



NATIONAL ACADEMY OF SCIENCES



JOHN
M^CCARTHY

1927-2011

A Biographical Memoir by

NILS J. NILSSON


© 2012 National Academy of Sciences

Any opinions expressed in this memoir are those of the author and do not necessarily reflect the views of the National Academy of Sciences.

JOHN MCCARTHY

September 4, 1927–October 24, 2011

BY NILS J. NILSSON



John McCarthy

JOHN MCCARTHY WAS BORN on September 4, 1927, in Boston, Massachusetts. John's father, John Patrick McCarthy, was an Irish Catholic immigrant. According to one biographer,¹ he "worked as a carpenter, a fisherman, and a union organizer," moving the family "from Boston to New York and then to Los Angeles" where John's father found work as an organizer for the Amalgamated Clothing Workers. John's mother, Ida Glatt, a Lithuanian Jewish immigrant, "worked as a journalist for *The Federated Press* wire service, then for a Communist newspaper and finally as a social worker." Both parents were active members of the Communist Party during the 1930s. John, along with his younger brother, Patrick, grew up in a politically charged household, but also one that encouraged learning

and critical thinking. In a recorded interview² John told me that even before he attended high school he got interested in science by reading a translation of a Russian popular science book for children, called *100,000 Whys*.

In Los Angeles, John attended and graduated from Belmont High School (two years early) in 1943. While in high school, he taught himself calculus from books used for freshmen and sophomores at nearby Caltech in Pasadena. In 1944 John was accepted as a mathematics undergraduate at Caltech, where he was able to skip the first two years of mathematics courses. A Wikipedia article³ about John states he “was reportedly expelled from Caltech for failure to attend physical education courses; he then served in the US Army and was readmitted to Caltech, receiving a B.S. in Mathematics in 1948.”

While taking some graduate courses at Caltech, John attended the Hixon Symposium on Cerebral Mechanisms in Behavior in September 1948. Several scientists and mathematicians at the intersection of mathematics, computation, and psychology gave talks there, and John became intrigued with the idea of developing machines that could think like people do.

After one graduate year at Caltech, McCarthy enrolled at Princeton University because he said that

it was the best place for mathematics, the subject he wanted to pursue. He told me⁴ that when he

arrived at Princeton in September of 1949, the [Mathematics] Department Chairman, Solomon Lefschetz, said to me, “What are you going to do your thesis on?” and I said, “I don’t know. I’ll think about it.” And so I immediately got myself a notebook, and started thinking about ideas for a thesis. Before the year was up, I had decided on a topic, and this made me relatively popular, because the usual thing is that the students pester the professors to suggest topics.

John’s topic involved a problem in partial differential equations. He finished his Ph.D. work under Lefschetz in 1951 with a dissertation entitled “Projection Operators and Partial Differential Equations.”

In addition to his thesis work at Princeton, McCarthy pursued his fascination with mechanizing intelligence. His early thoughts on that topic, actually forming while at Caltech, involved two interacting finite automata, one playing the role of a brain and the other playing the role of the environment. John von Neumann, of Princeton and the nearby Institute for Advanced Study, encouraged McCarthy to write up his ideas, but McCarthy thought that even if the “brain automaton” could be made to act intelligently, its internal structure wouldn’t be an explicit representation of human knowledge. He thought that somehow brains did explicitly

represent and reason about “knowledge,” and that’s what he wanted computers to be able to do.

McCarthy told one interviewer that he briefly joined the local Communist Party cell in 1949. It had two other members, a cleaning woman and a gardener. He quit the party soon afterward.

After finishing his Ph.D. work, McCarthy stayed on at Princeton for two years as an instructor. During the summer of 1952, he worked at Bell Labs, where he met Claude Shannon. They decided to collaborate on a volume of papers on automata. There were several submissions to the volume, which was published as *Automata Studies*.⁵ McCarthy was disappointed that most of the papers were on the theory of automata (including an important one by Stephen Kleene) and only one or two related at all to his chief interest, namely, machine intelligence.

It was at Princeton that McCarthy first met Marvin Minsky, who was beginning his graduate work there in 1951. They were both interested in mechanizing intelligence—one way or another. They collaborated over the next decade, but their approaches to making machines intelligent ultimately diverged widely.

In 1953 McCarthy returned to the west coast to take a position as an acting assistant professor of mathematics at Stanford University. Although he continued to think about machine intelligence, he claims he didn’t write anything about it during those early years at

Stanford. He continued his research on differential equations and analysis, but Stanford decided not to promote him. As he put it,⁶ “Stanford decided they’d keep two out of their three acting assistant professors, and I was the third.”

