

On John McCarthy's 80th Birthday, in Honor of his Contributions

Patrick J. Hayes and Leora Morgenstern

Abstract

John McCarthy's contributions to computer science and artificial intelligence are legendary. He invented Lisp, made substantial contributions to early work in timesharing and the theory of computation, and was one of the founders of artificial intelligence and knowledge representation. This article, written in honor of McCarthy's 80th birthday, presents a brief biography, an overview of the major themes of his research, and a discussion of several of his major papers.

Introduction

Fifty years ago, John McCarthy embarked on a bold and unique plan to achieve human-level intelligence in computers. It was not his dream of an intelligent computer that was unique, or even first: Alan Turing (Turing 1950) had envisioned a computer that could converse intelligently with humans back in 1950; by the mid 1950s, there were several researchers (including Herbert Simon, Allen Newell, Oliver Selfridge, and Marvin Minsky) dabbling in what would be called artificial intelligence. What distinguished McCarthy's plan was his emphasis on using mathematical logic both as a language for representing the knowledge that an intelligent machine should have, and as a means for reasoning with that knowledge. This emphasis on mathematical logic was to lead to the development of the *logician* approach to artificial intelligence, as well as to the development of the computer language Lisp.

John McCarthy's contributions to computer science and artificial intelligence are legendary. He revolutionized the use of computers with his innovations in timesharing; he invented Lisp, one of the longest-lived computer languages in use; he made substantial contributions to early work in the mathematical theory of computation; he was one of the founders of the field of artificial intelligence; and he foresaw the need for knowledge representation before the field of AI was even properly born.

John McCarthy turned 80 on September 4, 2007. This article is a celebration of that milestone. It in-

Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: John McCarthy

cludes a glimpse of McCarthy's life, an overview of the major themes of his work, and a discussion of several of his major papers. The aim is to introduce a more complete picture of McCarthy's long-term research to the AI Magazine readership. We hope it will help readers appreciate the range and depth of innovation in McCarthy's research.

Background

McCarthy was born in Boston in 1927 to John Patrick and Ida Glatt McCarthy, immigrants from, respectively, Ireland and Lithuania. The Depression started a few years after his birth; McCarthy's parents lost their house, and the family — which now included a second child — became briefly peripatetic. They lived for a short while in New York and then in Cleveland before finally settling in Los Angeles, where the senior McCarthy worked as an organizer for the Amalgamated Clothing Workers.

Like many child prodigies, John McCarthy was partly self-educated. Due to childhood illness, he began school a year late, but he quickly made up the time on his own, skipped several grades, and wound up graduating from high school two years early. When he was a teenager he developed an interest in mathematics and decided he wanted to go to the California Institute of Technology. At fifteen, he bought the calculus textbook then used in the Cal Tech freshman calculus course and taught it to himself in its entirety. In the hubris of youth, he applied only to Cal Tech, writing a one-sentence statement of purpose on his college application: "I intend to be a professor of mathematics." When he arrived at Cal Tech, he discovered that the textbook for the course had been changed. But the self-study paid off. Before classes started, he bumped into one of the instructors for freshman calculus, asked him a detailed question related to one of the textbook problems on which he had been working, and showed him his notebook with worked-out problems. "I don't want you in my class," that instructor said, and arranged for him to be given credit for freshman calculus. The same thing happened when he met the instructor for sophomore calculus. Ultimately, he was plunked down, at the age of 16, in graduate math class.

Birth of AI

After McCarthy received his Bachelors of Science from Cal Tech, he began his graduate studies there. For McCarthy, the intellectual ferment at the time was palpable. The excitement generated by the use of early computers to decode enemy messages — and thus help win World War II — was still running high. The theoretical foundations for computing had recently been laid by Church and Turing. There was a growing interest in the workings of the mind, in what would later be called cognitive science.

In September 1948, McCarthy went to the Hixon Symposium on Cerebral Mechanisms in Behavior (Jef-

fress 1951), a conference that joined together leading researchers in different areas related to cognitive science, including mathematicians Alan Turing and Claude Shannon, and psychologist Karl Lashley. As he listened to the discussions comparing computers and the brain, McCarthy had a watershed moment. From that time on, his chief interests related to the development of machines that could think like people. Indeed, some of the example problems that are present in McCarthy's papers, such as the monkeys and bananas problem, come from the Hixon Symposium: Karl Lashley described an ape pulling over a box and climbing on it in order to get a bunch of bananas, and McCarthy took that as a base level for intelligent reasoning.

