

Preface to Special Issue on Commonsense Reasoning

Andrew S. Gordon, University of Southern California, gordon@ict.usc.edu

Rob Miller, University College London, r.s.miller@ucl.ac.uk

Leora Morgenstern, Palo Alto Research Center, leora@parc.com

Gyorgy Turan, University of Illinois at Chicago and University of Szeged, gyt@uic.edu

A few years after the 1956 Dartmouth Summer Workshop (McCarthy et al., 1955; Moore, 2006), which first established artificial intelligence as a field of research, John McCarthy (1959) discussed the importance of explicitly representing and reasoning with commonsense knowledge to the enterprise of creating artificially intelligent robots and agents. McCarthy proposed that commonsense knowledge was best represented using formal logic, which he viewed as a uniquely powerful lingua franca that could be used to express and reason with virtually any sort of information that humans might reason with when problem solving. a stance he further developed and propounded in (McCarthy, 1963; McCarthy and Hayes, 1969). This approach, the *formalist* or *logic-based* approach to commonsense reasoning, was practiced by an increasing set of adherents over the next several decades (Lenat et al., 1986; Davis, 1990), and continues to be represented by the Commonsense Symposium Series, first held in 1991 (Lifschitz et al., 1991) and held biennially, for the most part, after that.

The commonsense reasoning landscape has changed considerably over the years. More than thirty years ago, Drew McDermott (1987) noted that correctly specifying commonsense knowledge within formal logic was an error-prone enterprise, pointing to the brittleness of existing formal theories of common sense and the difficulty of modifying such theories in the face of new information. Statistical approaches to AI have gained traction as they have proved to be successful at important tasks for speech recognition and natural language understanding (Manning, 1999; Jurafsky and Martin, 2014). Over the last decade, neural network techniques in particular have shown great promise in significantly increasing the level of performance not only on such tasks but also on previously unsolvable problems such as image recognition (Krizhevsky et al., 2017) and those involving complex planning tasks, such as playing Go (Silver et al., 2016). More recently, neural nets have succeeded beyond expectations on challenges such as the Winograd Schema Challenge

(Levesque et al., 2012), which was specifically designed to assess whether a system has commonsense knowledge and can perform commonsense reasoning. This success has been achieved despite the lack of large formalized commonsense knowledge bases or demonstrated ability of software systems to perform commonsense reasoning (Kocijan et al., 2020).

Yet neural network techniques are not a panacea. They are vulnerable to adversarial attacks (Bengio et al., 2017), magnify biases in training data, and are often prone to the sorts of errors that would be immediately obvious to a human and that indicate a serious lack of common sense (Marcus and Davis, 2019). Formal approaches to commonsense knowledge and reasoning, according to this view, will remain crucial to the endeavor to build artificial intelligences, perhaps as parts of *neurosymbolic* (Kautz 2020) systems that synthesize neural network approaches with formal approaches. DARPA is one example of a funding agency that is currently investing tens of millions of dollars in an attempt to build such systems, in such programs as Explainable AI (Gunning et al., 2019) and Machine Common Sense (Gunning, 2018).

In light of these developments, the Commonsense Symposium Series is currently being revamped to reflect the importance of the synthesis of neural and symbolic approaches. The most recent symposium in the series, held in November 2017, began to look forward to this synthesis, and highlighted invited talks by Sebastian Riedl and Murray Shanahan on their approaches to this endeavor. Riedel discussed how neural network algorithms on vector spaces could be used to model “soft” inference from assumptions, in many cases approximating a theorem prover, but with much greater efficiency. Shanahan presented methods to integrate deep learning and reinforcement learning for learning classical, formal relations and predicates, a precursor to the work described in (Gamelo and Shanahan, 2019).

The papers contained in this volume are for the most part extensions of a subset of papers presented at the 2017 meeting of the Commonsense Symposium Series (Gordon et al., 2018). They are briefly summarized below.