So, it was back across the country again. In February 1955 McCarthy took up a position as an assistant professor of mathematics at Dartmouth College in Hanover, New Hampshire. Even with his interest in things computational McCarthy didn’t actually try to program a computer until 1955. Around that time IBM decided to make an IBM 704 computer available for research and educational purposes. MIT, other New England colleges, and IBM itself would share its use. As the Dartmouth representative, McCarthy met Nathaniel Rochester, head of IBM’s Information Research Department in Poughkeepsie, New York. Rochester invited McCarthy to spend the summer of 1955 with his group at IBM.

At IBM that summer McCarthy and Rochester persuaded Claude Shannon and Marvin Minsky, then a Harvard junior fellow in mathematics and neurology, to join them in proposing a workshop to be held at Dartmouth during the following summer. McCarthy took the lead in writing the proposal and in organizing what was to be called a “Summer Research Project on Artificial Intelligence.” The proposal⁷ was submitted to the Rockefeller Foundation in August 1955. Among other things the proposal stated that the



McCarthy in December 1965.
Chuck Painter/Stanford News Service.

study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it ...For the present purpose the artificial intelligence problem is taken to be that of making a machine behave in ways that would be called intelligent if a human were so behaving.

Several people attended what turned out to be a rolling workshop at Dartmouth in the summer of 1956. McCarthy proposed the name “Artificial Intelligence” for the workshop because, as he said, “I had to call it something, so I called it ‘Artificial Intelligence,’ and I had a vague feeling that I’d heard the phrase before, but in all these years I have never been able to track it down.”⁸ McCarthy said later that not much was really accomplished at the workshop. But as Hayes and Morgenstern wrote,⁹

[It] served as a way of getting four researchers who were doing work in the field—McCarthy, Minsky, [Allen] Newell, and [Herb] 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.

Although others had written earlier about mechanizing intelligence (notably, Alan Turing in 1950), the Dartmouth workshop is generally considered to be the beginning of serious work in the field—a field to which McCarthy devoted his scientific career.

While at Dartmouth, McCarthy became a Sloan fellow in physical science (arranged by John Kemeny, then the chair of Dartmouth’s Department of Mathematics). He chose to spend his fellowship at MIT where he had access to an IBM computer and could interact with Marvin Minsky, then at nearby Harvard University. He never returned to Dartmouth, instead becoming an assistant professor of communication science at MIT in 1958.

McCarthy spent the summer of 1958 at IBM, and it was there that he first began to recognize the need for a new programming language—one that could support recursion and dynamic storage. (In recursive languages a program can invoke a [usually] simpler version of itself.) When he returned to MIT in the fall of 1958, he began work on a new language, one he called “LISP” (for “list processor”). Besides recursion, programs written in LISP could process arbitrary symbolic structures and could treat programs (written in LISP) as data (expressed as lists in LISP). (McCarthy said that he got the idea of list processing from Newell and Simon at the 1956 Dartmouth workshop, but that he didn’t like the language,

IPL, they were using.) John's 1960 paper "Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part I," established the theoretical foundations of LISP as a universal computational formalism.¹⁰ (Part II was never produced.) Comparing LISP to a universal Turing machine, McCarthy claimed that LISP was much more transparent in structure. LISP soon became the language of choice for AI research. Programs written in LISP have flown in a NASA spacecraft and are key parts of several practical AI systems.

In 1957 and 1958 at MIT, McCarthy provided the first suggestions for implementing timesharing—allowing a single computer to be simultaneously shared among several users.¹¹ These ideas eventually were embodied in 1962 in MIT's Compatible Time-Sharing System (CTSS) and also in later systems that McCarthy helped build at Bolt, Beranek and Newman and at Stanford University. CTSS led directly to the creation of Project MAC, which revolutionized computing at MIT. Although computers ultimately became sufficiently inexpensive that most people did not have to share computers, timesharing technology is still used to allow several programs of a single user to be simultaneously active.

In 1959 John invented a technique he called "garbage collection," in which random access memory is freed up by removing code that subsequent computations won't need. That technique was added to LISP



McCarthy in 1966, playing chess against a computer.

Chuck Painter/Stanford News Service



McCarthy in 1974.

Chuck Painter/Stanford News Service

and is still routinely used in Java and other programming languages. McCarthy also made substantial contributions to the algebraic languages ALGOL 58 and 60.