McCarthy began to consider the possibility of constructing intelligent computers. This was all at a rather abstract level, since at that point, no computers were available to him. He conceived of two interacting finite automata, one representing the human brain, and one representing the environment. When he arrived at Princeton in 1949, he spoke to John von Neumann about his ideas; von Neumann was interested and told him to write them up. But McCarthy, upon further reflection, was dissatisfied with his original idea. He realized that he needed to somehow represent the knowledge in the human brain, and that a pair of finite automata, even if they could accurately represent the behavior of a human brain interacting with the environment, didn't have enough structure to represent human knowledge. Although he hadn't yet named it, McCarthy had at this point already identified the need for knowledge representation. Twenty years later, McCarthy was to formalize this intuition when he made the distinction between a metaphysically adequate representation and an epistemologically adequate representation (McCarthy & Hayes 1969). (McCarthy's ideas on interacting automata were independently invented and developed by Larry Fogel some years later; see (Fogel 1962).)

McCarthy finished his Ph.D. in two years, on a problem in partial differential equations that he had discovered. He stayed on for a couple of years at Princeton as an instructor. It was during this time that he met Marvin Minsky, who began his graduate studies at Princeton just as McCarthy was starting his instructorship. McCarthy and Minsky discovered a shared passion for research into constructing intelligent machines; the two were to collaborate on many projects over the next decade. In 1952, McCarthy spent a summer working at Bell Labs. He approached Claude Shannon with the idea of collecting papers on the topic of intelligent machines. Claude Shannon, in addition to his seminal contributions in information theory, analog circuit design, and computational linguistics, had recently published a groundbreaking paper on how one could program a machine to play chess. Shannon agreed to collaborate on such a collection, but, recalls McCarthy, didn't want the title to be too provocative. As a result, they called the planned volume *Automata Studies* (Shannon & Mc-

Carthy 1956). As the papers began to come in, McCarthy (at this point an acting assistant professor at Stanford) was disappointed when he realized that the papers, while all related to automata studies in some way, had little to do with what he regarded as getting machines to do intelligent reasoning.

So he decided to nail the flag to the mast and be explicit about the sort of research and the sort of papers that he wanted to encourage. He therefore came up with “artificial intelligence” — a term that would make it clear that the goal was to construct machines that acted in a truly intelligent manner. It was this term that he used in 1955 when he began writing — with Minsky, Shannon, and Nathaniel Rochester, a researcher at IBM — the proposal to fund the first conference dedicated to the topic, the famous Dartmouth conference on artificial intelligence.

The Dartmouth conference — held where McCarthy was then on the faculty of mathematics — took place in the summer of 1956, and is often considered to be the start of organized research into artificial intelligence. One doesn’t create a field merely by naming it, of course, and McCarthy discovered that, as with the *Automata Studies* volume, there were some participants who came and spoke about their pet projects, whether or not they had anything to do with artificial intelligence. Still, this served as a way of getting four researchers who were doing work in the field — McCarthy, Minsky, Newell, and Simon — to meet and talk and plan for future research projects in artificial intelligence. The main accomplishment of the Dartmouth conference was not any particular idea or approach to AI, but the commitment of four researchers toward defining a discipline of artificial intelligence, and the bonds created between these colleagues.

The MIT Years

While McCarthy was at Dartmouth, John Kemeny, then chairman of Dartmouth’s math department, arranged for McCarthy to receive a one-year Sloan fellowship. This allowed him to spend the 1956-1957 academic year at any institution of his choosing. He decided on MIT, because it had access to an IBM computer. (The computer was physically located at MIT, but MIT had access to it for only one eight-hour shift each day. One eight-hour shift was reserved for other New England colleges and universities, and one eight-hour shift was reserved for IBM’s own use.) MIT also afforded close proximity to his colleague Marvin Minsky, who was at that point at Harvard. MIT offered McCarthy a faculty position once the Sloan fellowship ended, and McCarthy never returned to his Dartmouth post.

