehmann. Language model guided knowledge graph embeddings. *IEEE Access*, 10:76008–76020, 2022. doi:10.1109/ACCESS.2022.3191666.

5 Jesse Alama, Tom Heskes, Daniel Kühlwein, Evgeni Tsivtsivadze, and Josef Urban. Premise selection for mathematics by corpus analysis and kernel methods. *Journal of Automated Reasoning*, 52(2):191–213, 2014. doi:10.1007/s10817-013-9286-5.

6 Carlos E Alchourrón, Peter Gärdenfors, and David Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *The Journal of Symbolic Logic*, 50(2):510–530, 1985. doi:10.2307/2274239.

7 Mario Alviano, Wolfgang Faber, and Martin Gebser. Aggregate semantics for propositional answer set programs. *CoRR*, abs/2109.08662, 2021. arXiv:2109.08662, doi:10.48550/arXiv.2109.08662.

8 Medina Andresel, Julien Corman, Magdalena Ortiz, Juan L. Reutter, Ognjen Savkovic, and Mantas Simkus. Stable model semantics for recursive SHACL. In *WWW '20: The Web Conference 2020*, pages 1570–1580, 2020. doi:10.1145/3366423.3380229.

9 Renzo Angles, Marcelo Arenas, Pablo Barceló, Aidan Hogan, Juan L. Reutter, and Domagoj Vrgoc. Foundations of modern query languages for graph databases. *ACM Computing Surveys*, 50(5):68:1–68:40, 2017. doi:10.1145/3104031.

10 Charles Antaki and Ivan Leudar. Explaining in conversation: Towards an argument model. *European Journal of Social Psychology*, 22(2):181–194, 1992.

11 Marcelo Arenas, Pablo Barceló, Leonid Libkin, and Filip Murlak. *Foundations of Data Exchange*. Cambridge University Press, 2014. URL: http://www.cambridge.org/9781107016163.

12 Marcelo Arenas, Gonzalo I. Diaz, and Egor V. Kostylev. Reverse engineering SPARQL queries. In Jacqueline Bourdeau, Jim Hendler, Roger Nkambou, Ian Horrocks, and Ben Y. Zhao, editors, *Proceedings of the 25th International Conference on World Wide Web, WWW 2016*, pages 239–249. ACM, 2016. doi:10.1145/2872427.2882989.

13 Alessandro Artale and Enrico Franconi. Foundations of temporal conceptual data models. In Alexander Borgida, Vinay K. Chaudhri, Paolo Giorgini, and Eric S. K. Yu, editors, *Conceptual Modeling: Foundations and Applications - Essays in Honor of John Mylopoulos*, volume 5600 of *Lecture Notes in Computer Science*, pages 10–35. Springer, 2009. doi:10.1007/978-3-642-02463-4_2.

14 Alessandro Artale, Roman Kontchakov, Alisa Kovtunova, Vladislav Ryzhikov, Frank Wolter, and Michael Zakharyaschev. First-order rewritability and complexity of two-dimensional temporal ontology-mediated queries. *Journal of Artificial Intelligence Research*, 75:1223–1291, 2022. doi:10.1613/jair.1.13511.

15 Masataro Asai and Christian Muise. Learning neural-symbolic descriptive planning models via cube-space priors: The voyage home (to STRIPS). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 2676–2682. ijcai.org, 2020. doi:10.24963/ijcai.2020/371.

16 Katie Atkinson and Trevor J M Bench-Capon. Argumentation schemes in AI and law. *Argument and Computation*, 12(3):417–434, 2021. doi:10.3233/AAC-200543.

17 Abhijeet Awasthi, Sabyasachi Ghosh, Rasna Goyal, and Sunita Sarawagi. Learning from rules generalizing labeled exemplars. In *8th International Conference on Learning Representations, ICLR 2020*. OpenReview.net, 2020. URL: https://openreview.net/forum?id=SkeuexBtDr.

18 F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, Cambridge, second edition, 2007.

19 Franz Baader, Sebastian Brandt, and Carsten Lutz. Pushing the EL envelope. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30 - August 5, 2005*, pages 364–369. Professional Book Center, 2005. URL: http://ijcai.org/Proceedings/05/Papers/0372.pdf.

20 Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.

21 Franz Baader, Silvio Ghilardi, and Carsten Lutz. LTL over description logic axioms. *ACM Transactions on Computational Logic*, 13(3):21:1–21:32, 2012. doi:10.1145/2287718.2287721.

22 Franz Baader, Ian Horrocks, Carsten Lutz, and Ulrike Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017. URL: http://www.cambridge.org/de/academic/subjects/computer-science/knowledge-management-databases-and-data-mining/introduction-description-logic?format=PB#17zVGeWD2TZUeu6s.97.

23 Franz Baader, Ian Horrocks, and Ulrike Sattler. Description logics. In Frank van Harmelen, Vladimir Lifschitz, and Bruce W. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 135–179. Elsevier, 2008. doi:10.1016/S1574-6526(07)03003-9.

24 Franz Baader, Francesco Kriegel, Adrian Nuradiansyah, and Rafael Peñaloza. Making repairs in description logics more gentle. In Michael Thielscher, Francesca Toni, and Frank Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018*, pages 319–328. AAAI Press, 2018. URL: https://aaai.org/ocs/index.php/KR/KR18/paper/view/18056.

25 Samy Badreddine, Artur d'Avila Garcez, Luciano Serafini, and Michael Spranger. Logic tensor networks. *Artificial Intelligence*, 303:103649, 2022. doi:10.1016/j.artint.2021.103649.

26 Ivana Balazevic, Carl Allen, and Timothy M. Hospedales. Tucker: Tensor factorization for knowledge graph completion. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019*, pages 5184–5193. Association for Computational Linguistics, 2019. doi:10.18653/v1/D19-1522.

27 Mutsunori Banbara, Katsumi Inoue, Benjamin Kaufmann, Tenda Okimoto, Torsten Schaub, Takehide Soh, Naoyuki Tamura, and Philipp Wanko. teaspoon: solving the curriculum-based course timetabling problems with answer set programming. *Annals of Operations Research*, 275(1):3–37, 2019. doi:10.1007/s10479-018-2757-7.

28 Bita Banihashemi, Giuseppe De Giacomo, and Yves Lespérance. Abstraction of agents executing online and their abilities in the situation calculus. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, pages 1699–1706. ijcai.org, 2018. doi:10.24963/ijcai.2018/235.

29 Chitta Baral, Thomas Bolander, Hans van Ditmarsch, and Sheila McIlrath. Epistemic Planning (Dagstuhl Seminar 17231). *Dagstuhl Reports*, 7(6):1–47, 2017. doi:10.4230/DagRep.7.6.1.

30 Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors. *Handbook of Formal Argumentation*. College Publications, 2018.

31 Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10):675–700, 2007. Argumentation in Artificial Intelligence. doi:10.1016/j.artint.2007.04.004.

32 V. Belle and G. Lakemeyer. Reasoning about probabilities in unbounded first-order dynamical domains. In *IJCAI*, 2017. doi:10.24963/ijcai.2017/115.

33 Vaishak Belle. Logic meets probability: Towards explainable AI systems for uncertain worlds. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 5116–5120, 2017. doi:10.24963/ijcai.2017/733.

34 Vaishak Belle. Logic meets learning: From aristotle to neural networks. In Pascal Hitzler and Md Kamruzzaman Sarker, editors, *Neuro-Symbolic Artificial Intelligence - The State of the Art*, Frontiers in Artificial Intelligence and Applications, pages 78–102. IOS Press, Netherlands, jan 2022. doi:10.3233/FAIA210350.

35 Vaishak Belle, Thomas Bolander, Andreas Herzig, and Bernhard Nebel. Epistemic planning: Perspectives on the special issue. *Artificial Intelligence*, page 103842, 2022. doi:10.1016/j.artint.2022.103842.

36 Vaishak Belle and Brendan Juba. Implicitly learning to reason in first-order logic. *Advances in Neural Information Processing Systems*, 32, 2019. URL: https://proceedings.neurips.cc/paper/2019/hash/09fb05dd477d4ae6479985ca56c5a12d-Abstract.html.

37 Luigi Bellomarini, Emanuel Sallinger, and Georg Gottlob. The vadalog system: Datalog-based reasoning for knowledge graphs. *Proc. VLDB Endow.*, 11(9):975–987, 2018. doi:10.14778/3213880.3213888.

38 Michael Benedikt, Kristian Kersting, Phokion G. Kolaitis, and Daniel Neider. Logic and Learning (Dagstuhl Seminar 19361). *Dagstuhl Reports*, 9(9):1–22, 2020. doi:10.4230/DagRep.9.9.1.

39 Philippe Besnard and Anthony Hunter. *Elements of argumentation*, volume 47. MIT press Cambridge, 2008. URL: http://mitpress.mit.edu/books/elements-argumentation.

40 Philippe Besnard and Anthony Hunter. A review of argumentation based on deductive arguments. In Pietro Baroni, Dov Gabbay, Massimilino Giacomin, and Leendert Van der Torre, editors, *Handbook of Formal Argumentation*, pages 437–484. College Publications, 2018.

41 N. Bidoit and C. Froidevaux. Minimalism subsumes default logic and circumscription. In *Proceedings of IEEE Symposium on Logic in Computer Science, LICS-87*, pages 89–97. IEEE Press, 1987.

42 N. Bidoit and C. Froidevaux. Negation by default and unstratifiable logic programs. *Theoretical Computer Science*, 78(1, (Part B)):85–112, 1991. doi:10.1016/0304-3975(51)90004-7.

43 Meghyn Bienvenu. A short survey on inconsistency handling in ontology-mediated query answering. *Künstliche Intelligenz*, 34(4):443–451, 2020. doi:10.1007/s13218-020-00680-9.

44 Meghyn Bienvenu and Magdalena Ortiz. Ontology-mediated query answering with data-tractable description logics. In Wolfgang Faber and Adrian Paschke, editors, *Reasoning Web. Web Logic Rules - 11th International Summer School*, volume 9203 of *Lecture Notes in Computer Science*, pages 218–307. Springer, 2015. doi:10.1007/978-3-319-21768-0_9.

45 Meghyn Bienvenu, Balder ten Cate, Carsten Lutz, and Frank Wolter. Ontology-Based Data Access: A Study through Disjunctive Datalog, CSP, and MMSNP. *ACM Transactions on Database Systems*, 39(4):33:1–33:44, 2014. doi:10.1145/2661643.

46 D. Bobrow, editor. *Special Issue on Nonmonotonic Reasoning*, volume 13(1-2). Elsevier, 1980.

47 Thomas Bolander, Lasse Dissing, and Nicolai Herrmann. Del-based epistemic planning for human-robot collaboration: Theory and implementation. In Meghyn Bienvenu, Gerhard Lakemeyer, and Esra Erdem, editors, *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021*, pages 120–129, 2021. doi:10.24963/kr.2021/12.

48 Mathilde Boltenhagen, Thomas Chatain, and Josep Carmona. Optimized SAT encoding of conformance checking artefacts. *Computing*, 103(1):29–50, 2021. doi:10.1007/s00607-020-00831-8.

49 Glauber De Bona, John Grant, Anthony Hunter, and Sébastien Konieczny. Classifying inconsistency measures using graphs. *J. Artif. Intell. Res.*, 66:937–987, 2019. doi:10.1613/JAIR.1.11852.

50 Piero A. Bonatti. Rational closure for all description logics. *Artificial Intelligence*, 274:197–223, 2019. doi:10.1016/j.artint.2019.04.001.

51 Piero Andrea Bonatti, Iliana Mineva Petrova, and Luigi Sauro. Defeasible reasoning in description logics: An overview on DLN. In Giuseppe Cota, Marilena Daquino, and Gian Luca Pozzato, editors, *Applications and Practices in Ontology Design, Extraction, and Reasoning*, volume 49 of *Studies on the Semantic Web*, pages 178–193. IOS Press, 2020. doi:10.3233/SSW200043.

52 Andrei Bondarenko, Phan Minh Dung, Robert A Kowalski, and Francesca Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial intelligence*, 93(1-2):63–101, 1997. doi:10.1016/S0004-3702(97)00015-5.

53 Richard Booth and Thomas Andreas Meyer. Admissible and restrained revision. *Journal of Artificial Intelligence Research*, 26:127–151, 2006. doi:10.1613/jair.1874.

54 Richard Booth, Thomas Andreas Meyer, Ivan José Varzinczak, and Renata Wassermann. On the link between partial meet, kernel, and infra contraction and its application to horn logic. *Journal of Artificial Intelligence Research*, 42:31–53, 2011. doi:10.1613/jair.3364.

55 Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems*, pages 2787–2795, 2013. URL: https://proceedings.neurips.cc/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html.

56 Stefan Borgwardt, Jörg Hoffmann, Alisa Kovtunova, Markus Krötzsch, Bernhard Nebel, and Marcel Steinmetz. Expressivity of planning with horn description logic ontologies. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022*, pages 5503–5511. AAAI Press, 2022. doi:10.1609/AAAI.V36I5.20489.

57 Elena Botoeva, Boris Konev, Carsten Lutz, Vladislav Ryzhikov, Frank Wolter, and Michael Zakharyaschev. Inseparability and conservative extensions of description logic ontologies: A survey. In Jeff Z. Pan, Diego Calvanese, Thomas Eiter, Ian Horrocks, Michael Kifer, Fangzhen Lin, and Yuting Zhao, editors, *Reasoning Web: Logical Foundation of Knowledge Graph Construction and Query Answering - 12th International Summer School 2016*, volume 9885 of *Lecture Notes in Computer Science*, pages 27–89. Springer, 2016. doi:10.1007/978-3-319-49493-7_2.