In their paper *An Investigation of Parametrized Difference Revision Operators*, Theofanis Aravanis, Pavlos Peppas, and Mary-Anne Williams discuss belief revision, the area dealing with changing knowledge bases. The basic approach in this field is to identify properties that should be satisfied by knowledge change operators, to give general characterizations of operators satisfying the properties, and then to find classes of operators with those properties, and, in addition, having favorable computational properties. Their paper considers parametrized difference (PD) revision operators, a class introduced recently by Peppas and Williams. These

operators extend Dalal's operator, one of the basic early proposals, by allowing different epistemic values for the variables. New characterizations are given for PD operators, and it is shown that beyond the standard AGM postulates, these operators also satisfy the strong form of the relevance postulate. Furthermore, their complexity is the same as Dalal's. These features, along with some suggested further work, make parameterized difference revision operators a candidate for possible implementation in knowledge representation frameworks (such as Horn formulas, Answer Set Programming or Description Logics) handling belief change.

In their paper *A Probabilistic Interval-based Event Calculus for Activity Recognition*, Alexander Artikis, Evangelos Makris, and Georgios Paliouras build on their previous work in using the Event Calculus to compute probabilities of complex or 'long-term' activities given an input stream of simpler "short term" events. In this paper their approach is extended to define the notion of a probabilistic maximal interval, as opposed to a single time point, over which such a long term activity occurs. The authors describe and evaluate PIEC, a linear-time algorithm that computes all such credible intervals for a given data set of probabilistically specified input events, and their experiments identify common conditions under which PIEC out-performs time-point-based activity recognition. The use of the Event Calculus anchors this approach firmly in the realm of commonsense reasoning, allowing for succinct and intuitive definitions of complex activities by, for example, leveraging the Event Calculus' built-in commonsense law of inertia.

In *Valid Attacks in Argumentation Frameworks with Recursive Attacks*, Claudette Cayrol, Jorge Fandinno, Luis Fariñas del Cerro, and Marie-Christine Lagasque-Schiex extend previous work on theories of argumentation, an important framework for understanding how individuals use common sense to reason with conflicting information commonsense reasoning as well as for applications such as legal reasoning. Cayrol et al. consider abstract argumentation frameworks that facilitate representing *recursive* attacks, attacks whose target is other attacks, and study the problem of identifying the extension-dependent subset of the attack that carries the weight of the attack. As the authors show, their framework allows accurate representation and reasoning about several commonsense problems that previous frameworks were unable to handle.

In their contribution *LogAG: An Algebraic Non-Monotonic Logic for Reasoning with Graded Propositions*, Nourhan Ehab and Haythem Ismail investigate the belief revision problem in non-monotonic logics, exploring the use of ordered grades for

logical formulas as a means of resolving conflicts between propositions. Their proposed framework, LogAG, extends previous algebraic logical frameworks to encode preferences, priorities, trust, or certainties as grades assigned to classical logical formulae. The result is an expressive, weighted algebraic logic that can be utilized in an array of previously-studied tasks, offering a unifying approach to non-monotonic representation and reasoning.

In *Competing Hypotheses and Abductive Inference*, David Glass explores the nature of mutual exclusivity and independence between different hypotheses that entail the evidence in an abductive reasoning problem. Noting that adductive hypotheses need not be exclusive to be competitive, this paper advances an account of competition that depends not only on the hypotheses in question, but also the evidence under consideration. Acknowledging the difficulty of establishing mutual exclusivity among hypotheses, this work further explores the consequences of assuming exclusivity by means of three logical experiments. Stipulating a known correct model, these experiments proceed under different assumptions of exclusivity and independence among differing hypotheses, demonstrating how models can perform poorly when ignoring the overlap between the hypotheses.

In *What do you really want to do? Towards a Theory of Intentions for Human-Robot Collaboration*, Rocio Gomez, Mohan Sridharan, and Heather Riley focus on fusing two different approaches to planning: a logic-based approach that uses non-monotonic logic programming, and an approach that uses probabilistic models of uncertainty in sensing and actuation. This fusion captures the commonsensical way humans plan. The first approach captures the coarse-grained knowledge and reasoning that we typically use to construct a set of abstract actions, and that supports reasoning about satisfying agents' intentions (e.g., books are in libraries; if you need a book and you know it is in the library, go to the library), and the fine-grained sensor information and probabilistic reasoning that agents must use to construct an executable plan (e.g., procedures for navigation, object recognition, and obstacle avoidance). The fusion is facilitated by an architecture that centers on zooming, a meta-level process that allows mapping between the coarse-grained and fine-grained levels. After the non-monotonic logical reasoning level computes a plan of intentional abstract actions for each goal, plan execution proceeds by automatically zooming to the part of the fine-resolution transition diagram corresponding to the coarse-grained transition, and executing these concrete actions using probabilistic reasoning on sensor and effector information. Information flows freely between both levels. The authors' experiments demonstrate that including coarse-grained level results in greater computational efficiency and

accuracy, validating both the claims of the formal commonsense reasoning community, who have long argued (starting with McCarthy, 1958) that it is essential to capture the highly abstract models that humans use to express and reason about complex planning processes, and the claims of those who argue that formal methods are best used as part of a more complex commonsense architecture.