The six years that McCarthy spent at MIT were one of the most productive periods in his life. During this period, he conceived of the idea of timesharing, suggested using automated theorem proving for program verification (McCarthy 1963a), and invented Lisp. He also began what was to become his lifelong research

project: figuring out how to get computers to do commonsense reasoning. The work was radically different from that of his AI colleagues in at least two ways. First, previous work in AI had focused on getting a computer to replicate activities that are challenging for humans, such as playing chess and proving theorems of mathematics. In contrast, McCarthy was concerned with mundane and seemingly trivial tasks, such as constructing a plan to get to the airport. Second, he was proposing using the tools of mathematical logic for proving something other than theorems in mathematical domains. His paper, “Programs With Common Sense” (McCarthy 1959) (often referred to as the Advice-Taker paper), articulated the need for a computer to be able to perform commonsense reasoning and the need for a computer to have a formal representation of the commonsense knowledge that people use when they go about their everyday reasoning activities; argued that the representation of such knowledge, along with an inference method to reason with this knowledge, was an essential part of any artificial intelligence; and advocated for the feasibility of the project. Indeed, the presentation of this paper may be seen as the birth of the field of knowledge representation.

The paper generated much controversy when it was first given in Great Britain. Philosopher and linguist Yehoshua Bar-Hillel (Bar-Hillel, McCarthy, & Selfridge 1998) was concerned that McCarthy’s proposed project was infeasible. Bar-Hillel argued, among other things, that the example of formal commonsense reasoning given in the paper was oversimplified, and that any proper formalization of an example would necessarily be much longer and more complex. In retrospect, Bar-Hillel turned out to be correct on this point; the formalization of commonsense reasoning has proved to be a remarkably difficult enterprise. Nearly half a century after the presentation of that paper, researchers are still grappling with many of the underlying difficulties.

This was also the period in which McCarthy and Minsky (who had arrived at MIT in 1958) established the MIT AI Lab. As McCarthy tells this classic tale of ask-and-you-shall-receive, he and Minsky were talking in the hallway when he saw Jerome Weisner, then acting head of MIT’s Department of Electrical Engineering, walking down the hall. McCarthy buttonholed him and said, “Marvin and I want to have an artificial intelligence project.” Weisner asked, “Well, what do you need?” Minsky said, “A room, a secretary, a key punch, and two programmers.” Weisner responded, “How about 6 graduate students?” (Funding for these was available due to a recent prize awarded to the math department.) And thus the AI Lab was born.

McCarthy and Minsky did not always see eye-to-eye on their respective approaches to artificial intelligence. McCarthy became increasingly committed to the logicist approach to AI; Minsky came to believe that it was wrong-headed and infeasible (Minsky 1975). Nonetheless, both approaches continued to progress during this period.

The Stanford Years

In 1962, McCarthy received an offer of a faculty position from Stanford University's Department of Mathematics. McCarthy had worked there previously, in the early 1950s. Of three acting assistant professors, he was the one person whose contract had not been renewed. (This had precipitated his move to Dartmouth.) At that point an associate professor at MIT, he told the department chairman, George Forsyth, that he would accept no less than a full professorship. Forsyth, according to McCarthy, had quite a job convincing the Stanford administration to agree to hiring as full professor someone whose contract hadn't even been renewed less than a decade before, but he was enough of a visionary to realize the importance of computer science and artificial intelligence. In any case, McCarthy was only briefly in the Department of Mathematics. Stanford soon created a Department of Computer Science, and McCarthy was one of that department's original members.

Research in Knowledge Representation

It was at Stanford that McCarthy began working out, and has continued up to the present time to work out, most of his research in formal knowledge representation. Most of his research falls into four areas:

1. Reasoning about actions. This includes McCarthy's original work on situations (McCarthy 1963b), culminating in the situation calculus (McCarthy & Hayes 1969), as well as more recent extensions (McCarthy 2002); the discovery of several challenging problems for knowledge representation, such as the frame and qualification problems; and initial solutions to those problems.
2. Nonmonotonic reasoning. This includes McCarthy's development of domain circumscription (McCarthy 1977), predicate circumscription (McCarthy 1980), and formula circumscription (McCarthy 1986).
3. Issues related to reification. McCarthy has been a proponent of using the technique of *reification*, in which sentences and other complex constructs of first-order logic are mapped to *terms* of first-order logic in order to enable formalization of commonsense reasoning within first-order logic. This work includes his ongoing interest in contexts (McCarthy 1990; 1993; McCarthy & Buvac 1997).
4. Reasoning about knowledge. McCarthy's work in this area includes the discovery of the knowledge preconditions problem (McCarthy & Hayes 1969) and work on variations of classic puzzles involving common knowledge, such as the Three Wise Men or Muddy Children problem (Fagin *et al.* 1995).