58 R.J. Brachman and H.J. Levesque. Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, 3(2):78–93, 1987. doi:10.1111/j.1467-8640.1987.tb00176.x.

59 R.J. Brachman and H.J. Levesque. *Knowledge Representation and Reasoning*. Morgan-Kaufmann, 2004.

60 R.J. Brachman and H.J. Levesque. *Machines Like Us: Toward AI with Common Sense*. The MIT Press, Cambridge, MA, 2022.

61 R.J. Brachman and B.C. Smith. Special issue on knowledge representation. *SIGART Newsletter*, 70, 1980. doi:10.1145/1056751.1056752.

62 Ronald J. Brachman and James G. Schmolze. An Overview of the KL-ONE Knowledge Representation System. *Cognitive Science*, 9:171–216, 1985. doi:10.1207/s15516709cog0902_1.

63 Michael Brenner and Bernhard Nebel. Continual planning and acting in dynamic multiagent environments. *Auton. Agents Multi Agent Syst.*, 19(3):297–331, 2009. doi:10.1007/s10458-009-9081-1.

64 G. Brewka, I. Niemelä, and M. Truszczynski. Nonmonotonic reasoning. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 239–284. Elsevier Science, San Diego, USA, 2008. doi:10.1016/S1574-6526(07)03006-4.

65 Gerhard Brewka, Thomas Eiter, and Mirosław Truszczyński. Answer Set Programming at a Glance. *Communications of the ACM*, 54(12):92–103, 2011. doi:10.1145/2043174.2043195.

66 Gerhard Brewka, Thomas Eiter, and Miroslaw Truszczynski. Answer set programming: An introduction to the special issue. *AI Mag.*, 37(3):5–6, 2016. doi:10.1609/aimag.v37i3.2669.

67 Katarina Britz, Giovanni Casini, Thomas Meyer, Kody Moodley, Uli Sattler, and Ivan Varzinczak. Principles of KLM-style defeasible description logics. *ACM Transactions on Computational Logic*, 22(1), 2020. doi:10.1145/3420258.

68 Guy Van den Broeck, Ingo Thon, Martijn van Otterlo, and Luc De Raedt. Dtproblog: A decision-theoretic probabilistic prolog. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, AAAI'10, pages 1217–1222. AAAI Press, 2010. doi:10.1609/aaai.v24i1.7755.

69 Bruce Buchanan and Edward A. Feigenbaum. Dendral and Meta-dendral: Their Applications Dimension. In *Readings in artificial intelligence*, pages 313–322. Elsevier, 1981.

70 Bruce Buchanan and Edward Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* Addison Wesley, 1984.

71 Bruce Buchanan, Georgia Sutherland, and Edward A.Feigenbaum. Heuristic DENDRAL: A Program for Generating Explanatory Hypotheses. *Organic Chemistry*, 1969.

72 Wolfram Burgard, Armin B. Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner, and Sebastian Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1-2):3–55, 1999. doi:10.1016/S0004-3702(99)00070-3.

73 Paula-Andra Busoniu, Johannes Oetsch, Jörg Pührer, Peter Skocovsky, and Hans Tompits. Sealion: An eclipse-based IDE for answer-set programming with advanced debugging support. *Theory and Practice of Logic Programming*, 13(4-5):657–673, 2013. doi:10.1017/S1471068413000410.

74 Pedro Cabalar, Martín Diéguez, Torsten Schaub, and Anna Schuhmann. Metric temporal answer set programming over timed traces. In Georg Gottlob, Daniela Inclezan, and Marco Maratea, editors, *Logic Programming and Nonmonotonic Reasoning - 16th International Conference, LPNMR 2022*, volume 13416 of *Lecture Notes in Computer Science*, pages 117–130. Springer, 2022. doi:10.1007/978-3-031-15707-3_10.

75 Andrea Calì, Georg Gottlob, and Andreas Pieris. Towards more expressive ontology languages: The query answering problem. *Artificial Intelligence*, 193:87–128, 2012. doi:10.1016/j.artint.2012.08.002.

76 Francesco Calimeri, Wolfgang Faber, Martin Gebser, Giovambattista Ianni, Roland Kaminski, Thomas Krennwallner, Nicola Leone, Marco Maratea, Francesco Ricca, and Torsten Schaub. ASP-Core-2 input language format. *Theory and Practice of Logic Programming*, 20(2):294–309, 2020. doi:10.1017/S1471068419000450.

77 Francesco Calimeri, Davide Fuscà, Simona Perri, and Jessica Zangari. *I*-DLV: The new intelligent grounder of DLV. In Giovanni Adorni, Stefano Cagnoni, Marco Gori, and Marco Maratea, editors, *AI*IA 2016: XVth International Conference of the Italian Association for Artificial Intelligence*, volume 10037 of *Lecture Notes in Computer Science*, pages 192–207. Springer, 2016. doi:10.1007/978-3-319-49130-1_15.

78 Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *Journal of Automated Reasoning*, 39(3):385–429, 2007. doi:10.1007/s10817-007-9078-x.

79 Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Data complexity of query answering in description logics. *Artificial Intelligence*, 195:335–360, 2013. doi:10.1016/j.artint.2012.10.003.

80 Diego Calvanese, Giuseppe De Giacomo, and Marco Montali. Foundations of data-aware process analysis: a database theory perspective. In Richard Hull and Wenfei Fan, editors, *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2013*, pages 1–12. ACM, 2013. doi:10.1145/2463664.2467796.

81 Diego Calvanese, Giuseppe De Giacomo, Marco Montali, and Fabio Patrizi. First-order μ-calculus over generic transition systems and applications to the situation calculus. *Information and Computation*, 259(3):328–347, 2018. doi:10.1016/j.ic.2017.08.007.

82 Oana-Maria Camburu, Brendan Shillingford, Pasquale Minervini, Thomas Lukasiewicz, and Phil Blunsom. Make up your mind! adversarial generation of inconsistent natural language explanations. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online*, pages 4157–4165. Association for Computational Linguistics, 2020. doi:10.18653/v1/2020.acl-main.382.

83 Angelo Cangelosi and Minoru Asada. *Cognitive Robotics*. The MIT Press, may 2022. doi:10.7551/mitpress/13780.001.0001.

84 Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., and Tom M. Mitchell. Toward an architecture for never-ending language learning. In Maria Fox and David Poole, editors, *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010*. AAAI Press, 2010. doi:10.1609/AAAI.V24I1.7519.

85 David Carral, Irina Dragoste, Larry González, Ceriel J. H. Jacobs, Markus Krötzsch, and Jacopo Urbani. Vlog: A rule engine for knowledge graphs. In Chiara Ghidini, Olaf Hartig, Maria Maleshkova, Vojtech Svátek, Isabel F. Cruz, Aidan Hogan, Jie Song, Maxime Lefrançois, and Fabien Gandon, editors, *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference*, volume 11779 of *Lecture Notes in Computer Science*, pages 19–35. Springer, 2019. doi:10.1007/978-3-030-30796-7_2.

86 Michael Cashmore, Maria Fox, Derek Long, Daniele Magazzeni, Bram Ridder, Arnau Carrera, Narcís Palomeras, Natàlia Hurtós, and Marc Carreras. Rosplan: Planning in the robot operating system. In *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015*, pages 333–341, 2015. doi:10.1609/icaps.v25i1.13699.

87 Giovanni Casini, Thomas Meyer, and Ivan Varzinczak. Rational defeasible belief change. In Diego Calvanese, Esra Erdem, and Michael Thielscher, editors, *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020*, pages 213–222, 2020. doi:10.24963/kr.2020/22.

88 Federico Cerutti, Sarah A Gaggl, Matthias Thimm, and Johannes Wallner. Foundations of implementations for formal argumentation. In Pietro Baroni, Dov Gabbay, Massimilino Giacomin, and Leendert Van der Torre, editors, *Handbook of Formal Argumentation*, pages 688–767. College Publications, 2018.

89 Günther Charwat, Wolfgang Dvořák, Sarah A Gaggl, Johannes P Wallner, and Stefan Woltran. Methods for solving reasoning problems in abstract argumentation–a survey. *Artificial intelligence*, 220:28–63, 2015. doi:10.1016/j.artint.2014.11.008.

90 Vinay Chaudhri, Chaitanya Baru, Naren Chittar, Xin Dong, Michael Genesereth, James Hendler, Aditya Kalyanpur, Douglas Lenat, Juan Sequeda, Denny Vrandečić, and Kuansan Wang. Knowledge graphs: Introduction, history and, perspectives. *AI Magazine*, 43(1):17–29, 2022. doi:10.1609/aimag.v43i1.19119.

91 Jiaoyan Chen, Yuan He, Ernesto Jiménez-Ruiz, Hang Dong, and Ian Horrocks. Contextual semantic embeddings for ontology subsumption prediction. *CoRR*, abs/2202.09791, 2022. doi:10.48550/arXiv.2202.09791.

92 Juan Chen, Anthony G. Cohn, Dayou Liu, Sheng-sheng Wang, Jihong Ouyang, and Qiangyuan Yu. A survey of qualitative spatial representations. *Knowl. Eng. Rev.*, 30(1):106–136, 2015. doi:10.1017/S0269888913000350.

93 Amit Chopra, Leendert van der Torre, Harko Verhagen, and Serena Villata. *Handbook of normative multiagent systems*. College Publications, 2018.

94 Karel Chvalovský, Konstantin Korovin, Jelle Piepenbrock, and Josef Urban. Guiding an instantiation prover with graph neural networks. In Ruzica Piskac and Andrei Voronkov, editors, *Proceedings of 24th International Conference on Logic for Programming, Artificial Intelligence and Reasoning*, volume 94 of *EPiC Series in Computing*, pages 112–123. EasyChair, 2023. `doi:10.29007/tp23`.

95 Claudio Di Ciccio and Marco Montali. Declarative process specifications: Reasoning, discovery, monitoring. In Wil M. P. van der Aalst and Josep Carmona, editors, *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*, pages 108–152. Springer, 2022. `doi:10.1007/978-3-031-08848-3_4`.

96 Jens Claßen and James P. Delgrande. An account of intensional and extensional actions, and its application to belief, nondeterministic actions and fallible sensors. In Meghyn Bienvenu, Gerhard Lakemeyer, and Esra Erdem, editors, *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021*, pages 194–204, 2021. `doi:10.24963/kr.2021/19`.

97 Jens Claßen, Gerhard Lakemeyer, and Benjamin Zarrieß. Situation calculus meets description logics. In Carsten Lutz, Uli Sattler, Cesare Tinelli, Anni-Yasmin Turhan, and Frank Wolter, editors, *Description Logic, Theory Combination, and All That - Essays Dedicated to Franz Baader on the Occasion of His 60th Birthday*, volume 11560 of *Lecture Notes in Computer Science*, pages 240–265. Springer, 2019. `doi:10.1007/978-3-030-22102-7_11`.

98 Jens Claßen, Gabriele Röger, Gerhard Lakemeyer, and Bernhard Nebel. Platas - integrating planning and the action language golog. *Künstliche Intelligenz*, 26(1):61–67, 2012. `doi:10.1007/s13218-011-0155-2`.

99 Oana Cocarascu, Andria Stylianou, Kristijonas Čyras, and Francesca Toni. Data-empowered argumentation for dialectically explainable predictions. In *Proceedings of ECAI 2020*, pages 2449–2456. IOS Press, 2020. `doi:10.3233/FAIA200377`.

100 Roi Cohen, Mor Geva, Jonathan Berant, and Amir Globerson. Crawling the internal knowledge-base of language models. *CoRR*, abs/2301.12810, 2023. `doi:10.48550/arXiv.2301.12810`.

101 Anthony G. Cohn and Jochen Renz. Qualitative spatial representation and reasoning. In Frank van Harmelen, Vladimir Lifschitz, and Bruce W. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 551–596. Elsevier, 2008. `doi:10.1016/S1574-6526(07)03013-1`.

102 Alain Colmerauer and Philippe Roussel. The Birth of Prolog. In John A. N. Lee and Jean E. Sammet, editors, *History of Programming Languages Conference, HOPL-II*, pages 37–52. ACM, 1993. `doi:10.1145/154766.155362`.

103 Roberto Confalonieri, Oliver Kutz, and Diego Calvanese, editors. *Proceedings of the Workshop on Data meets Applied Ontologies in Explainable AI (DAO-XAI 2021) part of Bratislava Knowledge September (BAKS 2021), Bratislava, Slovakia, September 18th to 19th, 2021*, volume 2998 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2021. URL: `https://ceur-ws.org/Vol-2998`.

104 Andrew Cropper and Sebastijan Dumančić. Inductive logic programming at 30: A new introduction. *Journal of Artificial Intelligence Research*, 74, sep 2022. `doi:10.1613/jair.1.13507`.

105 Zhenhe Cui, Yongmei Liu, and Kailun Luo. A uniform abstraction framework for generalized planning. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event*, pages 1837–1844. ijcai.org, 2021. `doi:10.24963/ijcai.2021/253`.

106 Kristijonas Čyras, Tiago Oliveira, Amin Karamlou, and Francesca Toni. Assumption-based argumentation with preferences and goals for patient-centric reasoning with interacting clinical guidelines. *Argument & Computation*, 12(2):149–189, 2021. `doi:10.3233/AAC-200523`.

