

Guest Editors' Introduction

James Delgrande · Jérôme Lang

Received: 21 August 2014 / Accepted: 21 August 2014 / Published online: 18 September 2014
© Springer Science+Business Media Dordrecht 2014

This special issue presents a selection of papers in Knowledge Representation (KR) in Artificial Intelligence (AI), intended to illustrate the depth and breadth of current research in the area. It comes just over 25 years since a similar special issue of the *Journal of Philosophical Logic* appeared on the topic *Philosophical Logic and Artificial Intelligence* [15]. This latter special issue covered work addressing the use of logic, in one form or another, for representing and reasoning with knowledge. The papers of the 1988 special issue give a good indication of major themes of research at the time: reasoning about belief and knowledge [5], tractable logical reasoning [10], nonmonotonicity [11], and reasoning about action [13]. These topics can be considered to be centrally in KR; the current issue then also serves to indicate how the field has broadened and evolved over the last 25 years.

Traditionally the area of Knowledge Representation has been seen as “the field of study concerned with using formal symbols to represent a collection of propositions believed by some putative agent” [3]. (A further, key, aspect of KR is *reasoning* since, after all, KR is an area of Artificial Intelligence which in turn is concerned with the implementation of intelligent artifacts on a computer.) However, KR has grown to the extent that this traditional characterisation is overly limiting. A more general definition, which better fits the current state of the art, as reflected by the coverage of

J. Delgrande (✉)

School of Computing Science, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada
e-mail: jim@cs.sfu.ca

J. Lang

LAMSADE – CNRS and Universit Paris Dauphine, Paris, France
e-mail: lang@lamsade.dauphine.fr

the *Knowledge Representation and Reasoning* (KR) conferences, also considers the representation of the mental attitudes of an agent, including knowledge and belief, but also preferences and intentions. Hence KR is more appropriately seen as dealing with (symbolic, declarative) representations concerning the mental attitudes of an agent.

There are two questions regarding KR that we will review, first “*what is KR?*” and second (and of equal importance) “*why bother with KR?*”. The traditional answer to both these questions is given by what has been called the *knowledge representation hypothesis* ([14], page 2):

“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 a *declarative representation* of information or an agent’s *knowledge base*, or more generally, a representation of an agent’s *mental attitudes*, which bears a (semantic) relation to the domain in which the agent is located. Moreover, reasoning over these mental attitudes will play an essential, causal role in determining that agent’s behaviour.

AI is broadly concerned with the design and implementation of agents that exhibit intelligence to some degree or another. What distinguishes KR is that the relevant knowledge is explicitly given in a declarative form in its knowledge base. This can be compared, say, with a neural network, where a network may be said to be knowledgeable; however in this case the knowledge is *implicit*, whereas in a KR-based approach it is *explicit*. It is an open question as to how much of AI is, or needs to be, knowledge-based. Some areas, such as planning, diagnosis, or commonsense reasoning, are strongly knowledge-based. Others, such as motor control, and low-level vision and speech understanding are largely not. And yet other approaches may be a combination: for example in a game-playing program, the function that evaluates a static board configuration may be strongly declarative, whereas the game-playing strategy may be implicitly encoded in move-generating procedures. However, while it is an open question as to how much of intelligent behaviour is (or needs to be) knowledge-based, what is clear is that much of AI involves the creation of systems which reason over an explicit body of knowledge.

In a related but distinct role, KR is also concerned with the formalisation and analysis of (real-world, often commonsense) phenomena of interest. In constructing an artificial agent, one needs to deal with a wide range of phenomena such as reasoning about the effects of actions, revising one’s beliefs on learning new information, reasoning with typicality statements such as “birds fly”, or reasoning about other agents’ beliefs, preferences, or intentions. In such cases, a precise theory of such phenomena is required, and in such cases the appropriate tool, often by definition, is *logic*. “Logic” has to be understood here in a broad sense, as “formal methods allowing one to perform inferences”. So, on the one hand, logical theories have been proposed

in areas such as reasoning about action, reasoning about knowledge, nonmonotonic reasoning, and numerous others. On the other hand, formalisms such as Bayesian networks, conditional preference networks, or constraint satisfaction problems, which belong equally in KR, are perhaps not logics in the narrow sense of the word, but they all comprise a formal language in which declarative knowledge can be expressed in a compact way, a semantics, and they allow one to perform inferences.

