Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT
Artificial Intelligence 153 (2004) 1–12
Artificial Intelligence
www.elsevier.com/locate/artint

Introduction: Progress in formal commonsense reasoning

Ernest Davis [a], Leora Morgenstern [b,*]

[a] *Courant Institute, New York University, New York, NY 10012, USA*
[b] *IBM T.J. Watson Research Center, Hawthorne, NY 10532, USA*

1. Approaches to formal commonsense reasoning

This special issue consists largely of expanded and revised versions of selected papers of the Fifth International Symposium on Logical Formalizations of Commonsense Reasoning (Common Sense 2001), held at New York University in May 2001.[1,2] The Common Sense Symposia, first organized in 1991 by John McCarthy and held roughly biannually since, are dedicated to exploring the development of formal commonsense theories using mathematical logic.

Commonsense reasoning is a central part of human behavior; no real intelligence is possible without it. Thus, the development of systems that exhibit commonsense behavior is a central goal of Artificial Intelligence. It has proven to be more difficult to create systems that are capable of commonsense reasoning than systems that can solve "hard" reasoning problems. There are chess-playing programs that beat champions [5] and expert systems that assist in clinical diagnosis [32], but no programs that reason about how far one must bend over to put on one's socks. Part of the difficulty is the all-encompassing aspect of commonsense reasoning: any problem one looks at touches on many different types of knowledge. Moreover, in contrast to expert knowledge which is usually explicit, most commonsense knowledge is implicit. One of the prerequisites to developing commonsense reasoning systems is making this knowledge explicit.

John McCarthy [25] first noted this need and suggested using formal logic to encode commonsense knowledge and reasoning. In the ensuing decades, there has been much research on the representation of knowledge in formal logic and on inference algorithms to

* Corresponding author.
E-mail addresses: davise@cs.nyu.edu (E. Davis), leora@us.ibm.com (L. Morgenstern).

[1] Common Sense 2001 was chaired by the present authors, John McCarthy, and Ray Reiter.

[2] One of the papers in this issue, by Murray Shanahan, is an expanded and revised version of a paper presented at Common Sense 1998.

doi:10.1016/j.artint.2003.09.001

manipulate that knowledge. The arguments for a declarative knowledge representation—it allows the explicit representation of knowledge; it is modular; it supports modification far more easily than implicit, procedural knowledge—have gained credence not only among the AI community, but in the broad field of computer science. Basic principles of knowledge representation have been incorporated into the design of object-oriented languages and rule-based systems. But the formalization of commonsense reasoning remains an elusive goal. Seemingly trivial reasoning problems that, in McCarthy's words, can be carried out by any non-feeble minded human are still beyond the representational and reasoning abilities of existing theories and systems. (See [8] for examples.)

Progress has been slow because formalizing commonsense reasoning presents a variety of challenges. One must (1) develop a formal language that is sufficiently powerful and expressive; (2) capture the many millions of facts that people know and reason with;[3] (3) correctly encode this information as sentences in a logic; and (4) construct a system that will use its knowledge efficiently. The knowledge of these difficulties is as old as the endeavor itself. Forty-four years ago, linguist and logician Yehoshua Bar-Hillel (see the Discussion section of [25]) argued the inadequacy of standard deduction for planning in real-world domains in which circumstances can change—and perhaps presaged the need for default logics; pointed out the problems in formalizing even a simple relation like "At"; and questioned how a computer could choose from millions of facts the few facts relevant to a specific problem at hand.

There have been two strategies for tackling these very hard problems. On the one hand, researchers have methodically and painstakingly worked on foundational problems and have constructed small ("toy") formalizations to test their progress. On the other hand, there have been attempts to encode vast amounts of facts to enable broad commonsense reasoning. The Cyc Project [20,21], which aims to construct an encyclopedic knowledge base of all facts needed to perform commonsense reasoning, is the best known of these efforts. Other projects with a similar flavor include DARPA's HPKB project [33] which focused on creating very large and fast knowledge bases for commonsense reasoning in the domains of warfare and geography. In the subsequent discussion, we will refer to the first strategy as the foundational approach, and to its proponents and practitioners as foundationalists. Because of the renown of the Cyc project, we will refer to the second strategy as the Cyc approach. Cyc's adherents and practitioners are commonly called CycLists [44] and throughout this paper, we will extend this usage to all proponents of the second approach.