107 Kristijonas Čyras, Antonio Rago, Emanuele Albini, Pietro Baroni, and Francesca Toni. Argumentative xai: a survey. *arXiv preprint arXiv:2105.11266*, 2021. `doi:10.48550/arXiv.2105.11266`.
108 Kristijonas Cyras, Antonio Rago, Emanuele Albini, Pietro Baroni, and Francesca Toni. Argumentative XAI: A survey. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event*, pages 4392–4399. ijcai.org, 2021. `doi:10.24963/ijcai.2021/600`.
109 Bhavana Dalvi, Peter Jansen, Oyvind Tafjord, Zhengnan Xie, Hannah Smith, Leighanna Pipatanangkura, and Peter Clark. Explaining answers with entailment trees. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event*, pages 7358–7370. Association for Computational Linguistics, 2021. `doi:10.18653/v1/2021.emnlp-main.585`.
110 Adnan Darwiche. Bayesian Networks. In Frank Van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors, *Handbook of Knowledge Representation*, pages 467–509. Elsevier, 2008. `doi:10.1016/S1574-6526(07)03011-8`.
111 Adnan Darwiche and Judea Pearl. On the Logic of Iterated Belief Revision. *Artificial intelligence*, 89(1-2):1–29, 1997. `doi:10.1016/S0004-3702(96)00038-0`.
112 Ernest Davis. Physical reasoning. In Frank van Harmelen, Vladimir Lifschitz, and Bruce W. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 597–620. Elsevier, 2008. `doi:10.1016/S1574-6526(07)03014-3`.
113 Ernest Davis. Logical formalizations of commonsense reasoning: A survey. *Journal of Artificial Intelligence Research*, 59:651–723, 2017. `doi:10.1613/jair.5339`.
114 Ernest Davis and Gary Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM*, 58(9):92–103, aug 2015. `doi:10.1145/2701413`.
115 Randall Davis, Bruce Buchanan, and Edward Shortliffe. Production Rules as a Representation for a Knowledge-Based Consultation Program. *Artificial intelligence*, 8(1):15–45, 1977. `doi:10.1016/0004-3702(77)90003-0`.
116 Giuseppe De Giacomo, Riccardo De Masellis, Fabrizio Maria Maggi, and Marco Montali. Monitoring constraints and metaconstraints with temporal logics on finite traces. *ACM Transactions on Software Engineering Methodologies*, 31(4), jul 2022. `doi:10.1145/3506799`.
117 Giuseppe De Giacomo and Sasha Rubin. Automata-Theoretic Foundations of FOND Planning for LTLf and LDLf Goals. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, IJCAI'18, pages 4729–4735. AAAI Press, 2018. `doi:10.24963/ijcai.2018/657`.
118 Johan de Kleer, Alan Mackworth, and Raymond Reiter. Characterizing diagnoses and systems. *Journal of Artificial Intelligence Research*, 52(2-3):197–222, 1992. `doi:10.1016/0004-3702(92)90027-U`.
119 James P. Delgrande. What's in a default? Thoughts on the nature and role of defaults in nonmonotonic reasoning. In Gerhard Brewka, Victor W. Marek, and Miroslaw Truszczynski, editors, *Nonmonotonic Reasoning: Essays Celebrating its 30th Anniversary*. College Publications, 2011.
120 James P. Delgrande and P. Peppas. Belief revision in Horn theories. *Artificial Intelligence Journal*, 218:1–22, 2015. `doi:10.1016/j.artint.2014.08.006`.
121 James P. Delgrande, Pavlos Peppas, and Stefan Woltran. General belief revision. *Journal of the Association of Computing Machinery*, 65(5):1–34, 2018. `doi:10.1145/3203409`.
122 James P. Delgrande and R. Wassermann. Horn clause contraction functions. *Journal of Artificial Intelligence Research*, 48:475–511, nov 2013. `doi:10.1613/jair.4031`.
123 Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. Lifted rule injection for relation embeddings. In Jian Su, Xavier Carreras, and Kevin Duh, editors, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016*, pages 1389–1399. The Association for Computational Linguistics, 2016. `doi:10.18653/v1/d16-1146`.

124 Marc Denecker, Victor W. Marek, and Miroslaw Truszczyński. Uniform semantic treatment of default and autoepistemic logics. *Artificial Intelligence Journal*, 143(1):79–122, 2003. doi:10.1016/S0004-3702(02)00293-X.

125 Alin Deutsch, Richard Hull, Yuliang Li, and Victor Vianu. Automatic verification of database-centric systems. *ACM SIGLOG News*, 5(2):37–56, 2018. doi:10.1145/3212019.3212025.

126 *DLV* project, 2017+. URL: https://dlv.demacs.unical.it/home.

127 Carmine Dodaro, Philip Gasteiger, Kristian Reale, Francesco Ricca, and Konstantin Schekotihin. Debugging non-ground ASP programs: Technique and graphical tools. *Theory and Practice of Logic Programming*, 19(2):290–316, 2019. doi:10.1017/S1471068418000492.

128 A. Doxiadis and C.H. Papadimitriou. *Logicomix.* Bloomsbury, New York, US, 2009.

129 Richard Duda, John Gaschnig, and Peter Hart. Model Design in the PROSPECTOR Consultant System for Mineral Exploration. In *Readings in Artificial Intelligence*, pages 334–348. Elsevier, 1981.

130 Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995. doi:10.1016/0004-3702(94)00041-X.

131 Wolfgang Dvorák, Alexander Greßler, Anna Rapberger, and Stefan Woltran. The complexity landscape of claim-augmented argumentation frameworks. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*, pages 6296–6303. AAAI Press, 2021. doi:10.1609/AAAI.V35I7.16782.

132 Thomas Eiter and Georg Gottlob. Propositional circumscription and extended closed-world reasoning are π_2^p-complete. *Theoretical Computer Science*, 114(2):231–245, 1993. doi:10.1016/0304-3975(93)90073-3.

133 Esra Erdem, Muge Fidan, David F. Manlove, and Patrick Prosser. A general framework for stable roommates problems using answer set programming. *Theory and Practice of Logic Programming*, 20(6):911–925, 2020. doi:10.1017/S1471068420000277.

134 Esra Erdem, Michael Gelfond, and Nicola Leone. Applications of answer set programming. *AI Mag.*, 37(3):53–68, 2016. doi:10.1609/aimag.v37i3.2678.

135 Richard Evans and Edward Grefenstette. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61:1–64, 2018. doi:10.1613/jair.5714.

136 Patrick Eyerich, Robert Mattmüller, and Gabriele Röger. Using the context-enhanced additive heuristic for temporal and numeric planning. In Erwin Prassler, Johann Marius Zöllner, Rainer Bischoff, Wolfram Burgard, Robert Haschke, Martin Hägele, Gisbert Lawitzky, Bernhard Nebel, Paul-Gerhard Plöger, and Ulrich Reiser, editors, *Towards Service Robots for Everyday Environments - Recent Advances in Designing Service Robots for Complex Tasks in Everyday Environments*, volume 76 of *Springer Tracts in Advanced Robotics*, pages 49–64. Springer, 2012. doi:10.1007/978-3-642-25116-0_6.

137 Ronald Fagin, Phokion G. Kolaitis, Renée J. Miller, and Lucian Popa. Data exchange: semantics and query answering. *Theoretical Computer Science*, 336(1):89–124, 2005. doi:10.1016/j.tcs.2004.10.033.

138 Andreas A. Falkner, Gerhard Friedrich, Konstantin Schekotihin, Richard Taupe, and Erich Christian Teppan. Industrial applications of answer set programming. *Künstliche Intelligenz*, 32(2-3):165–176, 2018. doi:10.1007/s13218-018-0548-6.

139 Jorge Fandinno and Markus Hecher. Treewidth-aware complexity in ASP: not all positive cycles are equally hard. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*, pages 6312–6320. AAAI Press, 2021. URL: https://ojs.aaai.org/index.php/AAAI/article/view/16784, doi:10.1609/AAAI.V35I7.16784.

140 Onofrio Febbraro, Kristian Reale, and Francesco Ricca. ASPIDE: integrated development environment for answer set programming. In James P. Delgrande and Wolfgang Faber, editors, *Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011*, volume 6645 of *Lecture Notes in Computer Science*, pages 317–330. Springer, 2011. doi:10.1007/978-3-642-20895-9_37.

141 Paolo Felli, Alessandro Gianola, Marco Montali, Andrey Rivkin, and Sarah Winkler. Cocomot: Conformance checking of multi-perspective processes via SMT. In Artem Polyvyanyy, Moe Thandar Wynn, Amy Van Looy, and Manfred Reichert, editors, *Business Process Management - 19th International Conference, BPM 2021*, volume 12875 of *Lecture Notes in Computer Science*, pages 217–234. Springer, 2021. doi:10.1007/978-3-030-85469-0_15.

142 Johannes Klaus Fichte, Markus Hecher, and Mohamed A. Nadeem. Plausibility reasoning via projected answer set counting - A hybrid approach. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022*, pages 2620–2626. ijcai.org, 2022. doi:10.24963/ijcai.2022/363.

143 Nathanaël Fijalkow, Bastien Maubert, Aniello Murano, Sasha Rubin, and Moshe Vardi. Public and Private Affairs in Strategic Reasoning. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*, pages 132–140, aug 2022. doi:10.24963/kr.2022/14.

144 Richard E. Fikes and Nils J. Nilsson. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial intelligence*, 2(3-4):189–208, 1971. doi:10.1016/0004-3702(71)90010-5.

145 Michael Fisher. Temporal representation and reasoning. In Frank van Harmelen, Vladimir Lifschitz, and Bruce W. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 513–550. Elsevier, 2008. doi:10.1016/S1574-6526(07)03012-X.

146 Kenneth D. Forbus. Qualitative modeling. In Frank van Harmelen, Vladimir Lifschitz, and Bruce W. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 361–393. Elsevier, 2008. doi:10.1016/S1574-6526(07)03009-X.

147 Enrico Franconi and Gary Ng. The i.com tool for intelligent conceptual modeling. In *Proceedings of the 7th International Workshop on Knowledge Representation meets Databases (KRDB 2000)*, pages 45–53, 2000. URL: https://ceur-ws.org/Vol-29/04-franconi.ps.

148 Dov Gabbay, John Horty, Xavier Parent, Ron van der Meyden, and Leendert van der Torre, editors. *Handbook of Deontic Logic and Normative Systems*. College Publications, 2013.

149 Dov Gabbay, Agi Kurucz, Frank Wolter, and Michael Zakharyachev. *Many-dimensional modal logics: theory and applications*. Studies in Logic and the Foundations of Mathematics. Elsevier, 2003.

150 Krzysztof Gajowniczek, Yitao Liang, Tal Friedman, Tomasz Zabkowski, and Guy Van den Broeck. Semantic and generalized entropy loss functions for semi-supervised deep learning. *Entropy*, 22(3):334, 2020. doi:10.3390/e22030334.

151 P. Gärdenfors and D. Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65(2):197–245, 1994. doi:10.1016/0004-3702(94)90017-5.

152 Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Integrated task and motion planning. *CoRR*, abs/2010.01083, 2020. doi:10.48550/arXiv.2010.01083.

153 Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. *Answer Set Programming in Practice*. Springer, 2022.

154 Hector Geffner. Target languages (vs. inductive biases) for learning to act and plan. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022*, pages 12326–12333. AAAI Press, 2022. doi:10.1609/aaai.v36i11.21497.

155 Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard S. Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nat. Mach. Intell.*, 2(11):665–673, 2020. doi:10.1038/s42256-020-00257-z.

156 Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In Robert A. Kowalski and Kenneth A. Bowen, editors, *Logic Programming, Proceedings of the Fifth International Conference and Symposium*, pages 1070–1080. MIT Press, 1988.

157 Michael Gelfond and Vladimir Lifschitz. Representing Action and Change by Logic Programs. *The Journal of Logic Programming*, 17(2-4):301–321, 1993. doi:10.1016/0743-1066(93)90035-F.

158 Lise Getoor and Ben Taskar, editors. *Introduction to Statistical Relational Learning*. Adaptive Computation and Machine Learning. MIT Press, 2007.

159 Malik Ghallab and Hervé Laruelle. Representation and control in ixtet, a temporal planner. In *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems*, pages 61–67, 1994. URL: http://www.aaai.org/Library/AIPS/1994/aips94-011.php.

160 Giuseppe De Giacomo, Paolo Felli, Brian Logan, Fabio Patrizi, and Sebastian Sardiña. Situation calculus for controller synthesis in manufacturing systems with first-order state representation. *Artificial Intelligence*, 302:103598, 2022. doi:10.1016/j.artint.2021.103598.

161 Giuseppe De Giacomo, Luca Iocchi, Marco Favorito, and Fabio Patrizi. Restraining bolts for reinforcement learning agents. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 13659–13662. AAAI Press, 2020. doi:10.1609/aaai.v34i09.7114.

162 Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. Using ontologies for semantic data integration. In Sergio Flesca, Sergio Greco, Elio Masciari, and Domenico Saccà, editors, *A Comprehensive Guide Through the Italian Database Research Over the Last 25 Years*, volume 31 of *Studies in Big Data*, pages 187–202. Springer International Publishing, 2018. doi:10.1007/978-3-319-61893-7_11.