KR has been a distinct subarea of AI for roughly 35 years. If one wanted to specify a year, 1980 would be a good candidate: In 1980, a special issue of the *Artificial Intelligence Journal* [2] introduced several of the major approaches to nonmonotonic reasoning; as well in that year, a special issue of the SIGACT Newsletter [4] summarised research in KR to that point. By 1989 the field had matured to the extent that the first *International Conference on the Principles of Knowledge Representation and Reasoning* was held in Toronto; this conference remains the central specialised forum for research in the area. More recently, the *Handbook of Knowledge Representation* [16] provides a both broad and in-depth compendium and summary of research in the area.

This special issue then cannot hope to be comprehensive. Rather, our goal is to illustrate how the field has grown and evolved by presenting a selection of papers from diverse areas in KR. To this end, and to ensure breadth and representation of different areas, we invited submissions from distinguished researchers for this special issue, five of which appear here. Apart from being invited, all papers were handled as regular submissions: Each paper submitted for this collection was reviewed by three referees, each of whom was an expert on the topic of the paper; several papers were substantially revised in response to the comments made by the reviewers.

In the first paper in this collection, Franz Baader, Stefan Borgwardt, and Rafael Peñaloza examine decidability issues in combining fuzzy logic with the description logic \mathcal{ALC} . The area of *description logic* [1] addresses, roughly, the structure of concepts in terms of relations including subsumption and role relations; it has grown very significantly over the last two decades, becoming one of the major areas of research in KR. Similarly fuzzy logic has found application in dealing with imprecise and vague concepts. In this paper, the authors consider a central description logic, namely \mathcal{ALC} , and how it may be combined with fuzzy logic, parameterized by an arbitrary continuous triangular norm. They present a tableau algorithm for the basic combination but also show that in the presence of so-called general inclusion axioms the algorithm does not yield a decision procedure for consistency, and in fact show that, at least for the choice of the product t-norm, the problem becomes undecidable.

In the second paper, Jianmin Ji and Fangzhen Lin present what they call *position systems* for reasoning in dynamic domains. Historically there has been extensive research in reasoning about action; however most such work has focussed on the effects of actions, that is, on specifying a transition function, while constraints on the set of legal states of the world are only implicitly defined. In the present paper, the focus is on specifying the space of legal states in a transition system framework. In so doing, various planning domains are considered as examples.

For the third paper, we are very fortunate to have a contribution from Vladimir Lifschitz, who was also a contributor to the 1988 special issue. His paper “The Dramatic True Story of the Frame Default” is an expository article that takes as its point

of departure the solution to the frame problem proposed by Raymond Reiter in [12]. The frame problem is the problem of how to compactly and in a principled fashion state, after the execution of an action, what conditions are unaffected by the action. Reiter's solution involved expressing this condition in Default Logic. In his article, Lifschitz uses this idea to trace the evolution and development of several key notions in KR, including *abnormality minimization*, *explanation closure*, *negation as failure*, and culminating with one of the major success stories of KR, the *answer set programming* approach.

The paper by Matthias Westphal, Stefan Wölfl, Bernhard Nebel, and Jochen Renz deals with *qualitative reasoning* [6], in particular qualitative route descriptions. In their approach, agent models interpreting such descriptions are formalised using propositional dynamic logic. The authors discuss how information contained in map-like data can be used to obtain an agent-oriented qualitative representation. As well, they address issues of ambiguities in route descriptions, in which there are several alternatives to the agents processing them. Different types of agents are distinguished, are formalized using propositional dynamic logic, and their computational complexity analyzed.