Between 1959 and 1962 a group of MIT students, advised by McCarthy, developed a chess-playing

program. It was based on earlier programs for the IBM 704 written by McCarthy. One of the group members, Alan Kotok, described the program in his MIT bachelor's thesis.¹² The program, which played at the beginning amateur level, used a method devised by McCarthy called the alpha-beta procedure to speed up search. Here is John's description of how that procedure works,

The basic idea is this. Suppose I consider making a move, and then I consider that if I make this move, the opponent can capture my queen. Then I normally don't consider what else the opponent could do to me, because capturing the queen is a disaster sufficient that I shouldn't make the move that leads up to it."¹³

Even before the Dartmouth workshop McCarthy became interested in formulating a system that could be given statements (corresponding to English sentences) in some suitable language and then reason with these statements. At a 1958 conference in Teddington, England, he gave a paper titled "Programs with Commonsense" (often referred to as the Advice Taker paper) proposing just such a system.¹⁴ Although his proposal was called "half-baked" at the conference by the philosopher and linguist Yehoshua Bar-Hillel, its baking was to constitute significant parts of the subsequent AI literature and occupy the research lives of McCarthy and many others pursuing what has been called the "logician" approach to AI. "Indeed," Hayes and Morgenstern claimed "the

presentation of this paper may be seen as the birth of the field of knowledge representation."¹⁵

Complementing McCarthy's work in AI were his efforts at establishing a mathematical theory of computation. Although Turing and others had done earlier work on what could be computed and by what processes, McCarthy wanted to be able to treat computer programs as mathematical objects that could be proved "correct." That is, he wanted to prove mathematically that programs met the specifications laid down for them. He wrote three important papers on this topic, culminating in a paper describing the (first) proof of the correctness of a compiler.^{16,17,18}

In the fall of 1962 John was invited to join the Computer Sciences Division of the Mathematics Department at Stanford as a full professor. (He wrote on one of his Web pages that his "reason for having moved back to California from Massachusetts in 1962 was substantially because I prefer mild winters.") He joined Stanford's Department of Computer Science when it was formed in 1965 and started a new artificial intelligence project there. The project was funded by the Advanced Research Projects Agency (ARPA). During this time, he initiated the development of the first display-based timesharing system, called "Thor," which included many of the features found in modern personal computers.

McCarthy also continued to develop the chess program he had worked on at MIT, and in 1965 he challenged a group at the Moscow Institute for Theoretical and Experimental Physics to a match with their own program. Moves were exchanged by telegraph. Neither program did very well, but the Russian program won the match.

As McCarthy's various AI and related projects grew he formed the Stanford Artificial Intelligence Laboratory (SAIL). He was aided in running the lab by Lester Earnest, who initiated and worked on several projects there. With a series of contracts from ARPA the lab obtained a DEC PDP-6 computer, followed later by a PDP-10 and a KL-10.

For his own research at the lab McCarthy mainly followed up on the ideas first proposed in his earlier paper "Programs with Commonsense." He continued to hold that the knowledge needed by AI programs should be represented in declarative sentences (principally in a logical language) rather than being encoded within the programs that use that knowledge. As he put it, "Sentences can be true in much wider contexts than specific programs can be useful." He said, "For AI, logic is inevitable in the same sense that for physical systems calculus is inevitable."¹⁹

He and Patrick J. Hayes invented a formalism for reasoning about actions that they called the "situation calculus,"²⁰ in which a situation was a state of the world

and an action was regarded as a function that changed situations. Actions have preconditions and effects. Representing the preconditions faithfully involved what McCarthy called the "qualification problem." Representing the effects faithfully (which things change and which don't) involved what McCarthy called the "frame problem." Dealing with these problems in a logical setting required extensions to logic that allowed making reasonable guesses. McCarthy's principal contribution toward such extensions (aided by his colleague Vladimir Lifschitz) was a technique he called "circumscription."^{21,22} Circumscription allowed a kind of reasoning logicians called "non-monotonic." As Hayes and Morgenstern pointed out, "Showing that nonmonotonic reasoning could be formalized within an extension of first-order logic provided evidence that the logicist agenda was in fact feasible."²³ Most of McCarthy's AI work at Stanford focused on the problem of implementing commonsense reasoning, using first-order logic with these (and one or two other) extensions.