These areas overlap and influence one another. For example, as discussed below, the frame and qualification problems were the primary impetus for the development of circumscription.

The problems that McCarthy has chosen to work on all flow from his general goal of formalizing commonsense reasoning. Indeed, much of it can be seen to stem from the small example of commonsense reasoning that McCarthy discussed in his 1958 Advice Taker paper, that of a person planning to go to the airport while sitting at his desk, given that he has a car in his garage. McCarthy realized early on that virtually all commonsense reasoning involves reasoning about action and change (or the absence thereof). The invention and development of the situation calculus was meant to provide a framework that facilitated such reasoning.

The detailed formalization, within the situation calculus, of even simple problems — such as planning to build a tower of blocks — led to the discovery of a host of other problems. What Bar Hillel had feared in 1958 — that the formalization of enough knowledge to solve simple commonsense reasoning problems was, indeed, a difficult task — became evident. For example, McCarthy (working with Pat Hayes) realized that a simple plan to build a tower of three blocks raised unforeseen difficulties. If three blocks, A, B, and C, lay on a table, it seemed obvious that a plan to put B on C and A on B would succeed in building a tower. But when one formalizes this small problem, the question arises: how does one know that A is still on the table after placing B on C? This is an instance of the *frame problem*, which concerns the ability to represent and reason efficiently about what stays the same and what changes as actions are performed.

A related problem is the qualification problem, which concerns the ability to reason about the many conditions that must be true in order for an action to be performed successfully. For example, in the Missionaries and Cannibals problem, in which one reasons about how 3 cannibals and 3 missionaries can use 3 boats to safely cross a river, there are certain implicit qualifications in the rules stated for the problem, such as the fact that the boats do not leak and have oars, and that there is no bridge.

These problems eventually led McCarthy to believe that one must be able to go beyond classical logic, in which one reasons about things that are always true, to a default or nonmonotonic logic, in which one reasons about things that are typically true. For this purpose, McCarthy introduced *circumscription*, an extension of classical logic which allows one to prefer certain models of a theory.

McCarthy developed several increasingly powerful theories: *domain circumscription*, in which one prefers models which have minimal domains; *predicate circumscription*, in which one prefers models where certain predicates have minimal extensions; and *formula circumscription*, in which one prefers models where certain formulas have minimal extensions.

There were several other researchers investigating techniques for nonmonotonic reasoning during the late 1970s and early 1980s, including Ray Reiter (Reiter 1980) and Drew McDermott and Jon Doyle (McDer-

mott & Doyle 1980). McCarthy’s work can be distinguished in two ways. First, he aimed to stay as close as possible to classical logic. The principle of circumscription can be represented as a first-order axiom schema or as a second-order axiom. In contrast, McDermott and Doyle used a modal logic, and Reiter introduced an entirely new inference rule. Second, McCarthy has sought to tie his work in nonmonotonic reasoning to the specific applications in which he has been interested.

For example, McCarthy showed how to use circumscription over a set of abnormality predicates to formalize inheritance. To formalize, e.g., the facts that things typically don’t fly, but that birds typically do fly, but that, on the other hand, penguins typically do not fly, one could write the theory:

$$\begin{aligned} \forall x (Thing(x) \wedge \neg ab1(x) \Rightarrow \neg Flies(x)) \\ \forall x (Bird(x) \wedge \neg ab2(x) \Rightarrow Flies(x)) \\ \forall x (Penguin(x) \wedge \neg ab3(x) \Rightarrow \neg Flies(x)) \\ \forall x (Bird(x) \Rightarrow Thing(x)) \\ \forall x (Bird(x) \Rightarrow ab1(x)) \\ \forall x (Penguin(x) \Rightarrow Bird(x)) \\ \forall x (Penguin(x) \Rightarrow ab2(x)) \end{aligned}$$

If one adds the facts

Thing(Sylvester)

Bird(Tweety)

Penguin(Opus)

and circumscribes the predicates *ab1*, *ab2*, and *ab3*, one gets the desired result that Sylvester and Opus do not fly, while Tweety does fly.