The foundational and Cyc research programs focus on different subsets of Bar-Hillel's challenges. Foundationalists have concentrated on three issues: First, they have spent much effort developing sufficiently powerful and expressive alternatives to or extensions of classical logics. These include default logics that allow plausible reasoning [26,29,37], as well as logics of knowledge and belief [10], logics of obligation [35], and temporal reasoning [11,17,24]. Second, they have focused much attention on ensuring that the logical axioms they write down precisely model the facts they are attempting to formalize.

---

[3] Estimates on the number of facts people know range from ten million to several hundred million: see Erik Mueller's article at http://www.signiform.com/erik/pubs/cshumans.htmfor a discussion and list of references.

This is a process in which the careful teasing out of implicit knowledge, the delineation of boundary cases, and the recognition and construction of general rules are as important as determining how one can translate such knowledge into a formal system. Third, they have studied the ways in which their systems can efficiently make inferences.[4] Where foundationalists have fallen short is in the construction of large-scale theories. A typical foundational research project will focus on one or several toy problems, ranging in size from a few axioms to several dozen axioms.

The Cyc research program, on the other hand, focuses exactly on the construction of large-scale theories. CycLists have thus far constructed a knowledge base of several million commonsense facts, and plan to acquire at least several million more in the next few years. They are, of course, concerned about getting things right: the intention is to say a lot of facts and get as many right as possible. But absolute correctness is not a goal: it is expected that the representational language will be imperfect, that some data will be wrong, and that the translation of some facts into logic will go awry. Thus, they are not nearly so concerned with getting comprehensive solutions to fundamental representational problems [7]. They are also naturally concerned with efficient reasoning, because their theories are so large, but their methodologies to promote efficiency rest on fairly standard heuristics.[5]

There is less interaction that one might expect between the foundationalist and Cyc research groups. None of the papers in this special issue, for example, cite or mention Cyc at all. This disconnect is in large measure due to the fact that the Cyc camp has published very little about the representations they use since [21], now thirteen years out of date, and hardly as informative as one would have wished, even at the time. Even now that the Cyc Upper Ontology has been made available to the public,[6] the great majority of the knowledge base remains inaccessible.

Yet one cannot entirely explain the gap between the foundationalists and CycLists on Cyc's limited publication activity. Cyc is cited frequently by researchers working on large-scale ontologies and knowledge bases and knowledge engineering. Indeed, there has in recent years been increasing collaboration between the academic community and Cyc. Specifically, portions of Cyc's Upper Ontology were incorporated into the HPKB project [33] and served as the starting point for the Rapid Knowledge Formation project, which aimed to develop tools to enable subject matter experts to quickly enter domain knowledge [42].

Rather, the fundamental reason for the disconnect is probably a lack of shared interests. While both foundationalists and CycLists are nominally interested in formalizing commonsense reasoning, they are in fact working on different research problems, at least for the time being. Since both a solid foundation and large-scale formalizations are

[4] One might think that efficiency is not a particular concern due to the small size of most existing axiomatizations. In fact, however, theorem provers can run into efficiency problems even when dealing with relatively small axiom sets. Moreover, the computational complexity of some tasks like planning arises from the search through possible states, not the number of axioms in the planning theory.

[5] See Cyc's Ontology Engineers handbook, at http://www.cyc.com/doc/handbook/oe/06-el-hl-and-the-canonicalizer.html.

[6] See http://www.cyc.com/cyc-2-1/cover.html.

necessary for formalizing commonsense reasoning, one would hope that as these projects develop, their research interests will eventually converge.

2. Trends in formal commonsense reasoning

Commonsense reasoning is such a broad area that one might expect that researchers attempting to capture it would head off in hundreds or thousands of different directions. For better or worse, this has not happened. Research has tended to cluster in a few areas, which have remained relatively stable over the last few decades. These include default or nonmonotonic reasoning, logics of knowledge, and temporal reasoning, mentioned above, as well as spatial reasoning [36] and theories of belief, desire, and intention [6].