Last, Dongmo Zhang and Michael Thielscher develop and analyse a formal language for representing and reasoning about game strategies. This is carried out within an extension to GDL, the *Game Description Language* [8], a language for encoding rules of arbitrary games that can then be interpreted by general game-playing systems. They give extensions to GDL to allow the description of strategies, where a strategy is understood as a set of move recommendations. Implementations are given in terms of the *situation calculus* [9], a first-order theory for describing actions and their effects, and *answer set programming* [7].

The papers in this collection involve a wide range of areas in KR, including description logic, plausible reasoning, reasoning about action, nonmonotonic reasoning, qualitative reasoning, and reasoning in multi-agent environments. Notably, in 1988, at the time of the earlier special issue, all of these areas were either in their infancy, or else have evolved radically and in unpredictable directions. Noteworthy too is the fact that, whereas work in 1988 was confined largely to theoretical development and analysis, KR has matured to the extent that in several areas it is now being practically applied, and realistic implementations are being used to test and evaluate the underlying logical theories.

To conclude, we would like to thank the authors for their contributions, and we thank the reviewers for their excellent work. We would also like to thank Hans van Ditmarsch and the other editors of the *Journal of Philosophical Logic* for their encouragement and support for this project.

James Delgrande
Jérôme Lang

References

1. F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider (Eds.) (2007). *The Description Logic Handbook*, 2nd edn. Cambridge: Cambridge University Press.
2. Bobrow, D. (ed.) (1980). Special Issue on Nonmonotonic Reasoning. *13*(1-2) Elsevier.
3. Brachman, R.J., & Levesque, H.J. (2004). In *Knowledge Representation and Reasoning*. San Mateo: Morgan Kaufmann.
4. Brachman, R.J., & Smith, B.C. (1980). Special issue on knowledge representation. *SIGART Newsletter*, 70.
5. Fagin, R., & J.Y. Halpern. (1988). I'm ok if you're ok: On the notion of trusting communication. *Journal of Philosophical Logic*, *17*(4), 329–354.
6. Forbus, K. (2008). Qualitative modelling. In: F. van Harmelen, V. Lifschitz, B. Porter (Eds.), *Handbook of Knowledge Representation* (pp. 361–394). San Diego: Elsevier.
7. Gelfond, M., & Lifschitz, V. (1988). The stable model semantics for logic programming. In: Proceedings of the Fifth International Conference and Symposium of Logic Programming (ICLP'88) (pp. 1070–1080). Cambridge: MIT Press.
8. Genesereth, M., & Love, N. (2005). General game playing: Overview of the AAAI competition. *AI Magazine*, *26*, 62–72.
9. Levesque, H.J., Pirri, F., Reiter, R. (1998). Foundations for the situation calculus. *Linköping Electronic Articles in Computer and Information Science*, *3*(18).
10. Levesque, H.J. (1988). Logic and the complexity of reasoning. *Journal of Philosophical Logic*, *17*(4), 355–389.
11. Lifschitz, V. (1988). Circumscriptive theories: A logic-based framework for knowledge representation. *Journal of Philosophical Logic*, *17*(4), 391–441.
12. Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, *13*(1-2), 81–132.
13. Yoav Shoham (1988). Efficient reasoning about rich temporal domains. *Journal of Philosophical Logic*, *17*(4), 443–474.
14. Smith, B.C. (1982). Reflections and semantics in a procedural language. In: PhD thesis. Cambridge: MIT.
15. R. Thomason (ed.) (1988). Special Issue on Philosophical Logic and Artificial Intelligence. *17*(4) Kluwer Academic Publishers.
16. Harmelen, v.an.F., Lifschitz, V., Porter, B. (2008). In *Handbook of Knowledge Representation*. San Diego: Elsevier.