163 Giuseppe De Giacomo and Yves Lespérance. Goal formation through interaction in the situation calculus: A formal account grounded in behavioral science. In Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith, editors, *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20*, pages 294–302. International Foundation for Autonomous Agents and Multiagent Systems, 2020. doi:10.5555/3398761.3398800.

164 Giuseppe De Giacomo, Yves Lespérance, and Fabio Patrizi. Bounded situation calculus action theories. *Artificial Intelligence*, 237:172–203, 2016. doi:10.1016/j.artint.2016.04.006.

165 Giuseppe De Giacomo, Yves Lespérance, and Eugenia Ternovska. Elgolog: A high-level programming language with memory of the execution history. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 2806–2813. AAAI Press, 2020. doi:10.1609/AAAI.V34I03.5669.

166 Giuseppe De Giacomo, Fabrizio Maria Maggi, Andrea Marrella, and Fabio Patrizi. On the disruptive effectiveness of automated planning for LTLf-based trace alignment. In Satinder Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 3555–3561. AAAI Press, 2017. doi:10.1609/AAAI.V31I1.11020.

167 Giuseppe De Giacomo, Raymond Reiter, and Mikhail Soutchanski. Execution monitoring of high-level robot programs. In Anthony G. Cohn, Lenhart K. Schubert, and Stuart C. Shapiro, editors, *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 453–465. Morgan Kaufmann, 1998.

168 Giuseppe De Giacomo and Moshe Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In Francesca Rossi, editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 854–860. IJCAI/AAAI, 2013.

169 Giuseppe De Giacomo and Moshe Y. Vardi. Synthesis for LTL and LDL on finite traces. In Qiang Yang and Michael J. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, pages 1558–1564. AAAI Press, 2015. URL: `http://ijcai.org/Abstract/15/223`.

170 L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. Semantic characterization of rational closure: From propositional logic to description logics. *Artificial Intelligence*, 226:1–33, 2015. `doi:10.1016/j.artint.2015.05.001`.

171 Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. Hermit: An OWL 2 reasoner. *Journal of Automated Reasoning*, 2014. submitted. `doi:10.1007/s10817-014-9305-1`.

172 Birte Glimm and Yevgeny Kazakov. Classical algorithms for reasoning and explanation in description logics. In Markus Krötzsch and Daria Stepanova, editors, *Reasoning Web. Explainable Artificial Intelligence - 15th International Summer School 2019*, volume 11810 of *Lecture Notes in Computer Science*, pages 1–64. Springer, 2019. `doi:10.1007/978-3-030-31423-1_1`.

173 Pierpaolo Goffredo, Shohreh Haddadan, Vorakit Vorakitphan, Elena Cabrio, and Serena Villata. Fallacious argument classification in political debates. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence, IJCAI 2022*, pages 4143–4149, 2022. `doi:10.24963/ijcai.2022/575`.

174 Nikos Gorogiannis and Anthony Hunter. Instantiating abstract argumentation with classical logic arguments: Postulates and properties. *Artificial Intelligence*, 175(9-10):1479–1497, 2011. `doi:10.1016/j.artint.2010.12.003`.

175 Georg Gottlob. Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2(3):397–425, 1992. `doi:10.1093/logcom/2.3.397`.

176 Henrik Grosskreutz and Gerhard Lakemeyer. Probabilistic complex actions in golog. *Fundamenta Informaticae*, 57(2-4):167–192, 2003.

177 Susana Hahn, Orkunt Sabuncu, Torsten Schaub, and Tobias Stolzmann. Clingraph: ASP-based visualization. In Georg Gottlob, Daniela Inclezan, and Marco Maratea, editors, *Logic Programming and Nonmonotonic Reasoning - 16th International Conference, LPNMR 2022*, volume 13416 of *Lecture Notes in Computer Science*, pages 401–414. Springer, 2022. `doi:10.1007/978-3-031-15707-3_31`.

178 Shai Haim and Toby Walsh. Restart strategy selection using machine learning techniques. In Oliver Kullmann, editor, *Theory and Applications of Satisfiability Testing - SAT 2009, 12th International Conference, SAT 2009*, volume 5584 of *Lecture Notes in Computer Science*, pages 312–325. Springer, 2009. `doi:10.1007/978-3-642-02777-2_30`.

179 Joseph Y. Halpern. An analysis of first-order logics of probability. *Artificial Intelligence*, 46(3):311–350, 1990. `doi:10.1016/0004-3702(90)90019-V`.

180 Shibo Hao, Bowen Tan, Kaiwen Tang, Hengzhe Zhang, Eric P. Xing, and Zhiting Hu. Bertnet: Harvesting knowledge graphs from pretrained language models. *CoRR*, abs/2206.14268, 2022. `doi:10.48550/arXiv.2206.14268`.

181 Babak Bagheri Hariri, Diego Calvanese, Marco Montali, Giuseppe De Giacomo, Riccardo De Masellis, and Paolo Felli. Description logic knowledge and action bases. *Journal of Artificial Intelligence Research*, 46:651–686, 2013. `doi:10.1613/jair.3826`.

182 Amelia J. Harrison, Vladimir Lifschitz, and Fangkai Yang. The semantics of gringo and infinitary propositional formulas. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014*. AAAI Press, 2014. URL: `http://www.aaai.org/ocs/index.php/KR/KR14/paper/view/7966`.

183 Patrick J. Hayes and Peter F. Patel-Schneider, editors. *RDF 1.1 Semantics*. W3C Recommendation, 25 February 2014. Available at `https://www.w3.org/TR/rdf11-mt/`.

184 P.J. Hayes. In defense of logic. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 559–565, Cambridge, MA, 1977. URL: `http://ijcai.org/Proceedings/77-1/Papers/099.pdf`.

185 P.J. Hayes. The logic of frames. In D. Metzing, editor, *Frame Conceptions and Text Understanding*, pages 46–61. Walter de Gruyter and Co., 1979.

186 P.J. Hayes. Naive physics I: Ontology for liquids. In J.R. Hobbs and R.C. Moore, editors, *Formal Theories of the Commonsense World*, pages 71–108. Ablex, 1985.

187 P.J. Hayes. The second naive physics manifesto. In J.R. Hobbs and R.C. Moore, editors, *Formal Theories of the Commonsense World*, pages 1–36. Ablex, 1985.

188 Markus Hecher. Treewidth-aware reductions of normal ASP to SAT - is normal ASP harder than SAT after all? *Artificial Intelligence*, 304:103651, 2022. `doi:10.1016/j.artint.2021.103651`.

189 Fredrik Heintz, Gerhard Lakemeyer, and Sheila McIlraith. Cognitive Robotics (Dagstuhl Seminar 22391). *Dagstuhl Reports*, 12(9):200–219, 2023. `doi:10.4230/DagRep.12.9.200`.

190 C. Hewitt. Planner: A language for proving theorems in robots. In *International Joint Conference on Artificial Intelligence*, London, U.K., 1971.

191 Pascal Hitzler and Md. Kamruzzaman Sarker, editors. *Neuro-Symbolic Artificial Intelligence: The State of the Art*, volume 342 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2021. `doi:10.3233/FAIA342`.

192 Nick Hoernle, Rafael Michael Karampatsis, Vaishak Belle, and Kobi Gal. Multiplexnet: Towards fully satisfied logical constraints in neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5700–5709, 2022. `doi:10.1609/aaai.v36i5.20512`.

193 Jörg Hoffmann and Bernhard Nebel. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research*, 14:253–302, 2001. `doi:10.1613/jair.855`.

194 Till Hofmann, Tim Niemueller, Jens Claßen, and Gerhard Lakemeyer. Continual planning in golog. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 3346–3353. AAAI Press, 2016. `doi:10.1609/AAAI.V30I1.10414`.

195 Till Hofmann, Tarik Viehmann, Mostafa Gomaa, Daniel Habering, Tim Niemueller, and Gerhard Lakemeyer. Multi-agent goal reasoning with the CLIPS executive in the robocup logistics league. In Ana Paula Rocha, Luc Steels, and H. Jaap van den Herik, editors, *Proceedings of the 13th International Conference on Agents and Artificial Intelligence, ICAART 2021*, pages 80–91. SCITEPRESS, 2021. `doi:10.5220/0010252600800091`.

196 Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia D'amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. Knowledge graphs. *ACM Computing Surveys*, 54(4), jul 2021. `doi:10.1145/3447772`.

197 Matthew Horridge. *Justification based explanation in ontologies.* PhD thesis, University of Manchester, UK, 2011. URL: `http://www.manchester.ac.uk/escholar/uk-ac-man-scw:131699`.

198 Ian Horrocks. DAML+OIL: A description logic for the semantic web. *IEEE Data Eng. Bull.*, 25(1):4–9, 2002. URL: `http://sites.computer.org/debull/A02MAR-CD.pdf`.

199 John Horty. Epistemic oughts in stit semantics. *Ergo, an Open Access Journal of Philosophy*, 6, 2019.

200 G.E. Hughes and M.J. Cresswell. *An Introduction to Modal Logic.* Methuen and Co. Ltd., 1968.

201 Anthony Hunter. Towards a framework for computational persuasion with applications in behaviour change. *Argument & Computation*, 9(1):15–40, 2018. doi:10.3233/AAC-170032.

202 Anthony Hunter, Gabriele Kern-Isberner, Thomas Meyer, and Renata Wassermann. The role of non-monotonic reasoning in future development of artificial intelligence (Dagstuhl Perspectives workshop 19072). *Dagstuhl Reports*, 9(2):73–90, 2019. doi:10.4230/DagRep.9.2.73.

203 Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research*, 67:757–795, 2020. doi:10.1613/jair.1.11674.

204 Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. COMET-ATOMIC 2020: On symbolic and neural commonsense knowledge graphs. *CoRR*, abs/2010.05953, 2020. arXiv:2010.05953, doi:10.48550/arXiv.2010.05953.

205 Rodrigo Toro Icarte, Toryn Q. Klassen, Richard Anthony Valenzano, and Sheila A. McIlraith. Reward machines: Exploiting reward function structure in reinforcement learning. *Journal of Artificial Intelligence Research*, 73:173–208, 2022. doi:10.1613/jair.1.12440.

206 Benjamin Irwin, Antonio Rago, and Francesca Toni. Forecasting argumentation frameworks. *arXiv preprint arXiv:2205.11590*, 2022. doi:10.48550/arXiv.2205.11590.

207 Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):494–514, 2022. doi:10.1109/TNNLS.2021.3070843.

208 Kai Jia and Martin Rinard. Efficient exact verification of binarized neural networks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20. Curran Associates Inc., 2020. URL: https://proceedings.neurips.cc/paper/2020/hash/1385974ed5904a438616ff7bdb3f7439-Abstract.html.

209 E. Jiménez Ruiz, B. Cuenca Grau, I. Horrocks, and R. Berlanga. Supporting concurrent ontology development: Framework, algorithms and tool. *Data and Knowledge Engineering*, 70(1):146–164, 2011. doi:10.1016/j.datak.2010.10.001.

210 Y. Jin and M. Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007. doi:10.1016/j.artint.2006.11.002.

211 Zijian Jin, Xingyu Zhang, Mo Yu, and Lifu Huang. Probing script knowledge from pre-trained models. *CoRR*, abs/2204.10176, 2022. doi:10.48550/arXiv.2204.10176.

212 Jean Christoph Jung, Carsten Lutz, Hadrien Pulcini, and Frank Wolter. Logical separability of labeled data examples under ontologies. *Artificial Intelligence*, 313:103785, 2022. doi:10.1016/j.artint.2022.103785.

213 Mohimenul Kabir, Flavio O. Everardo, Ankit K. Shukla, Markus Hecher, Johannes Klaus Fichte, and Kuldeep S. Meel. ApproxASP - a scalable approximate answer set counter. In *Thirty-Sixth AAAI Conference on Artificial Intelligence*, pages 5755–5764. AAAI Press, 2022. doi:10.1609/AAAI.V36I5.20518.

214 Roland Kaminski, Javier Romero, Torsten Schaub, and Philipp Wanko. How to build your own ASP-based system?! *CoRR*, abs/2008.06692, 2020. doi:10.48550/arXiv.2008.06692.

215 Roland Kaminski and Torsten Schaub. On the foundations of grounding in answer set programming. *CoRR*, abs/2108.04769, 2021. doi:10.48550/arXiv.2108.04769.

216 H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In P. Gärdenfors, editor, *Belief Revision*, pages 183–203, Cambridge, 1992. Cambridge University Press.

217 Yevgeny Kazakov, Markus Krötzsch, and František Simančík. The incredible ELK: From polynomial procedures to efficient reasoning with $\mathcal{EL}$ ontologies. *Journal of Automated Reasoning*, 53(1):1–61, 2014. doi:10.1007/s10817-013-9296-3.

218 K. Kersting and L. De Raedt. Bayesian logic programs. In *ILP Work-in-progress reports*, 2000. URL: `https://ceur-ws.org/Vol-35/07-KerstingDeRaedt.ps`.

219 Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 5338–5348. PMLR, 2020. URL: `http://proceedings.mlr.press/v119/koh20a.html`.

220 Stanley Kok and Pedro Domingos. Learning markov logic networks using structural motifs. In *ICML*, pages 551–558, 2010. URL: `https://icml.cc/Conferences/2010/papers/502.pdf`.

221 D. Koller, A. Levy, and A. Pfeffer. P-classic: a tractable probablistic description logic. In *Proc. AAAI / IAAI*, pages 390–397, 1997. URL: `http://www.aaai.org/Library/AAAI/1997/aaai97-060.php`.