In *Revising Event Calculus Theories to Recover from Unexpected Observations*, Nikoleta Tsampanaki, Theodore Patkos, Giorgos Flouris, and Dimitris Plexousakis address the issue of conflict resolution between new observations and existing representations of narratives of events. Their work draws on ideas and principles in the field of belief revision and applies these to an epistemic version of the Event Calculus. The result is a formal framework able to model deterministic, dynamic causal domains and including a preference relation among candidate modifications to a given theory necessary to accommodate new information that appropriately respects standard belief revision principles such as the Principle of Minimal Change. They show how this framework can be implemented in ASP and embedded in an agent-based architecture that facilitates an ongoing reasoning and revision process as the agent receives input from and interacts with its environment. The flexibility of their implementation is demonstrated with a web interface that allows the user to upload domain examples, new observations and domain-specific update preferences and perform runs of the system based on these inputs.

References

Ernest Davis:

Representations of commonsense knowledge. Morgan Kaufmann series in representation and reasoning, Morgan Kaufmann 1990, ISBN 978-1-55860-033-1, pp. I-XXIII, 1-515

Daniel Jurafsky and James H. Martin. *Speech and Language Processing*, 3rd Edition, Prentice Hall, 2014.

Vid Kocijan, [Thomas Lukasiewicz](#), [Ernest Davis](#), [Gary Marcus](#), [Leora Morgenstern](#): A Review of Winograd Schema Challenge Datasets and Approaches. [CoRR abs/2004.13831](#) (2020)

Alex Krizhevsky, [Ilya Sutskever](#), [Geoffrey E. Hinton](#):

ImageNet classification with deep convolutional neural networks. [Commun. ACM 60\(6\)](#): 84-90 (2017)

Douglas B. Lenat, Mayank Prakash, Mary Shepherd:
CYC: Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks. *AI Magazine* 6(4): 65-85 (1986)

Hector J. Levesque, [Ernest Davis](#), [Leora Morgenstern](#):
The Winograd Schema Challenge. [KR 2012](#)

Vladimir Lifschitz, John McCarthy, and Leora Morgenstern (eds.). Proceedings, Logical Formalizations of Common Sense, AAAI Spring Symposium, 1991.

Christopher Manning, Foundations of Statistical Natural Language Processing, MIT Press, 1999.

John McCarthy: Programs with common sense, in Proceedings of the Teddington Conference on the Mechanization of Thought Processes, Her Majesty's Stationery Office, London, 1959, pp, 75-91.

John McCarthy. Situations, actions, and causal laws, 1963. Telch report, Stanford University, reprinted in M. Minsky (Ed.) Semantic Information Processing, MIT Press, Cambridge, Mass, 1968, pp. 410-417.

John McCarthy and Patrick J. Hayes. Some Philosophical Problems from the Standpoint of Artificial Intelligence, in . B. Meltzer, D. Michie (Eds.), Machine Intelligence, Vol. 4, Edinburgh University Press, Edinburgh, 1969, pp. 463-502.

Drew McDermott (1987) A critique of pure reason. *Computational Intelligence* 3(33), pp. 151–160

Moor, J. (2006). The Dartmouth College Artificial Intelligence Conference: The Next Fifty Years. AI Magazine, 27(4), 87.

Silver, David; Huang, Aja; Maddison, Chris J.; Guez, Arthur; Sifre, Laurent; Driessche, George van den; Schrittwieser, Julian; Antonoglou, Ioannis;

Panneershelvam, Veda; Lanctot, Marc; Dieleman, Sander; Grewe, Dominik; Nham, John; Kalchbrenner, Nal; Sutskever, Ilya; Lillicrap, Timothy; Leach, Madeleine; Kavukcuoglu, Koray; Graepel, Thore; Hassabis, Demis (28 January 2016).

"Mastering the game of Go with deep neural networks and tree search". *Nature*. 529 (7587): 484–489. doi:10.1038/nature16961.