Although not personally involved in it himself, McCarthy encouraged research in robotics at SAIL. McCarthy's reason for this interest was that he thought previous research in computer vision focused primarily on character and object *recognition* rather than on scene *description*. As he put it,

The foundations for graphical user interfaces and printers, computer typesetting and publishing, speech recognition, computer vision and robotics, computer music, and other technologies that are now parts of our everyday lives all got their start at SAIL's facilities in the Stanford foothills.

My idea was that you wanted description, not just discrimination, and my argument for that was that if you wanted to program a robot to pick up something, then it didn't just have to discriminate on the whole picture. It had to be able to locate the object, and represent its shape, so I decided to try for robotics."²⁴

One of the most successful robotic achievements at SAIL was the assembly of a Model T Ford racing water pump using a vision system and an electromechanical hand.²⁵ Another robotics project at Stanford used a wheeled vehicle equipped with stereovision to navigate around various obstacles (very slowly!) in the lab.²⁶

During the late 1960s and 1970s, SAIL was a shining example of what a community of very bright faculty, students, and staff could accomplish when provided with adequate funding, powerful computers, peripheral equipment (such as display devices and printers) and associated software. Much of the equipment and software needed by the projects at SAIL was developed at SAIL. John's philosophy in "managing" SAIL was to let a thousand flowers bloom. He often gave good advice to people about how to tackle a problem, but if they did it in a different way and made it work, they got no hassles from him. The foundations for graphical user interfaces and printers, computer typesetting and publishing, speech recognition, computer vision and robotics, computer music, and other technologies that are now parts of our everyday lives all got their start

at SAIL's facilities in the Stanford foothills. SAIL was also one of the first nodes on the ARPAnet, a precursor to the modern Internet. Over time SAIL produced many Ph.D.s and other graduates. Sixteen Turing awards from the Association for Computing Machinery were given to people who had been affiliated with SAIL.

McCarthy was always modest about his assessment of progress in AI, thinking that many discoveries and inventions would need to be made before we would have AI programs that reached general human levels of competence in thinking and reasoning. In his 1973 review of an appraisal of AI by Sir James Lighthill, McCarthy acknowledged that

artificial intelligence research has so far been only moderately successful; its rate of solid progress is perhaps greater than most social sciences and less than many physical sciences. This is perhaps to be expected, considering the difficulty of the problem."²⁷

He believed that we would have to know much more about how human intelligence works before being able to duplicate it in machines, writing that "[unfortunately we] understand human mental processes only slightly better than a fish understands swimming."

John was a philosopher in the old Greek sense of the word: one who loves wisdom. Like the ancient Greeks debating philosophy in the market place, John's arena was the electronic marketplace of bulletin boards, newsletters,

Along with his work in AI and computer science, McCarthy was deeply interested in the connections between AI, philosophy, and cognitive science.

and the Web. Some of his Usenet interaction can be found in the archives maintained in Google groups.²⁸ He was particularly passionate about free speech. When the Stanford administration curtailed access on some Stanford computers to the jokes newsletter called *rec.humor.funny*, John organized a petition to reinstate it. The administration finally backed down. Much of John's prolific writing is contained in memos and other material on his many Web pages, which are all accessible from <http://www-formal.stanford.edu/jmc>.

Along with his work in AI and computer science, McCarthy was deeply interested in the connections between AI, philosophy, and cognitive science. He wrote on one of his Web pages,²⁹ "It turns out that many philosophical problems take new forms when thought about in terms of how to design a robot." For example, to achieve human-level performance a program must

have built into it a concept of what knowledge is and how it is obtained. . . [Furthermore,] its designers will need an attitude to free will [and if] it is to do meta-level reasoning about what it can do, it needs an attitude of its own to free will . . . [and] if the program is to be protected from performing unethical actions, its designers will have to build in an attitude about that.

Taking these considerations into account necessarily requires taking positions on epistemology, consciousness, and ethics.

McCarthy, along with most AI researchers and many philosophers, believed that all animals, including humans, were deterministic machines. Some people worry, for example, that machines could not have free will. According to McCarthy, free will involves considering different courses of action and having the ability to choose among them. He summarized his position about human free will by quoting his daughter, Sarah, who said at age four, "I can, but I won't." Even some chess programs have this kind of free will (even though they are completely deterministic). In two memos McCarthy claimed that useful robots would also need free will.³⁰

About consciousness, McCarthy was of the opinion that "thinking about consciousness with a view to designing it provides a new approach to some of the problems of consciousness studied by philosophers." In particular, "From the AI point of view, consciousness must be regarded as a collection of interacting processes rather than the unitary object of much philosophical speculation."³¹ But he also argued that "robots should not be programmed to have emotions or to behave so as to have emotion ascribed to them."³²