McCarthy also suggested using circumscription to handle the frame problem, by formulating the principle of inertia — that fluents typically don’t change when actions are performed — using abnormality predicates. His formulation turned out to be overly simplistic, and as a result, led to incorrect conclusions; see, e.g., the blocks-world example of Lifschitz (Lifschitz 1986) and the Yale Shooting problem (Hanks & McDermott 1987). The difficulty is that the principle of inertia can apply to multiple *fluents*, or properties, not all of which can be simultaneously minimized; minimizing change can thus lead to multiple models, some of which are unintuitive. This difficulty does not arise in theories in which there is both an explicit theory of causation and a correspondingly more realistic formalization of the principle of inertia (e.g., (Lifschitz 1987)).

From the standpoint of logicist AI, the development of formal theories such as circumscription that enabled nonmonotonic reasoning was crucial for the survival of the logicist agenda. Critics had correctly pointed out that much commonsense reasoning was nonmonotonic in nature and could not be formalized within a pure first-order logic. Showing that nonmonotonic reasoning could be formalized within an extension of first-order logic provided evidence that the logicist agenda was in fact feasible. Indeed, much of McCarthy’s work during his years at Stanford has focused on showing that the concerns that others have had about the logicist agenda can be addressed in formal logic.

Leadership and Recognition

During his years at Stanford, John McCarthy advised more than 30 Ph.D. students; he is listed in the Mathematics Genealogy Database, at <http://www.genealogy.ams.org>, as having 175 academic descendants.¹ The number of researchers who have collaborated with him or who have been influenced, through personal contact with him, or by his writings and his vision, is considerably greater. For example, Lifschitz’s body of work on circumscription from the 1980s and 1990s was directly influenced by his contact with McCarthy during his years at Stanford; Reiter’s book *Knowledge in Action* (Reiter 2001) is rooted in the reworking and extension of McCarthy’s original situation calculus; and Bob Kowalski cites his early work on the Event Calculus (Kowalski & Sergot 1986) as being heavily influenced by McCarthy’s situation calculus. The centers of logicist AI today, at the University of Toronto, at the University of Texas at Austin, at Linköping University in Sweden, at Imperial College in London, and at many other universities around the world, owe much to McCarthy’s groundbreaking ideas.

McCarthy’s influence is due not only to the strength of his ideas but also to his personal qualities. He is the ultimate optimist, and his belief in the power of formal logic is infectious. His generosity of spirit has nurtured many a new researcher in the area of logicist AI. His philosophy is to let as much research as possible flourish; his delight in hearing about new work in the field is evident.

Although McCarthy has remained mostly at the sidelines of academic politics, he has been active in organizations close to his research interests. He founded SAIL, the Stanford AI Laboratory shortly after he came to Stanford, and served as president of AAAI from 1984-1985. In 1991, hoping to reverse the trend of logicist AI toward producing metalevel results rather than object-level theories, he founded and organized the first Symposium on Logical Formalizations of Commonsense Reasoning. The symposium is now held every two years.

In addition, McCarthy has used his academic position to further humanitarian causes, particularly during the years when the Soviet Union existed. McCarthy, who learned Russian from foreign language records in his twenties, and, as a graduate student, translated a text in differential equations from Russian into English, made several visits to the Soviet Union starting in 1965. In 1968, he taught for two months in Akademgorodok, on Novosibirsk’s outskirts, and in Novosibirsk itself. In 1975, he was instrumental in getting cybernetics researcher and refusenik Alexander Lerner permission from Soviet officials to attend and talk at the 4th International Joint Conference on Artificial Intelligence

¹This is likely an underestimate, since several of McCarthy’s Ph.D. students are not listed in the database; moreover, in general not all an individual’s academic descendants are reported to this database.

(IJCAI) in Tbilisi, Georgia. In the 1980s he smuggled a fax and copier machine to linguist and Soviet dissident Larisa Bogoraz.

John McCarthy's many awards include the Turing Award (1971), the first IJCAI Award for Research Excellence (1985), the Kyoto Prize (1988), the National Medal of Science (1990), and the Benjamin Franklin Medal in Computer and Cognitive Sciences (2003). He is a founding fellow of AAAI (1990), and a member of the American Academy of Arts and Sciences (1974), the National Academy of Engineering (1987), and the National Academy of Sciences (1989).

Retirement

John McCarthy officially retired from Stanford on January 1, 2001. The last of his Ph.D students, Eyal Amir and Aarati Parmar Martino, defended their dissertations in 2002 and 2003, respectively.