222 Boris Konev, Dirk Walther, and Frank Wolter. Forgetting and uniform interpolation in large-scale description logic terminologies. In Craig Boutilier, editor, *IJCAI*, pages 830–835, 2009. URL: `http://ijcai.org/Proceedings/09/Papers/142.pdf`.

223 S. Konieczny and R. Pino Pérez. Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002. `doi:10.1093/logcom/12.5.773`.

224 Sébastien Konieczny. On the Difference between Merging Knowledge Bases and Combining them. In Anthony G. Cohn, Fausto Giunchiglia, and Bart Selman, editors, *Proceedings od the 7th International Conference on Principles of Knowledge Representation and Reasoning, KR 2000*, pages 135–144. Morgan Kaufmann, 2000.

225 Robert Kowalski. Predicate Logic as Programming Language. In *IFIP congress*, volume 74, pages 569–544, 1974.

226 Robert Kowalski and Donald Kuehner. Linear Resolution with Selection Function. *Artificial Intelligence*, 2(3-4):227–260, 1971. `doi:10.1016/0004-3702(71)90012-9`.

227 Robert A. Kowalski and Marek J. Sergot. A Logic-based Calculus of Events. *New Gener. Comput.*, 4(1):67–95, 1986. `doi:10.1007/BF03037383`.

228 D. Kozen. A probabilistic PDLL. *Journal of Computer and System Sciences*, 30(2):162–178, 1985. `doi:10.1016/0022-0000(85)90012-1`.

229 S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2):167–207, 1990. `doi:10.1016/0004-3702(90)90101-5`.

230 S.A. Kripke. *Naming and Necessity.* Harvard University Press, Cambridge, Mass., 1980.

231 Markus Krötzsch, Maximilian Marx, Ana Ozaki, and Veronika Thost. Attributed description logics: Reasoning on knowledge graphs. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, pages 5309–5313. ijcai.org, 2018. `doi:10.24963/ijcai.2018/743`.

232 Markus Krötzsch, František Simančík, and Ian Horrocks. Description logics. *IEEE Intelligent Systems*, 29(1):12–19, 2014. `doi:10.1109/MIS.2013.123`.

233 Benjamin Kuipers, Edward Feigenbaum, Peter E. Hart, and Nils J. Nilsson. Shakey: from conception to history. *AI Magazine*, 38(1):88–103, 2017. `doi:10.1609/aimag.v38i1.2716`.

234 Ondrej Kuzelka, Jesse Davis, and Steven Schockaert. Learning possibilistic logic theories from default rules. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*, pages 1167–1173. IJCAI/AAAI Press, 2016. URL: `http://www.ijcai.org/Abstract/16/169`.

235 Jonas Kvarnström and Patrick Doherty. Talplanner: A temporal logic based forward chaining planner. *Annals of Mathematics and Artificial Intelligence*, 30(1-4):119–169, 2000. `doi:10.1023/A:1016619613658`.

236 B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015. doi:10.1126/science.aab3050.

237 Thomas Léauté and Brian C. Williams. Coordinating agile systems through the model-based execution of temporal plans. In *Proceedings, The Twentieth National Conference on Artificial Intelligence*, pages 114–120, 2005. URL: http://www.aaai.org/Library/AAAI/2005/aaai05-019.php.

238 Joohyung Lee and Zhun Yang. LPMLN, weak constraints, and p-log. In Satinder Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 1170–1177. AAAI Press, 2017. doi:10.1609/aaai.v31i1.10686.

239 Kevin Lee, Thomas Andreas Meyer, Jeff Z. Pan, and Richard Booth. Computing maximally satisfiable terminologies for the description logic $\mathcal{ALC}$ with cyclic definitions. In Bijan Parsia, Ulrike Sattler, and David Toman, editors, *Proceedings of the 2006 International Workshop on Description Logics (DL2006)*, volume 189 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2006. URL: https://ceur-ws.org/Vol-189/submission_29.pdf.

240 Daniel Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992. doi:10.1016/0004-3702(92)90041-U.

241 Tuomo Lehtonen, Johannes P Wallner, and Matti Järvisalo. Declarative algorithms and complexity results for assumption-based argumentation. *Journal of Artificial Intelligence Research*, 71:265–318, 2021. doi:10.1613/jair.1.12479.

242 Maurizio Lenzerini. Data integration: A theoretical perspective. In *Proceedings of the Twenty-first ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 233–246, 2002. doi:10.1145/543613.543644.

243 Hector J. Levesque and Gerhard Lakemeyer. Cognitive robotics. In Frank van Harmelen, Vladimir Lifschitz, and Bruce W. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 869–886. Elsevier, 2008. doi:10.1016/S1574-6526(07)03023-4.

244 Hector J. Levesque, Raymond Reiter, Yves Lespérance, Fangzhen Lin, and Richard B. Scherl. GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31(1-3):59–83, 1997. doi:10.1016/S0743-1066(96)00121-5.

245 Na Li, Zied Bouraoui, and Steven Schockaert. Ontology completion using graph convolutional networks. In Chiara Ghidini, Olaf Hartig, Maria Maleshkova, Vojtech Svátek, Isabel F. Cruz, Aidan Hogan, Jie Song, Maxime Lefrançois, and Fabien Gandon, editors, *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference*, volume 11778 of *Lecture Notes in Computer Science*, pages 435–452. Springer, 2019. doi:10.1007/978-3-030-30793-6_25.

246 Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 3530–3537. AAAI Press, 2018. doi:10.1609/AAAI.V32I1.11771.

247 Tao Li and Vivek Srikumar. Augmenting neural networks with first-order logic. In Anna Korhonen, David R. Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019*, pages 292–302. Association for Computational Linguistics, 2019. doi:10.18653/v1/p19-1028.

248 Jia Hui Liang, Vijay Ganesh, Pascal Poupart, and Krzysztof Czarnecki. Learning rate based branching heuristic for SAT solvers. In Nadia Creignou and Daniel Le Berre, editors, *Theory and Applications of Satisfiability Testing - SAT 2016 - 19th International Conference*, volume 9710 of *Lecture Notes in Computer Science*, pages 123–140. Springer, 2016. doi:10.1007/978-3-319-40970-2_9.

249 Yuliya Lierler. Constraint answer set programming: Integrational and translational (or smt-based) approaches. *CoRR*, abs/2107.08252, 2021. doi:10.48550/arXiv.2107.08252.

250 Yuliya Lierler, Marco Maratea, and Francesco Ricca. Systems, engineering environments, and competitions. *AI Mag.*, 37(3):45–52, 2016. doi:10.1609/aimag.v37i3.2675.

251 Vladimir Lifschitz. *Answer Set Programming*. Springer, 2019. doi:10.1007/978-3-030-24658-7.

252 Songtuan Lin, Alban Grastien, and Pascal Bercher. Towards automated modeling assistance: An efficient approach for repairing flawed planning domains. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI 2023)*. AAAI Press, 2023. doi:10.1609/aaai.v37i10.26418.

253 Felix Lindner, Robert Mattmüller, and Bernhard Nebel. Evaluation of the moral permissibility of action plans. *Artificial Intelligence*, 287, 2020. doi:10.1016/j.artint.2020.103350.

254 Daxin Liu and Gerhard Lakemeyer. Reasoning about beliefs and meta-beliefs by regression in an expressive probabilistic action logic. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021*, pages 1951–1958. ijcai.org, 2021. doi:10.24963/ijcai.2021/269.

255 Daniel Lowd and Amirmohammad Rooshenas. Learning markov networks with arithmetic circuits. In *AISTATS*, 2013. URL: http://proceedings.mlr.press/v31/lowd13a.html.

256 Matt Luckcuck, Marie Farrell, Louise A. Dennis, Clare Dixon, and Michael Fisher. Formal specification and verification of autonomous robotic systems: A survey. *ACM Computing Surveys*, 52(5):100:1–100:41, 2019. doi:10.1145/3342355.

257 Thomas Lukasiewicz. Expressive probabilistic description logics. *Artificial Intelligence*, 172(6):852–883, 2008. doi:10.1016/j.artint.2007.10.017.

258 Carsten Lutz, Frank Wolter, and Michael Zakharyaschev. Temporal description logics: A survey. In Stéphane Demri and Christian S. Jensen, editors, *15th International Symposium on Temporal Representation and Reasoning, TIME 2008*, pages 3–14. IEEE Computer Society, 2008. doi:10.1109/TIME.2008.14.

259 Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. DeepProbLog: Neural Probabilistic Logic Programming. *Advances in Neural Information Processing Systems*, 31, 2018. URL: https://proceedings.neurips.cc/paper/2018/hash/dc5d637ed5e62c36ecb73b654b05ba2a-Abstract.html.

260 Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Neural probabilistic logic programming in DeepProbLog. *Artificial Intelligence*, 298:103504, 2021. doi:10.1016/j.artint.2021.103504.

261 V. Marek and M. Truszczynski. Stable Semantics for Logic Programs and Default Theories. In Ewing L. Lusk and Ross A. Overbeek, editors, *Proceedings of the North American Conference on Logic Programming, NACLP 1989*, pages 243–256. MIT Press, 1989.

262 Victor W. Marek and Miroslaw Truszczynski. *Nonmonotonic Logic: Context-dependent Reasoning*. Springer, 1993.

263 Victor W. Marek and Miroslaw Truszczynski. Stable models and an alternative logic programming paradigm. In Krzysztof R. Apt, Victor W. Marek, Mirek Truszczynski, and David Scott Warren, editors, *The Logic Programming Paradigm - A 25-Year Perspective*, Artificial Intelligence, pages 375–398. Springer, 1999. doi:10.1007/978-3-642-60085-2_17.

264 Giuseppe Marra and Ondrej Kuzelka. Neural markov logic networks. In Cassio P. de Campos, Marloes H. Maathuis, and Erik Quaeghebeur, editors, *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event*, volume 161 of *Proceedings of Machine Learning Research*, pages 908–917. AUAI Press, 2021. URL: https://proceedings.mlr.press/v161/marra21a.html.

265 J. McCarthy. Epistemological problems in artificial intelligence. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1038–1044, Cambridge, MA, 1977. URL: http://ijcai.org/Proceedings/77-2/Papers/094.pdf.

266 J. McCarthy. First order theories of individual concepts and propositions. In D. Michie, editor, *Machine Intelligence 9*, pages 129–147. Edinburgh University Press, 1979.
267 J. McCarthy. Circumscription – a form of non-monotonic reasoning. *Artificial Intelligence Journal*, 13:27–39, 1980. `doi:10.1016/0004-3702(80)90011-9`.
268 J. McCarthy and P.J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In D. Michie and B. Meltzer, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.
269 John McCarthy. Programs with Common Sense, 1959. URL: `http://jmc.stanford.edu/articles/mcc59/mcc59.pdf`.
270 John McCarthy. Situations, actions, and causal laws. Technical report, Stanford University, Department of Computer Science, 1963.
271 John McCarthy and Patrick J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969. reprinted in McC90.
272 Drew McDermott. Nonmonotonic Logic II: Nonmonotonic Modal Theories. *Journal of the ACM (JACM)*, 29(1):33–57, 1982. `doi:10.1145/322290.322293`.
273 Drew McDermott and Jon Doyle. Non-monotonic Logic I. *Artificial intelligence*, 13(1-2):41–72, 1980. `doi:10.1016/0004-3702(80)90012-0`.
274 John P. McDermott. R1: an Expert in the Computer Systems Domain. In Robert Balzer, editor, *Proceedings of the 1st Annual National Conference on Artificial Intelligence, AAAI 1980*, pages 269–271. AAAI Press/MIT Press, 1980. URL: `http://www.aaai.org/Library/AAAI/1980/aaai80-076.php`.
275 Deborah L. McGuinness and Frank Van Harmelen. OWL Web Ontology Language Overview. *W3C recommendation*, 10(10):2004, 2004.
276 Thomas Andreas Meyer, Kevin Lee, and Richard Booth. Knowledge integration for description logics. In Manuela M. Veloso and Subbarao Kambhampati, editors, *Proceedings, The Twentieth National Conference on Artificial Intelligence*, pages 645–650. AAAI Press / The MIT Press, 2005. URL: `http://www.aaai.org/Library/AAAI/2005/aaai05-101.php`.
277 Marvin Minsky. A Framework for Representing Knowledge. MIT-AI Laboratory Memo 306, June, 1974, Reprinted in *The Psychology of Computer Vision*, P. Winston (Ed.), McGraw-Hill, 1975.
278 Raquel Mochales and Marie-Francine Moens. Argumentation mining. *Artificial Intelligence and Law*, 19:1–22, 2011. `doi:10.1007/s10506-010-9104-x`.
279 Sanjay Modgil and Henry Prakken. A general account of argumentation with preferences. *Artificial Intelligence*, 195:361–397, 2013. `doi:10.1016/j.artint.2012.10.008`.
280 Sanjay Modgil and Henry Prakken. Abstract rule-based argumentation. In Pietro Baroni, Dov Gabbay, Massimilino Giacomin, and Leendert Van der Torre, editors, *Handbook of Formal Argumentation*, pages 287–364. College Publications, 2018.
281 Kodylan Moodley, Thomas Meyer, and Uli Sattler. DIP: A defeasible-inference platform for OWL ontologies. In Meghyn Bienvenu, Magdalena Ortiz, Riccardo Rosati, and Mantas Simkus, editors, *Informal Proceedings of the 27th International Workshop on Description Logics*, volume 1193 of *CEUR Workshop Proceedings*, pages 671–683. CEUR-WS.org, 2014. URL: `http://ceur-ws.org/Vol-1193/paper_44.pdf`.
282 Robert C. Moore. Possible-world semantics for autoepistemic logic. In *Proceedings of the Non-Monotonic Reasoning Workshop*, pages 344–354. American Association for Artificial Intelligence (AAAI), 1984.
283 Robert C Moore. Semantical Considerations on Nonmonotonic Logic. *Artificial Intelligence*, 25(1):75–94, 1985. `doi:10.1016/0004-3702(85)90042-6`.