There has been substantial progress in each of these areas. For example, research in nonmonotonic reasoning, in it earliest days [26,29,37], focused mostly on developing logics that allowed one to express default rules and perform default inference. Later researchers gave semantics to these logics [31], explored the connection between logic programming and nonmonotonic logic [30,34], and investigated new paradigms for default reasoning, such as consequence relations [18] and belief revision [1]. In addition, researchers have thoroughly explored the integration of default reasoning with theories of action [13,38,41], investigated the connections among existing systems of default reasoning [9,15,16], and considered how nonmonotonic reasoning and planning might best be implemented and used for applications [12,22].

This is indicative of the research trends that have recently emerged: the integration of different fields of commonsense reasoning, the development of unified, expressive theories, the development of implemented systems, and the examination of somewhat larger problems than have previously been considered. These trends are evident in the papers appearing in this issue.

Many papers in this collection owe a direct debt to problems first proposed by John McCarthy, e.g., in [25,27]. This was not true twelve years ago when the Common Sense Symposia first started, nor even five years ago—perhaps because at that time the formal commonsense community was still struggling with preliminary but difficult foundational problems. The progress that has been made in areas such as nonmonotonic reasoning and temporal reasoning has enabled the community to turn its attention toward creating larger and more flexible commonsense reasoning systems that incorporate some of the spirit of the Advice Taker. McCarthy was far ahead of his time; the formal commonsense community may now be ready to catch up with his vision.

3. Papers in the special issue

One of the criticisms of the foundational camp is that so many researchers develop their own theories to solve a particular problem, even when similar theories already exist. The result is a large number of theories, mostly incomparable, each suited to some problem, but none suited to a broad class of problems. This is particularly true of the reasoning-about-actions community, where the Situation Calculus [24], the Event Calculus [17,41], and the Fluent Calculus [43] are only a few of the many formalisms currently available. There have been some efforts to compare various popular formalisms [40], but ultimately one must choose a specific, narrow formalism. Whether or not language shapes thought, it is certainly the case that one's theory is shaped by the ontology of the formalism that one chooses. The state of the art is that when a researcher selects a temporal logic, he has already committed himself to significant decisions about the theory that he will create. For example, the Event Calculus commits one to a linear time structure in which time points are arranged in a total order; the Situation Calculus commits one to a branching time structure in which time points are arranged in a partial order.

Bennett and Galton advance the state of the art one large step further. They present a very expressive language (Versatile Event Logic or VEL) for representing temporal relationships and events. VEL incorporates a variety of different approaches to temporal reasoning within a single framework. For example, existing formalisms view events either as transitions between states or occurrences over intervals or suitably modified syntactic units, and force the user to explicitly refer to (respectively) either time points or time intervals or propositional tenses. In contrast, VEL allows the user to choose whichever representation is most suitable for his needs, and further, to simultaneously use multiple representations if desired. In their generous and flexible framework, Bennett and Galton also provide some unusual features, such as a modal operator to describe alternative histories. This feature can facilitate hypothetical and counterfactual reasoning. Bennett and Galton suggest that their framework can be used as a lingua franca for comparing different temporal formalisms, and demonstrate that the Situation Calculus and the Event Calculus can be represented within VEL. This is a modest claim: VEL may be most useful when researchers use it as their primary temporal formalism and exploit its flexibility to develop more realistic and powerful theories of action than are possible using existing logics.

In the past two decades, researchers have struggled to formalize the notion of causation. The aim has been to develop a theory of causation that supports inference in multiple directions (e.g. from cause to effect, from effect to cause, from behavior to object properties, etc. [39]). This entails, among other things, solving several temporal reasoning problems: determining which fluents, or time-varying properties, do *not* change as the result of an action, without having to explicitly list all such properties (the frame problem), determining the indirect effects of an action, without having to explicitly list all such effects (the ramification problem), and concisely representing the conditions necessary for the successful performance of an action (the qualification problem).

Researchers early on [26,28] conjectured that default logics could help solve these problems, partly because statements of the form "typically X is the case" can be used to concisely represent phenomena of change and causation, and default logics are designed to express such statements. For example, default logics make it easy to say that it is typically the case that if you strike a match, it will light, and that it is typically the case that if you park your car somewhere, it will remain there until you next drive it. However, a naive integration of default logic into temporal logic can be inadequate, particularly with respect to reasoning about conflicting defaults. Consider for example, the principle of inertia, which states that properties typically remain the same over time. Two instances of the principle of inertia—animals typically stay alive and guns typically stay loaded—conflict in a scenario in which a gun is loaded and after a short time fired point-blank at a turkey's head. The difficulty (the Yale Shooting problem [14]) is that naive default temporal logics do not support the preference of the expected models (the gun stays loaded, and the turkey dies) over the unexpected models (the turkey stays alive, entailing that the gun must have become unloaded), and thus do not support any prediction about the turkey's status.