From McCarthy's point of view, retirement has meant more time for research and writing papers. He has been active in research for nearly sixty years, and he has been grappling with the central problems of artificial intelligence and commonsense reasoning for more than half a century. But he founded a difficult field of research. It is no easy task to imbue a machine with enough knowledge to do commonsense reasoning. Fifty years have not been nearly enough to finish the job, so McCarthy — along with the many AI researchers who, inspired by his early vision, now follow his logicist philosophy — continues to work at it.

APPENDIX (can this be put in as a sidebar?)

Reading John McCarthy: The Top Eleven — A Selected Annotated Bibliography

Listed below are eleven of John McCarthy's papers, selected for their enduring technical importance and/or their historical interest, along with a precis and comments. Most of John McCarthy's works are available on his website, at <http://www-formal.stanford.edu/jmc>. Six of the papers in this list — 2, 3, 4, 5, 6, and 7 — can also be found in McCarthy's collected papers on commonsense reasoning (Lifschitz 1990).

1. John McCarthy: "Recursive Functions of Symbolic Expressions and Their Computation By Machine," *Communications of the ACM*, 1960. This is the original paper describing Lisp. It introduces a methodology in which a program in a language could also be considered data for a(nother) program in that language, a capability essential for reflective reasoning. Many of the examples used in standard Lisp textbooks come directly from this paper.

2. John McCarthy: "Programs With Common Sense," *Proceedings, Teddington Conference on the Mechanization of Thought Processes*, 1958.

This paper was groundbreaking first, for proposing that commonsense reasoning (e.g., figuring out a way to get to the airport) was as worthy of study as seemingly difficult intellectual problems (e.g. playing chess); second, for suggesting that commonsense knowledge and reasoning could be expressed in first-order logic; and third, in recognizing the centrality of action and causality in commonsense reasoning. McCarthy deliberately avoided the representation of time in this paper. He believed that the proper representation of time would be isomorphic to the real numbers, but that knowledge of the theory of real numbers, developed only in the 19th century, could hardly be considered part of commonsense reasoning.

3. John McCarthy: "Situations, Actions, and Causal Laws," unpublished Stanford Technical Report, 1963. (This paper can be found, tacked onto the end of "Programs With Common Sense," in Marvin Minsky's *Semantic Information Processing*, MIT Press, 1968.) This is the first paper on the situation calculus (though it was not yet so named), a language that facilitates representation of actions and change. McCarthy had come to realize that he needed to distinguish fluents, properties that change over time. The notion of a situation, a (complete) description of the world at a particular moment in time, allowed him to represent fluents but still not explicitly mention time.

4. John McCarthy and Patrick J. Hayes: "Some Philosophical Problems from the Standpoint of Artificial Intelligence," *Machine Intelligence 4*, 1969. The central argument of this paper is that in order to successfully develop programs that can reason intelligently, we will have to deal with many of the problems that have concerned philosophers for centuries; and that, moreover, the demands of AI can put a new spin on old problems. The major contributions of the paper are first, a complete presentation of the situation calculus; second, a discussion of several areas of philosophy, including logics of knowledge, modal logics, and tense logics, which are relevant for artificial intelligence; and third, the presentation of several crucial problems for knowledge representation. These problems include the *frame problem*, the problem of efficiently determining what remains the same in a changing world, and what later came to be called the *knowledge preconditions problem*, the problem of determining what an agent (or set of agents) needs to know in order to perform an action.

5. John McCarthy: "Epistemological Problems of Artificial Intelligence," *Proceedings, IJCAI 1977*. This paper presents a number of problems and themes that have become central to research in knowledge representation. The paper includes the first discussion of the *qualification problem*, which concerns how one can represent the implicit qualifications in our general knowledge. It also includes one of the first discussions of circumscription, an extension of first-order logic which allows one to prefer certain models of a first-

order theory. Circumscription was developed in order to permit nonmonotonic reasoning, that is, the ability to jump to conclusions and later retract some of those conclusions.

6. John McCarthy, “Circumscription: A Form of Nonmonotonic Reasoning,” *Artificial Intelligence*, 1980. This paper gives a detailed account of predicate circumscription, the most commonly used form of circumscription. Predicate circumscription minimizes the extension of certain predicates; that is, models with minimal extensions of those predicates are preferred. The circumscriptive formula is given as a first-order schema.