284 Boris Motik, Bernardo Cuenca Grau, Ian Horrocks, Zhe Wu, Achille Fokoue, and Carsten Lutz, editors. *OWL 2 Web Ontology Language: Profiles.* W3C Recommendation, 27 October 2009. URL: `http://www.w3.org/TR/owl2-profiles/`.

285 Stephen Muggleton. Inductive logic programming. *New Generation Computing*, 8(4):295–318, 1991. `doi:10.1007/BF03037089`.

286 Stephen H. Muggleton and Luc De Raedt. Inductive logic programming: Theory and methods. *Journal of Logic Programming*, 19/20:629–679, 1994. `doi:10.1016/0743-1066(94)90035-3`.

287 Nina Narodytska. Formal verification of deep neural networks. In Nikolaj S. Bjørner and Arie Gurfinkel, editors, *2018 Formal Methods in Computer Aided Design, FMCAD 2018*, page 1. IEEE, 2018. `doi:10.23919/FMCAD.2018.8603017`.

288 Nina Narodytska, Shiva Prasad Kasiviswanathan, Leonid Ryzhyk, Mooly Sagiv, and Toby Walsh. Verifying properties of binarized deep neural networks. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *AAAI*, pages 6615–6624. AAAI Press, 2018. `doi:10.1609/AAAI.V32I1.12206`.

289 A.C. Nayak, M. Pagnucco, and P. Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146(2):193–228, 2003. `doi:10.1016/S0004-3702(03)00017-1`.

290 Mojtaba Nayyeri, Chengjin Xu, Mirza Mohtashim Alam, Jens Lehmann, and Hamed Shariat Yazdi. Logicenn: A neural based knowledge graphs embedding model with logical rules. *IEEE Transactions on Pattern Analysis of Machince Intelligence*, 45(6):7050–7062, 2023. `doi:10.1109/TPAMI.2021.3121646`.

291 Yavor Nenov, Robert Piro, Boris Motik, Ian Horrocks, Zhe Wu, and Jay Banerjee. Rdfox: A highly-scalable RDF store. In Marcelo Arenas, Óscar Corcho, Elena Simperl, Markus Strohmaier, Mathieu d'Aquin, Kavitha Srinivas, Paul Groth, Michel Dumontier, Jeff Heflin, Krishnaprasad Thirunarayan, and Steffen Staab, editors, *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference*, volume 9367 of *Lecture Notes in Computer Science*, pages 3–20. Springer, 2015. `doi:10.1007/978-3-319-25010-6_1`.

292 Allen Newell, J. C. Shaw, and Herbert A. Simon. Empirical Explorations of the Logic Theory Machine: a Case Study in Heuristic. In Morton M. Astrahan, editor, *Western Joint Computer Conference: Techniques for Reliability, IRE-AIEE-ACM 1957 (Western)*, pages 218–230. ACM, 1957. `doi:10.1145/1455567.1455605`.

293 Allen Newell and Herbert A. Simon. The Logic Theory Machine-A Complex Information Processing System. *IRE Trans. Inf. Theory*, 2(3):61–79, 1956. `doi:10.1109/TIT.1956.1056797`.

294 Ilkka Niemelä. Logic programs with stable model semantics as a constraint programming paradigm. *Annals of Mathematics and Artificial Intelligence*, 25(3-4):241–273, 1999. `doi:10.1023/A:1018930122475`.

295 Ilkka Niemelä and Patrik Simons. Smodels - An Implementation of the Stable Model and Well-Founded Semantics for Normal LP. In *Proceedings of the 4th International Conference on Logic Programming and Nonmotonic Reasoning, LPNMR 1997*, volume 1265 of *Lecture Notes in Computer Science*, pages 421–430. Springer, 1997. `doi:10.1007/3-540-63255-7_32`.

296 Nils J. Nilsson. Shakey the robot, 1984. SRI International, Menlo Park, California.

297 Nils J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28(1):71–87, 1986. `doi:10.1016/0004-3702(86)90031-7`.

298 Ahmed Nouman, Volkan Patoglu, and Esra Erdem. Hybrid conditional planning for robotic applications. *Int. J. Robotics Res.*, 40(2-3), 2021. `doi:10.1177/0278364920963783`.

299 Z. Ognjanovic and M. Raškovic. Some first-order probability logics. *Theoretical Computer Science*, 247(1–2):191–212, 2000. `doi:10.1016/S0304-3975(98)00341-7`.

300 Jukka Pajunen and Tomi Janhunen. Solution enumeration by optimality in answer set programming. *CoRR*, abs/2108.03474, 2021. `doi:10.48550/arXiv.2108.03474`.

301 Paolo Pareti and George Konstantinidis. A review of SHACL: from data validation to schema reasoning for RDF graphs. In *Reasoning Web. Declarative Artificial Intelligence - 17th International Summer School 2021*, pages 115–144, 2021. doi:10.1007/978-3-030-95481-9_6.

302 Judea Pearl. Bayesian networks: A model of self-activated memory for evidential reasoning. In *Proceedings of the 7th Conference of the Cognitive Science Society*, pages 15–17, 1985.

303 Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.

304 Judea Pearl. Graphical models for probabilistic and causal reasoning. In *Quantified Representation of Uncertainty and Imprecision*, pages 367–389. Springer, 1998.

305 Rafael Peñaloza. Axiom pinpointing. In Giuseppe Cota, Marilena Daquino, and Gian Luca Pozzato, editors, *Applications and Practices in Ontology Design, Extraction, and Reasoning*, volume 49 of *Studies on the Semantic Web*, pages 162–177. IOS Press, 2020. doi:10.3233/SSW200042.

306 Maximilian Pensel and Anni-Yasmin Turhan. Reasoning in the defeasible description logic ϵ - computing standard inferences under rational and relevant semantics. *International Journal of Approximate Reasoning*, 103:28–70, 2018. doi:10.1016/J.IJAR.2018.08.005.

307 P. Peppas. Belief revision. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 317–359. Elsevier Science, San Diego, USA, 2008. doi:10.1016/S1574-6526(07)03008-8.

308 Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. Language models as knowledge bases? In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP*, pages 2463–2473. Association for Computational Linguistics, 2019. doi:10.18653/v1/D19-1250.

309 Giulio Petrucci, Marco Rospocher, and Chiara Ghidini. Expressive ontology learning as neural machine translation. *Journal of Web Semantics*, 52-53:66–82, 2018. doi:10.1016/j.websem.2018.10.002.

310 *Potassco* project, 2011+. https://potassco.org. URL: https://github.com/potassco/.

311 Henry Prakken and Rosa Ratsma. A top-level model of case-based argumentation for explanation: Formalisation and experiments. *Argument & Computation*, 13(2):159–194, 2022. doi:10.3233/AAC-210009.

312 Marcelo O. R. Prates, Pedro H. C. Avelar, Henrique Lemos, Luís C. Lamb, and Moshe Y. Vardi. Learning to solve np-complete problems: A graph neural network for decision TSP. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI*, pages 4731–4738. AAAI Press, 2019. doi:10.1609/aaai.v33i01.33014731.

313 Anna Queralt, Alessandro Artale, Diego Calvanese, and Ernest Teniente. Ocl-lite: Finite reasoning on UML/OCL conceptual schemas. *Data Knowledge Engineering*, 73:1–22, 2012. doi:10.1016/j.datak.2011.09.004.

314 M Ross Quillian. Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities. *Behavioral Science*, 12(5):410–430, 1967.

315 L. De Raedt and K. Kersting. Probabilistic inductive logic programming. In *Probabilistic Inductive Logic Programming*, pages 1–27, 2008. doi:10.1007/978-3-540-78652-8_1.

316 L. De Raedt, A. Kimmig, and H. Toivonen. Problog: A probabilistic prolog and its application in link discovery. In *Proc. IJCAI*, pages 2462–2467, 2007. URL: http://ijcai.org/Proceedings/07/Papers/396.pdf.

317 Marco Ragni, Gabriele Kern-Isberner, Christoph Beierle, and Kai Sauerwald. Cognitive logics - features, formalisms, and challenges. In Giuseppe De Giacomo and Jérôme Lang, editors, *ECAI 2020 - 24th European Conference on Artificial Intelligence*, volume 325, pages 2931–2932. IOS Press, 2020. doi:10.3233/FAIA200459.

318 Antonio Rago, Oana Cocarascu, Christos Bechlivanidis, and Francesca Toni. Argumentation as a Framework for Interactive Explanations for Recommendations. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, pages 805–815, sep 2020. doi:10.24963/kr.2020/83.

319 Manfred Reichert and Barbara Weber. *Enabling Flexibility in Process-Aware Information Systems - Challenges, Methods, Technologies*. Springer, 2012. doi:10.1007/978-3-642-30409-5.

320 R. Reiter. On closed world data bases. In H. Gallaire and J. Minker, editors, *Logic and Databases*, pages 55–76. Plenum Press, 1978. doi:10.1007/978-1-4684-3384-5_3.

321 R. Reiter. A logic for default reasoning. *Artificial Intelligence Journal*, 13(1-2):81–132, 1980. doi:10.1016/0004-3702(80)90014-4.

322 R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, Cambridge, MA, 2001.

323 Raymond Reiter. On Closed World Data Bases. In Herve Gallaire and Jack Minker, editors, *Logic and Data Bases*, pages 55–76. Plenum Press, 1978. doi:10.1007/978-1-4684-3384-5_3.

324 Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi, editors, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144. ACM, 2016. doi:10.1145/2939672.2939778.

325 Momina Rizwan, Volkan Patoglu, and Esra Erdem. Human robot collaborative assembly planning: An answer set programming approach. *Theory and Practice of Logic Programming*, 20(6):1006–1020, 2020. doi:10.1017/S1471068420000319.

326 Mark Roberts, Laura M. Hiatt, Alexandra Coman, Dongkyu Choi, Benjamin Johnson, and David W. Aha. Actorsim, A toolkit for studying cross-disciplinary challenges in autonomy. In *2016 AAAI Fall Symposia, Arlington, Virginia, USA, November 17-19, 2016*. AAAI Press, 2016. URL: https://aaai.org/papers/14130-14130-actorsim-a-toolkit-for-studying-cross-disciplinary-challenges-in-autonomy/.

327 John Alan Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, 1965. doi:10.1145/321250.321253.

328 Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, third edition, 2010.

329 S. Sanghai, P. Domingos, and D. Weld. Relational Dynamic Bayesian Networks. *Journal of Artificial Intelligence Research*, 24:759–797, 2005. doi:10.1613/jair.1625.

330 Abulhair Saparov and He He. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. *CoRR*, abs/2210.01240, 2022. doi:10.48550/arXiv.2210.01240.

331 Roger C. Schank and Robert P. Abelson. Scripts, Plans and Knowledge. In *The 4th International Joint Conference on Artificial Intelligence, IJCAI 1975*, pages 151–157, 1975. URL: http://ijcai.org/Proceedings/75/Papers/021.pdf.

332 Klaus Schild. A correspondence theory for terminological logics: Preliminary report. In John Mylopoulos and Raymond Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 466–471, Sydney, Australia, 1991. URL: http://ijcai.org/Proceedings/91-1/Papers/072.pdf.

333 Jodi Schneider, Tudor Groza, and Alexandre Passant. A review of argumentation for the social semantic web. *Semantic Web*, 4(2):159–218, 2013. doi:10.3233/SW-2012-0073.

334 Grigori Schwarz and Miroslaw Truszczynski. Minimal knowledge problem: A new approach. *Artificial Intelligence*, 67(1):113–141, 1994. doi:10.1016/0004-3702(94)90013-2.