There has been a flurry of solutions to these temporal reasoning problems, and an emerging consensus on the fundamental concepts that support these solutions. For example, it is now understood that a proper formalization of the principle of inertia, which is central to a solution to the frame problem, requires a sufficiently rich theory of causation, and that formalizing constraints between fluents can help reason about ramifications [43]. Often, however, such solutions have appeared in piecemeal fashion.

Giunchiglia, Lee, Lifschitz, McCain, and Turner present a theory of causal reasoning that integrates many of the insights that have been learned from this research. In their nonmonotonic causal logic, one can distinguish between a fact being true and a fact having a cause. There is a strong connection between a fact being caused and a fact obtaining (being true at some point): a fact is caused iff the fact obtains. (The "if" part of this principle can be explicitly disabled for specific facts, if desired.) Giunchiglia et al. show how their theory can be used to elegantly formalize the principle of inertia and constraints. They further introduce the action description language $\mathcal{C}+$, a formal language for describing transition systems, and define a correspondence between action descriptions and causal theories.

Giunchiglia et al. also deal with implementational issues. The development of implemented systems for reasoning about actions has received increasing attention in recent years as theories of action have matured, enabling researchers to study concrete examples. Giunchiglia et al. demonstrate how problem statements in $\mathcal{C}+$ can be translated into the Causal Calculator [23] and thereby demonstrate how solutions to some well-known toy problems of commonsense reasoning can be computed. Their work thus integrates an action language, a theory of causal reasoning, and an implemented system for reasoning about action and causation into one unified package.

The enumeration and elaboration of and detailed investigation into toy problems has a rich and worthy history in AI. Toy problems such as Missionaries and Cannibals and Blocks World have generated a wealth of groundbreaking research. McCarthy compares the use of such toy problems to the use of drosophila in biological research, and argues [27] that they are essential to progress in formal commonsense reasoning. He believes that especially because commonsense reasoning is so complex and difficult to formalize, it is important to study simpler cases that allow researchers to focus on the particular research problems they are trying to solve. For example, the simplified Blocks World allows researchers to concentrate on planning issues and temporal reasoning problems and ignore motion planning and sensing.

Nevertheless, many members of the research community believe that in addition to toy problems, researchers should also look at somewhat larger problems. They believe that toy problems can often simplify a real-world problem to such an extent that nearly all interesting aspects are altered, and research results are not of much use. For example, certain variants of the Yale shooting problem are so weakened that they can be solved without any formalization of causation. Moreover, because most toy problems focus on only one or two aspects of commonsense reasoning, solving them does not entail exploring how different commonsense problems interact with one another, or how one can integrate solutions. In addition, exclusively examining toy problems promotes insularity by reducing researchers' motivation to examine commonsense problems that arise in the real world, and prevents researchers from discovering if their formalizations will scale up to larger problems.

There has therefore been an attempt to enlarge "drosophila" from toy-problem size. During the last few years, several mid-sized challenge problems have been proposed. Erik Sandewall has suggested for commonsense research two microworlds with fairly rigorous specifications, a simplified traffic world and a simplified zoo world.[7] A sampling of commonsense reasoning problems, such as reasoning about staking a plant, reasoning about cracking an egg, and reasoning about surprising a friend with a birthday present, can be found on the Common Sense Problem Page.[8] Some of these problem descriptions list a number of variants. In general an acceptable representation for such a problem should also handle a wide range of variants; this helps guarantee that the representation is not too narrowly tailored to the specific problem. (McCarthy [27] calls this the principle of *elaboration tolerance.*) For example, variants of the staking problem include situations where the stake is not pushed into the ground; where the stake is very far from the plant; where the gardener attempts to use a heavy chain on a small plant, and so on.