7. John McCarthy, “Applications of Circumscription to Formalizing Common Sense Knowledge,” *Artificial Intelligence*, 1986. This paper extends the idea of predicate circumscription to the more general form of formula circumscription. It presents a second-order circumscriptive axiom, rather than the first-order schema originally used. The paper suggests several ways of applying circumscription to commonsense reasoning problems, such as inheritance and the frame problem.

8. John McCarthy and Sasa Buvac, “Formalizing Context: (Expanded Notes),” *Computing Natural Language*, 1997. This is the third and most comprehensive of four papers that McCarthy wrote, over the course of a decade, about the idea of reasoning using contexts. McCarthy sought to model the ease with which people reason relative to particular contexts. For example, one can reason about facts that are true in the context of a Sherlock Holmes story (e.g., Holmes’s feud with Moriarty) although they are not true in the actual world. Various technical mechanisms are introduced to deal with contexts, such as *lifting formulas*, which relate propositions and terms in subcontexts to more general propositions and terms in outer contexts. Lifting formulas are needed to combine or *transcend* contexts.

One of McCarthy’s Ph.D. students during this period, R.V. Guha, who wrote his thesis (Guha 1991) on reasoning with contexts, incorporated contexts into the Cyc knowledge base (Lenat & Guha 1990). They remain an important part of the Cyc project.

9. John McCarthy, “Elaboration Tolerance,” *Working Papers, Commonsense-1998*. A commonsense theory typically contains general commonsense information as well as a formalization of specific circumstances. A theory is said to be *elaboration tolerant* if modifying information about specific circumstances requires making only moderate changes to the general commonsense theory. This paper discusses various types of elaborations, using the Missionaries and Cannibals problem as a motivating example.

10. Tom Costello and John McCarthy, “Useful Counterfactuals,” *Electronic Transactions of Artificial Intelligence*, 1999. A *useful counterfac-*

tual is one which, if true, potentially teaches a useful lesson: contrast the useful counterfactual *If a car had been approaching while you switched lanes, you would have crashed* with the useless counterfactual *If wishes were horses, beggars would ride*. An agent can learn from useful counterfactuals without experiencing adverse events; and can use them to do hypothetical reasoning while constructing a plan. This paper discusses the evaluation of such useful counterfactuals, using the notion of a *Cartesian counterfactual*, which allows giving a precise meaning to Lewis’s and Stalnaker’s (Lewis 1973; Stalnaker 1984) imprecise but intuitively appealing notion of a *most similar* possible world.

11. John McCarthy, “Actions and Other Events in Situation Calculus,” *Proceedings, KR-2002*. This paper extends the original McCarthy-Hayes situation calculus in several significant ways, in order to handle natural events, concurrency, and combined linear and branching time.

References

- Bar-Hillel, Y.; McCarthy, J.; and Selfridge, O. 1998. Discussion of the paper: Programs with common sense. In Lifschitz, V., ed., *Formalizing Common Sense*. Intellect. 17–20.
- Costello, T., and McCarthy, J. 1999. Useful counterfactuals. *Electronic Transactions on Artificial Intelligence* 3(A):51–76.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning About Knowledge*. Cambridge, Massachusetts: The MIT Press.
- Fogel, L. J. 1962. Autonomous automata. *Industrial Research* 4:14–19.
- Guha, R. V. 1991. *Contexts: A Formalization and Some Applications*. Ph.D. Dissertation, Stanford University. Also available as technical report STAN-CS-91-1399-Thesis and as MCC Technical Report ACt-CYC-423-91.
- Hanks, S., and McDermott, D. V. 1987. Nonmonotonic logic and temporal projection. *Artificial Intelligence* 33(3):379–412.
- Jeffress, L. A. 1951. *Cerebral Mechanisms in Behavior: The Hixon Symposium*. New York: Wiley.
- Kowalski, R. A., and Sergot, M. J. 1986. A logic-based calculus of events. *New Generation Computing* 4(1):67–95.
- Lenat, D. B., and Guha, R. V. 1990. *Building Large Knowledge Based Systems: Representation and Inference in the Cyc Project*. Reading, Massachusetts: Addison Wesley.
- Lewis, D. 1973. *Counterfactuals*. Blackwell.
- Lifschitz, V. 1986. Pointwise circumscription: Preliminary report. In *AAAI*, 406–410.
- Lifschitz, V. 1987. Formal theories of action (preliminary report). In *IJCAI*, 966–972.