335 Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011. doi:10.1016/j.artint.2010.04.003.
336 Yoav Shoham. A semantical approach to nonmonotonic logics. In *Proceedings of Logics in Computer Science*, pages 275–279, Ithaca, New York, 1987.
337 Edward Shortliffe. *Computer-based Medical Consultations: MYCIN*. Addison-Wesley, 1976.
338 Evren Sirin, Bijan Parsia, Bernardo Cuenca Grau, Aditya Kalyanpur, and Yarden Katz. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, 5(2):51–53, 2007. doi:10.1016/j.websem.2007.03.004.
339 B.C. Smith. *Reflections and Semantics in a Procedural Language*. PhD thesis, MIT, 1982.
340 Lars-Phillip Spiegel, Gabriele Kern-Isberner, and Marco Ragni. Rational inference patterns. In Abhaya C. Nayak and Alok Sharma, editors, *PRICAI 2019: 16th Pacific Rim International Conference on Artificial Intelligence*, volume 11670 of *Lecture Notes in Computer Science*, pages 405–417. Springer, 2019. doi:10.1007/978-3-030-29908-8_33.
341 W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume II, pages 105–134. Kluwer Academic Publishers, 1988.
342 Andreas Steigmiller, Thorsten Liebig, and Birte Glimm. Konclude: System description. *Journal of Web Semantics*, 27:78–85, 2014. doi:10.1016/j.websem.2014.06.003.
343 Erik Strumbelj and Igor Kononenko. Explaining prediction models and individual predictions with feature contributions. *Knowledge Information Systems*, 41(3):647–665, 2014. doi:10.1007/s10115-013-0679-x.
344 Heiner Stuckenschmidt, Christine Parent, and Stefano Spaccapietra, editors. *Modular Ontologies: Concepts, Theories and Techniques for Knowledge Modularization*, volume 5445 of *Lecture Notes in Computer Science*. Springer, 2009. doi:10.1007/978-3-642-01907-4.
345 Rudi Studer, V.Richard Benjamins, and Dieter Fensel. Knowledge engineering: Principles and methods. *Data & Knowledge Engineering*, 25(1):161–197, 1998. doi:10.1016/S0169-023X(97)00056-6.
346 D. Suciu, D. Olteanu, C. Ré, and C. Koch. *Probabilistic databases*, volume 3 of *Synthesis Lectures on Data Management*. Morgan & Claypool Publishers, 2011. doi:10.2200/S00362ED1V01Y201105DTM016.
347 Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. Rotate: Knowledge graph embedding by relational rotation in complex space. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL: https://openreview.net/forum?id=HkgEQnRqYQ.
348 Tommi Syrjänen and Ilkka Niemelä. The smodels system. In *Proceedings of the 6th International Conference on Logic Programming and Nonmotonic Reasoning, LPNMR 2001*, volume 2173 of *Lecture Notes in Computer Science*, pages 434–438. Springer, 2001. doi:10.1007/3-540-45402-0_38.
349 Moritz Tenorth and Michael Beetz. Representations for robot knowledge in the knowrob framework. *Artificial Intelligence*, 247:151–169, 2017. doi:10.1016/j.artint.2015.05.010.
350 Simone Teufel, Advaith Siddharthan, and Colin Batchelor. Towards domain-independent argumentative zoning: Evidence from chemistry and computational linguistics. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 1493–1502, 2009. URL: https://aclanthology.org/D09-1155/.
351 Michael Thielscher. Ramification and causality. *Artificial Intelligence Journal*, 89(1-2):317–364, 1997. doi:10.1016/S0004-3702(96)00033-1.

352 Sebastian Thrun, Michael Beetz, Maren Bennewitz, Wolfram Burgard, Armin B. Cremers, Frank Dellaert, Dieter Fox, Dirk Hähnel, Charles R. Rosenberg, Nicholas Roy, Jamieson Schulte, and Dirk Schulz. Probabilistic algorithms and the interactive museum tour-guide robot minerva. *International Journal of Robotics Research*, 19(11):972–999, 2000. doi:10.1177/02783640022067922.

353 Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics.* Intelligent robotics and autonomous agents. MIT Press, 2005.

354 Ilaria Tiddi and Stefan Schlobach. Knowledge graphs as tools for explainable machine learning: A survey. *Artif. Intell.*, 302:103627, 2022. doi:10.1016/J.ARTINT.2021.103627.

355 David S. Touretzky. *The Mathematics of Inheritance Systems.* Morgan Kaufmann, 1986.

356 Théo Trouillon, Christopher R. Dance, Éric Gaussier, Johannes Welbl, Sebastian Riedel, and Guillaume Bouchard. Knowledge graph completion via complex tensor factorization. *Journal of Machine Learning Research*, 18:130:1–130:38, 2017. URL: http://jmlr.org/papers/v18/16-563.html.

357 Johan van Benthem. Epistemic logic and epistemology: The state of their affairs. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 128(1):49–76, 2006. URL: http://www.jstor.org/stable/4321713.

358 Wil M. P. van der Aalst. *Process Mining - Data Science in Action, Second Edition.* Springer, 2016. doi:10.1007/978-3-662-49851-4.

359 Frank van Harmelen, Vladimir Lifschitz, and Bruce W. Porter, editors. *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence.* Elsevier, 2008. URL: https://www.sciencedirect.com/science/bookseries/15746526/3.

360 Jean van Heijenoort. *From Frege to Gödel: A Source Book in Mathematical Logic, 1879-1931.* Cambridge, MA, USA: Harvard University Press, 1967.

361 Linde Vanbesien, Maurice Bruynooghe, and Marc Denecker. Analyzing semantics of aggregate answer set programming using approximation fixpoint theory. *Theory and Practice of Logic Programming*, 22(4):523–537, 2022. doi:10.1017/S1471068422000126.

362 Serena Villata and Elena Cabrio. Five years of argument mining: a data-driven analysis. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI 2018.* International Joint Conferences on Artificial Intelligence Organization, 2018. doi:10.24963/ijcai.2018/766.

363 Denny Vrandečić and Markus Krötzsch. Wikidata: A free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85, sep 2014. doi:10.1145/2629489.

364 W3C OWL Working Group. OWL 2 web ontology language document overview. URL, 2009. URL: http://www.w3.org/TR/owl2-overview/.

365 Przemyslaw Andrzej Walega, Michal Zawidzki, and Bernardo Cuenca Grau. Finite materialisability of datalog programs with metric temporal operators. *Journal of Artificial Intelligence Research*, 76, 2023. doi:10.1613/jair.1.14040.

366 Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29:2724–2743, 2017. doi:10.1109/TKDE.2017.2754499.

367 *Wasp* project, 2013+. URL: http://alviano.github.io/wasp/.

368 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *CoRR*, abs/2201.11903, 2022. arXiv:2201.11903, doi:10.48550/arXiv.2201.11903.

369 Antonius Weinzierl, Richard Taupe, and Gerhard Friedrich. Advancing lazy-grounding ASP solving techniques - restarts, phase saving, heuristics, and more. *Theory and Practice of Logic Programming*, 20(5):609–624, 2020. doi:10.1017/S1471068420000332.

370 Alfred North Whitehead and Bertrand Arthur William Russell. *Principia Mathematica; 2nd ed.* Cambridge University Press, Cambridge, 1927. URL: https://cds.cern.ch/record/268025.

371 Sarah Wiegreffe and Ana Marasovic. Teach me to explain: A review of datasets for explainable natural language processing. In Joaquin Vanschoren and Sai-Kit Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, virtual*, 2021. URL: https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/698d51a19d8a121ce581499d7b701668-Abstract-round1.html.

372 Hong Wu, Zhe Wang, Kewen Wang, and Yi-Dong Shen. Learning Typed Rules over Knowledge Graphs. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*, pages 494–503, aug 2022. doi:10.24963/kr.2022/51.

373 Guohui Xiao, Diego Calvanese, Roman Kontchakov, Domenico Lembo, Antonella Poggi, Riccardo Rosati, and Michael Zakharyaschev. Ontology-based data access: A survey. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, pages 5511–5519. ijcai.org, 2018. doi:10.24963/ijcai.2018/777.

374 Guohui Xiao, Davide Lanti, Roman Kontchakov, Sarah Komla-Ebri, Elem Güzel Kalayci, Linfang Ding, Julien Corman, Benjamin Cogrel, Diego Calvanese, and Elena Botoeva. The virtual knowledge graph system ontop. In Jeff Z. Pan, Valentina A. M. Tamma, Claudia d'Amato, Krzysztof Janowicz, Bo Fu, Axel Polleres, Oshani Seneviratne, and Lalana Kagal, editors, *The Semantic Web - ISWC 2020 - 19th International Semantic Web Conference*, volume 12507 of *Lecture Notes in Computer Science*, pages 259–277. Springer, 2020. doi:10.1007/978-3-030-62466-8_17.

375 Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. A semantic loss function for deep learning with symbolic knowledge. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 5498–5507. PMLR, 2018. URL: http://proceedings.mlr.press/v80/xu18h.html.

376 Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015*, 2015. doi:10.48550/arXiv.1412.6575.

377 Zhun Yang, Adam Ishay, and Joohyung Lee. NeurASP: Embracing neural networks into answer set programming. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI*, pages 1755–1762, 2020. doi:10.24963/ijcai.2020/243.

378 Safdar Zaman, Gerald Steinbauer, Johannes Maurer, Peter Lepej, and Suzana Uran. An integrated model-based diagnosis and repair architecture for ros-based robot systems. In *2013 IEEE International Conference on Robotics and Automation*, pages 482–489, 2013. doi:10.1109/ICRA.2013.6630618.

379 Benjamin Zarrieß. Complexity of projection with stochastic actions in a probabilistic description logic. In Michael Thielscher, Francesca Toni, and Frank Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018*, pages 514–523. AAAI Press, 2018. URL: https://aaai.org/ocs/index.php/KR/KR18/paper/view/18050.

380 Rowan Zellers, Jiasen Lu, Ximing Lu, Youngjae Yu, Yanpeng Zhao, Mohammadreza Salehi, Aditya Kusupati, Jack Hessel, Ali Farhadi, and Yejin Choi. MERLOT RESERVE: neural script knowledge through vision and language and sound. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022*, pages 16354–16366. IEEE, 2022. doi:10.1109/CVPR52688.2022.01589.

381 Maximilian Zellner, Ali E Abbas, David V Budescu, and Aram Galstyan. A survey of human judgement and quantitative forecasting methods. *Royal Society Open Science*, 8(2):201187, 2021.
382 Honghua Zhang, Liunian Harold Li, Tao Meng, Kai-Wei Chang, and Guy Van den Broeck. On the paradox of learning to reason from data. *CoRR*, abs/2205.11502, 2022. doi:10.48550/arXiv.2205.11502.
383 Zheng Zhang, Levent Yilmaz, and Bo Liu. A critical review of inductive logic programming techniques for explainable AI. *CoRR*, abs/2112.15319, 2021. doi:10.48550/arXiv.2112.15319.
384 Yizheng Zhao and Renate A. Schmidt. FAME: an automated tool for semantic forgetting in expressive description logics. In Didier Galmiche, Stephan Schulz, and Roberto Sebastiani, editors, *Automated Reasoning - 9th International Joint Conference, IJCAR*, volume 10900 of *Lecture Notes in Computer Science*, pages 19–27. Springer, 2018. doi:10.1007/978-3-319-94205-6_2.
385 Zhiqiang Zhuang, Zhe Wang, Kewen Wang, and Guilin Qi. DL-Lite contraction and revision. *Journal of Artificial Intelligence Research*, 56:328–378, 2016. doi:10.1613/jair.5050.

A Appendix: History of Knowledge Representation and Reasoning

Today it is broadly acknowledged that knowledge is a fundamental aspect of intelligence. This was arguably first stated by McCarthy and Hayes in their seminal paper "Some Philosophical Problems from the Standpoint of AI [268]. They wrote there: "*[...] intelligence has two parts, which we shall call the* epistemological *and the* heuristic*. The epistemological part is the representation of the world in such a form that the solution of problems follows from the facts expressed in the representation. The heuristic part is the mechanism that on the basis of the information solves the problem and decides what to do.*" The *epistemological* component corresponds to what we understand today as the problem of *representation*, and the *heuristic* part is the *reasoning* problem. This stance was later restated by Brian Smith [339] in the *Knowledge Representation Hypothesis*, discussed in the introduction.

It seems appropriate then to take the year 1969, the year when the paper by McCarthy and Hayes was published, as marking the birth of KR. That paper not only pinpointed the role of KR in AI but also identified many of the key concerns of this area, all of which remain pertinent today. They are (quoting from the paper):

1. *What kind of general representation of the world will allow the incorporation of specific observations and new scientific laws as they are discovered?*
2. *Besides the representation of the physical world what other kinds of entities have to be provided for? For example, mathematical systems, goals, states of knowledge.*
3. *How are observations to be used to get knowledge about the world, and how are the other kinds of knowledge to be obtained? In particular what kinds of knowledge about the system's own state of mind are to be provided for?*
4. *In what kind of internal notation is the system's knowledge to be expressed?*

While some earlier developments such as Newell and Simon's *General Problem Solver* [293, 292], and Robinson's *Resolution Principle* anticipated important KR concerns, it was the work by McCarthy and Hayes that first explicitly stated the challenges defining the field.[17] That said, research in KR as a mature area of AI is commonly taken as being marked by an

[17] Other early work in the area is summarised in Brachman and Smith's 1980 survey [61] of KR.

Artificial Intelligence Journal Special Issue on Non-monotonic Reasoning in 1980 [46]. In 1989 the Principles of Knowledge Representation and Reasoning Conference was founded, providing a dedicated, specialised forum for research in the area.[18]

Over the years KR has grown into a mature field of cross-disciplinary ideas, with theoretical results and implemented systems inspired by foundational concerns, as well as by practical considerations stemming from the needs of concrete applications. The primary focus was on developing languages for knowledge representation based on first-order logic, sometimes extended by non-logical connectives and modal operators and, in some cases, assigned a non-classical semantics. However, driven to a large degree by the need for efficient reasoning, important early alternative approaches to knowledge representation emerged as well, most notably frames [277], scripts [331], production rule systems, and graphical formalisms such as conceptual graphs, inheritance networks, semantic networks, and Bayesian networks, the latter to model uncertainty. Today, knowledge representation languages rooted in these early proposals form the core of KR.