The potential advantages of working on such mid-sized problems include the development of core, reusable theories of commonsense reasoning, testing existing theories as they are integrated into axiomatizations of mid-sized problems, and the discovery of new representational issues and problems that might not appear in an artificially small toy problem.

This issue contains two papers describing their approaches to such mid-sized benchmark problems. *Akman, Erdogan, Lee, Lifschitz, and Turner* present axiomatizations of both the zoo and traffic worlds. They axiomatize these microworlds in $\mathcal{C}+$, the action description language discussed above by Giunchiglia et al., thus demonstrating $\mathcal{C}+$'s expressiveness and its ability to handle the frame, ramification, and other temporal reasoning problems. The new representational issues that arose included representing continuous motion using integer arithmetic, and representing change in the absence of action. For example, if a car is already on a road and there is no specific action to stop or leave the road, the car will continue and its position will change. The paper proposes ways to represent such facts using dynamic laws that do not contain action constants.

Shanahan gives a detailed axiomatization of the egg-cracking domain. The problem, which appears on the Common Sense Problem Page, is to characterize the correct procedure of cracking an egg and transferring its contents to a bowl. Variants include situations in which the bowl is made of loose-leaf paper, the bowl is upside down, and the agent tries the egg-cracking procedure with a hard-boiled egg. The problem itself is quite a complex one, since cracking an egg involves reasoning about so many domains of physical reasoning: containment and parts, materials, collisions, liquids, and vessels. Shanahan develops simple re-usable theories for each of these domains. He integrates this theory with the extension of the Event Calculus that he had earlier developed to solve

---

[7] See http://www.ida.liu.se/ext/etai/lmw/.

[8] See http://www-formal.stanford.edu/leora/commonsense.

the frame problem [41] to arrive at an elegant axiomatization of the egg-cracking domain. Moreover, Shanahan demonstrates that his theory can be modified, with little effort, to handle the listed variants of the egg-cracking problem, as well as several that are not listed. Shanahan argues, however, that one cannot forever do armchair axiomatizations. He contends that it is necessary to at some point check one's formalizations in the real world, for example, by running an implementation of one's axiomatization on a robot. In the egg-cracking axiomatization, he therefore attempts, as much as possible, to work toward the ultimate goal of having a set of bottom-level predicate and function symbols that can be anchored through robot sensors and actuators.

The need to closely couple formal axiomatizations with an implemented robot that senses and acts in the real world is the driving force behind the work of *Amir and Maynard-Zhang*. They present an integration of Brooks' subsumption architecture [4] with logical representations and formal reasoning algorithms. Like Brooks, they suggest decomposing a domain along behavioral lines. Thus, their architecture includes layers for high-level motion planning, local action planning, destination seeking, and obstacle avoidance. Each layer consists of (1) the set of axioms describing the layer's behavior, (2) the sensory input and output latches, which accept input axioms from the sensors and from higher layers, (3) the set of goal sentences which determines, via proof and instantiation, the layer's behavior, and (4) the default assumptions, used to implement (via circumscription) nonmonotonic reasoning. A layer operates by collecting data from sensors and inputs from higher level theories and trying to prove the layer's goal from the theory and the default assumptions. Goals are transmitted to the layer below, or to the robot manipulators. Amir and Maynard-Zhang's system has been implemented on a Nomad 200 robot which travels to different rooms in a multi-story building. The architecture affords great flexibility: it is possible to correct the robot's behavior at run-time by giving it a fact encoded in logical form. This comes very close to McCarthy's dream of the Advice Taker.

McCarthy has long argued that a commonsense theory should be elaboration tolerant. That is, it should be relatively easy to add new facts to a theory, or new details to a problem statement, without having to rework large parts of one's existing theory. McCarthy has taken the classic Missionaries and Cannibals problem (MCP), in which three missionaries try to ferry three cannibals across a river without being eaten [2]. In the original problem, there is one boat which holds two people, and if the cannibals ever outnumber the missionaries, the missionaries get eaten. One can encode the problem as a sequence of triples, each of which indicates the number of missionaries, cannibals, and boats on the left bank of the river. But this representation is exceedingly brittle. It cannot support a solution to a problem that has extra facts about the river or boats, for example. McCarthy [27] proposes 19 elaborations, including the 4 person MCP (4 missionaries and 4 cannibals; this is unsolvable), the conversion MCP (three missionaries can convert 1 cannibal if they get him alone), and the big cannibal MCP (one cannibal is so large that he cannot fit into a boat with anyone else).