- Lifschitz, V. 1990. *Formalizing Common Sense: Papers by John McCarthy*. Ablex.
- McCarthy, J., and Buvac, S. 1997. Formalizing context: Expanded notes. In Aliseda, A.; van Glabbeek, R.; and Westerstahl, D., eds., *Computing Natural Language*. Stanford University. Also available as Stanford Technical Note STAN-CS-TN-94-13.
- McCarthy, J., and Hayes, P. J. 1969. Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, B., and Michie, D., eds., *Machine Intelligence 4*. Edinburgh: Edinburgh University Press. 463–502.
- McCarthy, J. 1959. Programs with common sense. In *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, 75–91. London: Her Majesty's Stationery Office.
- McCarthy, J. 1960. Recursive functions of symbolic expressions and their computation by machine. *Communications of the ACM* 3(4):184–195.
- McCarthy, J. 1963a. A basis for a mathematical theory of computation. In *Computer Programming and Formal Systems*. North-Holland.
- McCarthy, J. 1963b. Situations, actions, and causal laws. Technical report, Stanford University.
- McCarthy, J. 1977. Epistemological problems of artificial intelligence. In *IJCAI*, 1038–1044.
- McCarthy, J. 1980. Circumscription: A form of non-monotonic reasoning. *Artificial Intelligence* 13(1–2):23–79.
- McCarthy, J. 1986. Applications of circumscription to common sense reasoning. *Artificial Intelligence* 28(1):89–116.
- McCarthy, J. 1990. Generality in artificial intelligence. In Lifschitz, V., ed., *Formalizing Common Sense*. Ablex. 226–236.
- McCarthy, J. 1993. Notes on formalizing context. In *IJCAI*, 555–562.
- McCarthy, J. 1998. Elaboration tolerance. In *Working Papers of the Fourth International Symposium on Logical Formalizations of Commonsense Reasoning, Commonsense-1998*.
- McCarthy, J. 2002. Actions and other events in situation calculus. In Fensel, D.; Giunchiglia, F.; McGuinness, D.; and Williams, M., eds., *Proceedings of KR-2002*, 615–628.
- McDermott, D. V., and Doyle, J. 1980. Non-monotonic logic i. *Artificial Intelligence* 13(1-2):41–72.
- Minsky, M. 1975. A framework for representing knowledge. In Winston, P. H., ed., *The Psychology of Computer Vision*. McGraw-Hill. Also available as MIT-AI Lab Memo 306.
- Reiter, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13(1–2):81–132.
- Reiter, R. 2001. *Knowledge in Action*. Cambridge, Massachusetts: MIT Press.
- Shannon, C. E., and McCarthy, J. 1956. *Automata Studies*. Princeton University Press.
- Stalnaker, R. 1984. *Inquiry*. MIT Press.
- Turing, A. M. 1950. Computing machinery and intelligence. *Mind* 59:433–460.

Pat Hayes received a B.A. in mathematics from Cambridge University and a Ph.D. in Artificial Intelligence from Edinburgh. He has held academic positions at the University of Essex (England), the University of Illinois, and as the Luce Professor of cognitive science at the University of Rochester. He has been a visiting scholar at Universite de Geneve and the Center for Advanced Study in the Behavioral Sciences at Stanford, and has directed applied AI research at Xerox-PARC, SRI, and Schlumberger, Inc.. Pat has served as secretary of AISB, chairman and trustee of IJCAI, associate editor of Artificial Intelligence, a governor of the Cognitive Science Society, and president of AAAI. His current research interests include knowledge representation and automatic reasoning, especially the representation of space and time; the semantic web; ontology design; and the philosophical foundations of AI and computer science. He can be reached at phayes@ihmc.us.

Leora Morgenstern is a senior research scientist at the IBM T.J. Watson Research Center in Hawthorne, NY. She has taught at Brown University, Columbia University, and the University of Pennsylvania. She has a B.A. in mathematics and philosophy from the City College of New York, and a Ph.D. in computer science from the Courant Institute of Mathematics at New York University. Her research interests include reasoning about knowledge and belief, nonmonotonic reasoning, and multi-agent planning. She is particularly interested in applying knowledge representation techniques to commercial applications, and has extended techniques in semantic networks, nonmonotonic inheritance, and intelligent agents for the development of expert systems for benefits inquiry in medical insurance, and intelligent recommendation systems for banking, insurance, telephony, and sales. She can be reached at leora@steam.stanford.edu and leora@us.ibm.com.