A.1 Logic and KR

The choice of logic as a formalism for knowledge representation was not arbitrary. We have already mentioned McCarthy and Hayes's 1969 paper; as well a series of papers by McCarthy [265, 266] and Hayes [184, 185, 187, 186] implicitly or explicitly argued for this position. A large part of knowledge is declarative and consists of statements describing entities and relationships between them. And incorporating new knowledge is often as simple as adding a new set of statements to the existing one. The language of logic, first-order logic in particular, is well suited for modelling declarative knowledge. Statements describing what is known can be modelled as logical formulas. The Tarskian semantics of first-order logic provides a connection between a collection of formulas – a body of knowledge about an application domain – and structures, which could be interpreted as abstractions of that application domain with all its relevant entities and relationships. First-order logic is modular and its statements can be arranged in hierarchies. Further, the concept of a formal proof in first-order logic offers a way to derive new formulas from those given or established earlier. And logic already proved its mettle helping establish formal foundations for the rich and rigor-demanding edifice of mathematics.

Not surprisingly then, the influence of first-order logic can be seen in the *Logic Theorist* by Newell, Simon and Shaw [293, 292], sometimes referred to as the first AI program, and the specification of the *Advice Taker* by McCarthy [269]. The discovery of the resolution proof method by Alan Robinson [327] and subsequent development of automated reasoning further strengthened the position of logic as the formalism of choice for KR. That being said, the early work on logic-based approaches to KR already pointed to limitations of first-order logic. First, as it was already known since Gödel, first-order logic is undecidable. Moreover, even under restrictions on the language that result in decidability, determining provability (or entailment) is often inherently hard. Logic-based AI programs could serve as proofs of concept or as a theoretical specification, but not as deployed systems working under time constraints. Second, some aspects of knowledge proved stubbornly difficult to formalise in first-order logic. In particular, commonsense reasoning, which McCarthy identified early on as a critical capability of intelligent systems [269], turned out to be a source of several

[18] `https://kr.org/pastconfs.html`

challenging problems hard to capture within the first-order framework. Third, as argued by several researchers, Minsky most notable among them, logic and its automated reasoning machinery simply was not a good model of how the human brain works.

A.2 Objections to First-Order Logic as a KR Language

We will now consider these objections to first-order logic as KR language in more detail, as each led to major research directions within KR.

A.2.1 Undecidability

The undecidability of first-order logic naturally pushed researchers to look for fragments of first-order logic in which effective reasoning was possible.

Propositional Logic. Propositional logic is one such fragment. For its most commonly considered clausal version, resolution refutation was proved to be complete and sound. However, the problem of deciding whether a proof exists is coNP-complete while the dual problem of satisfiability is NP-complete. Thus, there is little hope for the existence of fast (polynomial-time bounded) algorithms for this problem. Nevertheless, because of impressive advances made in propositional satisfiability, the role of propositional logic in KR has been growing. It has left a strong mark among others in diagnosis, abductive reasoning, belief revision and update, as well as in planning. It has suggested or directly provided reasoning algorithms for answer-set programming, and spawned generalisations such as constraint satisfaction.

Production Rules. Production rules are another formalism proposed for KR that is closely related to a fragment of first-order logic known as Horn logic. Because of their restricted form, reasoning with production rules can be highly efficient. They were used as the basis of expert systems, one of the most successful class of AI programs built in the first golden era of AI in the 1970s and 1980s [115]. This line of research started with expert systems such as MYCIN [337, 70], DENDRAL [71, 69], PROSPECTOR [129], and R1 (Xcon) [274], and evolved into building general purpose expert system shells. Important outcomes of this research include identifying knowledge acquisition as a major bottleneck and a challenge for KR, and establishing knowledge engineering as a subarea of KR.

Production rules emphasised the importance of the Horn fragment of first-order logic. Specialised resolution proof methods such as the SLD resolution [226] made it possible to think about logic as a declarative programming language [225]. On the practical side, it led to the development of programming languages such as PLANNER [190] and PROLOG [102]. On the theoretical side it opened the key question of the semantics of programs (especially programs with negation, which was allowed by the designers of those languages but only given a procedural interpretation), and helped establish the field of logic programming. Originally, logic programming and KR were developing side-by-side with few interactions between the two communities. Eventually, however, strong connections between the two fields were discovered. One of the outcomes of this coming together is answer set programming [65], one of the most successful computational KR systems.

Description Logics. Another line of research emerged from the realization of a trade-off between the efficiency of reasoning and expressiveness of the language [58]. It gave rise to an area in KR known as description logics [20, 23]. Its origins can be traced to graphical models such as semantic networks [314] and structured languages such as KL-ONE [62].

However, it truly flourished once these earlier efforts were mapped onto fragments of first-order logic. This led to a careful and systematic analysis of the complexity and expressiveness and resulted in a broad range of description logics differing from each other in these respects [20, 22]. A major push behind the development of description logics came from the needs of the semantic web – the world wide web with its content converted to machine interpretable formats. It resulted in the semantic web languages DAML, DAML+OIL [198] and OWL [275]. Another important influence of description logics can be found in the areas of databases and information systems.

KR Languages for Application Domains. Finally, we note that much of the research in KR was driven by specialised application domains. One of the most influential ones concerns modelling and reasoning about systems that evolve as a result of actions. It was identified as an area of interest in AI by McCarthy in a technical report published by the Stanford University in 1963 [270]. In that report McCarthy introduced the *situation calculus*, a first-order logic language designed to model actions and their effects. This language was further elaborated by McCarthy and Hayes [268] and developed into a full theory by Reiter [322]. The first two papers also introduced the key concerns to be addressed: in particular, the famous *frame problem* as well as two other fundamental problems for KR, the *ramification problem* and the *qualification problem*. Addressing these problems, especially the frame problem, constituted major milestones in the development of KR. Elegant solutions were proposed not only for the original framework of the situation calculus but also in a related *event calculus* [227], and in non-classical formalisms for reasoning about action based on logic programming with the stable model semantics [157].

Systems that change as actions are executed subsume planning problems. Automated planning is as old as AI and some of the most influential early AI work concerns planning (for example, the STRIPS language [144] and the robot Shakey [296, 233]). Planning in general, and planning to address the needs of autonomous agents (robots), has been one of the core areas of AI. Many concerns in planning are those of representing knowledge about the world, and reasoning about that world as it changes. Hence, there are significant overlaps between KR and planning. Nevertheless, today planning and KR are two separate areas of AI.

A.2.2 Addressing the Commonsense Reasoning Challenge

The second major objection to first-order logic as a KR formalism stems from difficulties researchers discovered when attempting to use first-order logic to formalise commonsense reasoning. Humans can naturally handle rules with exceptions, reason without full information, represent and use normative statements, revise or update their knowledge (belief) base when new information becomes available, they can reason qualitatively about physical phenomena in the context of space and time, and they can represent and reason about uncertainty. Not surprisingly, from the outset and continuing to today, researchers have argued that formalising and automating common sense is necessary for artificial intelligence. Yet, it is not possible to do so strictly within first-order logic. This realization motivated several subareas of KR. We have already summarised current research and challenges in nonmonotonic reasoning and belief revision; below we explore these areas, and their history, in a bit more detail.

The types of reasoning mentioned above involve various non-logical reasoning principles such as adopting the view that the world is, in some fashion, "normal" or making assumptions based on the likelihood of a property or event. These non-logical principles make it possible that deduced properties may become invalid in light of new information. This is clearly not the case for inference supported by first-order logic. The need then arose for formalising such non-logical principles into an inference mechanism in a principled fashion.

The first such principle was the *closed world assumption* identified and studied by Reiter [323]. Further developments came in 1980, when McCarthy introduced *circumscription* [267], which incorporated the principle of reasoning in the context of minimal models only, Reiter introduced default logic [321], which integrated non-classical inference rules (defaults) conditioned upon what could reasonably be assumed as possible, while McDermott and Doyle introduced a modal logic allowing for introspection [273]. This last approach was found lacking. However, soon thereafter improved proposals for non-standard reasoning with modal theories were made by McDermott [272], who introduced a single fixpoint principle to capture non-classical inferences for all normal modal logics, and Moore, who proposed *autoepistemic logic* [283]. Circumscription, default logic and modal logics of McDermott, Doyle and Moore gave rise to the area of non-monotonic logics that soon became one of the main areas in KR. A significant research effort expanded in that area explained the relationships between different non-monotonic logics, and established a comprehensive picture of their expressive power and complexity [262]. It also demonstrated the effectiveness of non-monotonic logics in modelling abductive and diagnostic reasoning tasks [118]. An important contribution of default logic is the semantics it suggested for logic programs with negation. That semantics, known as the stable-model semantics, has been broadly accepted and is the foundation of answer-set programming mentioned above.

A different attempt to address non-monotonicity in reasoning was provided by Kraus, Lehmann, and Magidor [229]. In what has become known as the *KLM approach* to non-monotonic reasoning, the authors proposed sets of axioms against which any non-monotonic consequence relation can be tested. Although this work was initially focused on non-monotonic consequence for classical propositional logic, it was subsequently extended to study non-monotonic consequence over conditional propositional logics [240]. This extension has given rise to a flurry of activity in designing non-monotonic reasoning systems for versions of conditional logics going beyond propositional logic.

While logic was the most extensively studied framework to capture non-monotonic reasoning, a significant effort was expanded in 1980s to extend semantic networks to handle exceptions and defaults. This research effort evolved into the area of inheritance networks [355]. As inheritance networks could be mapped into (non-monotonic) logics, this line of research gradually waned. It may receive soon renewed attention with the growing role of knowledge graphs [207] in KR.

Human commonsense reasoning allows us to gracefully update or revise our knowledge and beliefs once new information becomes available. Formalising this capability is essential for intelligent behavior in a changing world, when a knowledge (belief) base has to be updated to reflect a new situation, or revised, when the situation does not change but "better" (more accurate) information is gained. The field that grew out of these considerations is called belief revision but it includes also considerations that are more accurately referred to as belief update. The major milestone for the area is the paper by Alchourrón, Gärdenfors and Makinson [6]. That paper proposed a set of postulates, now known as the AGM postulates, that any belief revision operator should satisfy. It was a way to address the unsettling realization that belief revision could not be reduced to a single revision operator. In fact, numerous operators were proposed that, arguably, adequately handled the task of revising a knowledge base by a formula representing a new piece of information. The paper was followed by papers characterizing, often in a constructive way, classes of operators satisfying all (or some) of the AGM postulates. In another interesting development, Gärdenfors and Makinson showed that KLM-style non-monotonic reasoning and AGM belief revision are two sides of the same coin [151]. Informally, the link between the two formalisms is easy to explain: B is a non-monotonic consequence of A whenever B is a classical consequence of the knowledge base obtained from a revision of A.

Additional concerns addressed or partially addressed over the years included iterated revision, a major extension of the original framework that focused on a single act of revising [341, 111], and a more general problem of knowledge merging [224]. The field of belief update was identified as separate from belief revision by Katsuno and Mendelzon [216], who started it by proposing a set of postulates for belief update operators. In many respects its development followed that of belief revision.

Qualitative reasoning is another type of commonsense reasoning humans are good at. It concerns reasoning about "continuous" domains such as space and time. While phenomena in such domains are most effectively modelled by systems of (partial) differential equations, few of us maintain such systems in our mind and solve them when the need arises. Instead, we have developed "commonsense" representations of such domains that allow us to reason efficiently about time, space, motion and actions that might affect them, and with sufficient accuracy to function [146, 112]. The main thrust of this work is to develop ontologies of relevant concepts and develop reasoning techniques based on these ontologies. In this way, the area can be mapped onto (a fragment) of first-order logic and can be developed exploiting logical means. Two important subareas of qualitative reasoning focus on spatial [101] and temporal [145] reasoning, respectively.

A.2.3 Beyond Logic

Much of what we discussed falls firmly in the domain of qualitative approaches based on the language of logic. However, quantitative approaches also played a major role in KR research. One of the most prominent examples is the formalism of *Bayesian networks* [302, 303, 110]. A Bayesian network consists of a directed acyclic graph representing variables of the problem as nodes and direct influences among variables as edges, and of a collection of conditional probability tables associated with each variable. The power of Bayesian networks comes from an observation that valid inferences about probabilities of events or contributions of certain events to cause events that occurred can be derived without constructing explicitly joint distribution functions on all events involved. To make Bayesian networks an effective KR tool, researchers developed methods to represent Bayesian networks, to compile them into arithmetic circuits and propositional formulas, as well as algorithms to perform exact and approximate inference. Another key research theme concerned methods to construct Bayesian networks by standard knowledge engineering methods involving domain experts, by deriving them from specifications exploring their structural regularity by learning. Today, Bayesian networks form a dominant tool for reasoning with uncertain and incomplete information. An important aspect of Bayesian networks has been their ability to model causality.

Finally, we note that not all knowledge is declarative. Much of what we know is procedural or algorithmic, and stored as such procedures in our brains. Using this knowledge involves little if any reasoning, just recalling from memory the right procedure for the occasion. Knowledge needed to accomplish routine, frequently repeated tasks, in particular, scene and speech understanding fall into this category. The ways human brains act to accomplish these tasks are not fully understood. However, they inspired the so called *connectionist* models, best known as neural networks, as a possible framework for not only representing algorithmic knowledge but also for learning these representations from observations (examples). This approach has been outside of the scope of KR for much of its history. This is beginning to change. On the one hand, machine learning and neural network representations form that part of AI that admittedly ignited imagination of both research communities and the community at large. It is now by far the most vibrant branch of AI dwarfing all others. On the other hand, it is not by itself sufficient to reach the point when machines will act intelligently in ways humans do. For that these systems need knowledge and it is now evident that further progress in AI depends on successful integration of machine learning and knowledge representation.