Gustafsson and Kvarnström show how one can use techniques from object-oriented programming to construct theories that are elaboration tolerant. They argue that the key to elaboration tolerance is proper organization, which is facilitated by using the object-oriented paradigm. One should initially design the domain formalization at a high level of abstraction and organize the domain so that one can add details later on. In their theory, domain entities are represented as objects, which are organized into classes. As in classical object-oriented programming, methods are associated with objects. However, a method in this paradigm is not a sequence of code; it is a set of formulas that must be satisfied when the method is invoked. Inheritance allows re-use of appropriate methods, while adding, or when necessary, overriding existing attributes, constraints, and methods allows new information to be added. For example, the big cannibal MCP is modeled by adding a new class of Big Cannibals that extends the original Cannibals class; a new constraint is then added to this class so that whenever a big cannibal is on a boat, he is alone in the boat. Gustafsson and Kvarnström show that all solvable MCP elaborations that McCarthy proposed can be handled in their new paradigm.

Despite the trend toward larger benchmark problems, most axiomatizations that formalists produce are still quite small. *Gordon* argues that for the broad domain of reasoning about rational action (and a fortiori for the entire domain of commonsense reasoning), it is unreasonable to expect that existing techniques will ever result in comprehensive formalizations. He contends that the amount of research required to correctly formalize even one basic concept is so large that the scope of the endeavor cannot be supported by the AI community. This is, of course, close to the Cyc camp's arguments. But Gordon's solution is quite different. He describes a method that he used to sketch out a compendium of planning strategies. This method consists of four steps: first, collecting strategies, using expert texts, interviews, and interpretive observation; second, translating each strategy into a pre-formal representation, a stylized and restricted subset of English; third, identifying the words and phrases corresponding to concepts that would need to be formalized in a logical axiomatization; and fourth, organizing representational terms and areas. For example, one warfare strategy identified is giving false information to one's enemy. Some of the concepts included in the pre-formal representation for this strategy are adversarial relationships, cooperative plans, information, and false beliefs, concepts which appear in other strategies not only in the domain of warfare, but also in the business world and in the (anthropomorphized) animal kingdom.

Gordon has made contributions in both representation and methodology. He has demonstrated how to collect and organize a large amount of commonsense information in a relatively short period of time. His work is intentionally broad rather than deep. Many of the representational terms he identifies will need to be further decomposed into primitive concepts in a formal representation. Nevertheless, the expectation is that the pre-formal representation will facilitate the development of the structure of the formal theory. Moreover, converting a collection of facts in pre-formal representation into a formal theory may be easier, or at least more principled, than the two techniques used today: sitting down to write down in formal logic the commonsense facts one knows, or constructing a microworld that corresponds to what one knows and subsequently axiomatizing that microworld. For the domain of planning strategies, Gordon has transformed "what it is that one knows" before one sets out to write formal axioms. His methodology should enable us to make this transformation for other domains as well.

Because of the computational complexity of theorem proving, logicists have long been concerned with the efficiency of their theories. This concern is evident in several of the papers already discussed. For example, one of the motivations for Amir and Maynard-Zhang's subsumption architecture is divide-and-conquer: because each layer has fewer axioms, theorem proving is quicker. One goal of formal commonsense reasoning is making reasoning more efficient by declaratively encoding heuristic information [25]. That is, it should be possible to tell a planning system that one strategy is better than another. This approach is explored in detail by *Sierra-Santibáñez*. She shows that a set of simple but powerful meta-rules for action selection suffices to enable her program to find better plans more quickly both in the blocks world and in a logistics domain.

Efficiency is also one of the motivating concerns of *Benferhat, Kaci, Le Berre, and Williams*. They consider the problem of revising one's beliefs as one gets more information. While most algorithms for standard and iterated belief revision mandate throwing out inconsistent information, Benferhat et al. propose a method to weaken inconsistent information by turning a piece of inconsistent information into one disjunct of a larger statement. This method turns out to be equivalent, but more efficient than, the lexicographical approach, a coherence-based approach where an inconsistent knowledge base is replaced by a set of maximally preferred consistent subbases [3,19].

4. Dedication

Ray Reiter (1939–2002) was one of the founding fathers of formal AI. A recipient of the IJCAI Award for Research Excellence in 1993, he was widely recognized for his leadership in the areas of deductive databases, logic programming, truth maintenance systems, and temporal reasoning. His seminal work on default logic [37] along with the work of [29] and [26] established the field of nonmonotonic reasoning.

Reiter extended and popularized the situation calculus, which prior to his work had been widely considered to be too inexpressive to be useful. He used the situation calculus and the concept of goal regression to provide one of the first comprehensive solutions to the frame problem. In doing so, he jump-started a fruitful and active line of inquiry into temporal reasoning. His recent book detailing this exploration, *Knowledge in Action* [38], is an exemplar of the best work of foundationalist researchers: rigorous yet accessible; formal yet grounded in application. He was a leader in the knowledge representation community, helped establish the conferences on Principles of Knowledge Representation and Reasoning, and chaired the first of those conferences in 1989. He was an active participant in the Common Sense Symposia, and served as one of the program chairs of Common Sense 2001, the symposium that gave rise to this special issue.

Ray Reiter's memory lives on in all those in the AI community who knew him. His influence lives on in the excellent group of Ph.D. students he advised, and in everyone who reads his papers. He is sorely missed by all those he touched. We dedicate this issue to his memory.

References

[1] C.E. Alchourrón, P. Gardenförs, D. Makinson, On the logic of theory change: Partial meet contraction and revision functions, J. Symbolic Logic 50 (2) (1985) 510–530.

[2] S. Amarel, On representation of problems of reasoning about action, in: D. Michie (Ed.), Machine Intelligence, vol. 3, Edinburgh University Press, Edinburgh, 1971, pp. 131–171.

[3] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, H. Prade, Inconsistency management and prioritized syntax-based entailment, in: R. Bajcsy (Ed.), Proceedings of IJCAI-93, Chambéry, France, Morgan Kaufmann, San Francisco, CA, 1993, pp. 640–647.
[4] R.A. Brooks, Intelligence without representation, Artificial Intelligence 47 (1–3) (1991) 139–159.
[5] M. Campbell, A.J. Hoane Jr., F.H. Hsu, Deep blue, Artificial Intelligence 134 (1–2) (2002) 57–83.
[6] P.R. Cohen, H.J. Levesque, Intention is choice with commitment, Artificial Intelligence 42 (2–3) (1990) 263–309.
[7] B.J. Copeland, CYC: A case study in ontological engineering, Electron. J. Anal. Philos. 5 (1997).
[8] E. Davis, The naive physics perplex, AI Magazine 19 (3) (1998) 51–79.
[9] M. Denecker, V.W. Marek, M. Truszczynski, Uniform semantic treatment of default and autoepistemic logics, Artificial Intelligence 143 (1) (2003) 79–122.
[10] R. Fagin, J.Y. Halpern, Y. Moses, M.Y. Vardi, Reasoning About Knowledge, MIT Press, Cambridge, MA, 1995.
[11] A. Galton, Time and change for AI, in: D.M. Gabbay, C.J. Hogger, J.A. Robinson (Eds.), Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 4, Epistemic and Temporal Reasoning, Oxford University Press, Oxford, 1995, pp. 175–240.
[12] M. Gelfond, V. Lifschitz, Compiling circumscriptive theories into logic programs, in: Proceedings of AAAI-88, St. Paul, MN, AAAI Press/MIT Press, Cambridge, MA, 1988, pp. 455–459.
[13] M. Gelfond, V. Lifschitz, Representing action and change by logic programs, J. Logic Programming 17 (2–4) (1993) 301–321.
[14] S. Hanks, D.V. McDermott, Nonmonotonic logic and temporal projection, Artificial Intelligence 33 (3) (1987) 379–412.
[15] K. Konolige, On the relation between default and autoepistemic logic, Artificial Intelligence 35 (3) (1988) 343–382.
[16] K. Konolige, On the relation between autoepistemic logic and circumscription, in: N.S. Sridharan (Ed.), Proceedings of IJCAI-89, Detroit, MI, Morgan Kaufmann, San Francisco, CA, 1989, pp. 1213–1218.
[17] R.A. Kowalski, M.J. Sergot, A logic-based calculus of events, New Generation Comput. 4 (1) (1986) 67–95.
[18] S. Kraus, D.J. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artificial Intelligence 44 (1–2) (1990) 167–207.
[19] D.J. Lehmann, Another perspective on default reasoning, Ann. Math. Artificial Intelligence 15 (1) (1995) 61–82.
[20] D.B. Lenat, Cyc: A large-scale investment in knowledge infrastructure, Comm. ACM 38 (11) (1995) 32–38.
[21] D.B. Lenat, R.V. Guha, Building Large Knowledge Based Systems, Addison Wesley, Reading, MA, 1990.
[22] V. Lifschitz, Answer set programming and plan generation, Artificial Intelligence 138 (1–2) (2002) 39–54.
[23] N. McCain, Causality in commonsense reasoning about actions, Ph.D. Dissertation, University of Texas at Austin, TX, 1997.
[24] J. McCarthy, P.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.
[25] J.L. McCarthy, Programs with common sense, in: Proceedings of the Teddington Conference on the Mechanization of Thought Processes, Her Majesty's Stationary Office, London, 1959, pp. 75–91.
[26] J.L. McCarthy, Circumscription: A form of non-monotonic reasoning, Artificial Intelligence 13 (1–2) (1980) 23–79.
[27] J.L. McCarthy, Elaboration tolerance, in: Working Papers of the Fourth International Symposium on Logical Formalizations of Commonsense Reasoning, Common Sense 98, 1998.
[28] D.V. McDermott, A temporal logic for reasoning about processes and plans, Cognitive Sci. 6 (1982) 101–155.
[29] D.V. McDermott, J. Doyle, Non-monotonic logic I, Artificial Intelligence 13 (1–2) (1980) 41–72.
[30] J. Minker, An overview of nonmonotonic reasoning and logic programming, J. Logic Programming 17 (2–4) (1993) 95–126.
[31] R.C. Moore, Semantical considerations on nonmonotonic logic, Artificial Intelligence 25 (1) (1985) 75–94.
[32] B.N. Nathwani, K. Clarke, M.C. Pike, S.P. Azen, Evaluation of an expert system on lymph node pathology, Human Pathology 28 (9) (1997) 1097–1110.

[33] A. Pease, V.K. Chaudhri, F. Lehmann, A. Farquhar, Practical knowledge representation and the DARPA high performance knowledge bases project, in: A.G. Cohn, F. Giunchiglia, B. Selman (Eds.), Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR2000), Morgan Kaufmann, San Francisco, CA, 2000, pp. 717–724.

[34] L.M. Pereira, J.N. Aparicío, J.J. Alferes, Non-monotonic reasoning with logic programming, J. Logic Programming 17 (2–4) (1993) 227–263.

[35] H. Prakken, Two approaches to the formalisation of defeasible deontic logic, Studia Logica 57 (1) (1996) 73–90.

[36] D.A. Randell, Z. Cui, A.G. Cohn, A spatial logic based on regions and connection, in: B. Nebel, C. Rich, W.R. Swartout (Eds.), Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR1992), Morgan Kaufmann, San Francisco, CA, 1992, pp. 165–176.

[37] R. Reiter, A logic for default reasoning, Artificial Intelligence 13 (1–2) (1980) 81–132.

[38] R. Reiter, Knowledge in Action, MIT Press, Cambridge, MA, 2001.

[39] E. Sandewall, Features and Fluents: The Representation of Knowledge about Dynamical Systems, vol. I, Oxford University Press, Oxford, 1994.

[40] E. Sandewall, Panel on ontologies for action and change: Issues and questions, ETAI News J. Reasoning about Actions and Change 1 (3) (1997).

[41] M. Shanahan, Solving the Frame Problem, MIT Press, Cambridge, MA, 1997.

[42] G. Tecuci, M. Boicu, M. Bowman, D. Marcu, An innovative application from the DARPA knowledge bases programs: Rapid development of a course-of-action critiquer, AI Magazine 22 (2) (2000) 43–61.

[43] M. Thielscher, Ramification and causality, Artificial Intelligence 89 (1997) 317–364.

[44] C. Thompson, The know-it-all machine, Lingua Franca 11 (6) (2001).