Although in his early life McCarthy had progressive and liberal views, he later became quite conservative. He argued with environmentalists over many matters, including nuclear energy and population. He strongly maintained that "human material progress is

desirable and sustainable.” On his Web pages devoted to that topic he writes about “energy in general, nuclear energy, solar energy, food supply, population, fresh water supply, forests and wood supply, global engineering, pollution, biodiversity, various menaces to human survival, the role of ideology in discussing these matters” and other problems. He loved to tweak people whom he suspected of being less than enthusiastic about technology and progress. If you were going to debate him on these issues you had to have your facts in order because he did.

McCarthy made several visits to the Soviet Union, learned to speak Russian, and developed friendships with several computer scientists there. Hayes and Morgenstern write,³³

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

John was scrupulously honest in his claims about his work. AI as a field has been accused occasionally of

overselling itself, but that certainly wasn’t McCarthy’s style. In stark contrast with the bold advances described in many of John’s papers and memos are many rather modest qualifications, such as “We don’t know a completely satisfactory way of doing this” and “—but we don’t presently see how to do it,” and “We are only part way to our goal of —”

McCarthy disliked bureaucracy. He thought that at least 98 percent of any bureaucracy could be eliminated with a consequent 98 percent reduction in its follies. John’s own enterprises reflected his view of a streamlined (or absent) administration. When he and Ralph Gorin set up the first timeshared computer system for all Stanford students to use in the mid-1980s, it was called LOTS, for Low Overhead Time-Sharing System. It was run by a part-time student or two.

John McCarthy’s many awards included the Association for Computing Machinery’s Turing Award (1971), the first International Joint Conference on Artificial Intelligence Award for Research Excellence (1985), the Kyoto Prize (1988), the National Medal of Science (1990), the Benjamin Franklin Medal in Computer and Cognitive Sciences (2003), and memberships in the American Academy of Arts and Sciences (1974), the National Academy of Engineering (1987), and the National Academy of Sciences (1989). He was a founding fellow and past president of the Association for

the Advancement of Artificial Intelligence. He was named the Charles M. Pigott Professor of Engineering by Stanford in 1987.

His many Ph.D. students, and the ones he advised, include (in alphabetical order) Eyal Amir, Ruzena Bajcsy, Randall Davis, Cordell Green, Ramanathan V. Guha, Donald Kaplan, Kurt Konolige, Barbara Huberman Liskov, Robert Moore, Francis Morris, Aarati Parmar Martino, David E. Wilkins, and Raj Reddy. Two of these (Liskov and Reddy) were subsequently honored with the ACM Turing Award.

About religion John wrote,

Responding to Richard Dawkins's pestering his fellow atheists to "come out," I mention that I am indeed an atheist. To count oneself as an atheist one need not claim to have a proof that no gods exist. One need merely think that the evidence on the god question is in about the same state as the evidence on the werewolf question.³⁴

John McCarthy died on October 24, 2011, from complications of heart disease at his home in Stanford, California. He is survived by his third wife, Carolyn Talcott of Stanford; two daughters, Susan McCarthy of San Francisco and Sarah McCarthy of Nevada City, California.; a son, Timothy Talcott McCarthy of Stanford; a brother, Patrick, of Los Angeles; two grandchildren, Kitty McCarthy of San Francisco and Joseph Gunther of New

York City; and his first wife, Martha Coyote. His second wife, Vera Watson, died in 1978 in a mountain-climbing accident attempting to scale Annapurna in Nepal.

John McCarthy's genius, puckish humor, and presence, along with his provocations to think more deeply, will be greatly missed by his colleagues, family, and many friends. Fortunately there is still much to be mined from his Web pages at <http://www-formal.stanford.edu/jmc/>.

NOTES

1. D. E. Shasha and C. A. Lazere. *Out of Their Minds: The Lives and Discoveries of 15 Great Computer Scientists*. New York: Springer, 1998.
2. Transcript of “Oral History of John McCarthy,” interview by Nils Nilsson on Sept. 12, 2007, at and for the Computer History Museum, Mountain View, Calif.
3. From [http://en.wikipedia.org/wiki/John_McCarthy_\(computer_scientist\)](http://en.wikipedia.org/wiki/John_McCarthy_(computer_scientist)).
4. Transcript of “Oral History of John McCarthy,” op. cit.
5. C. E. Shannon and J. McCarthy, eds. *Automata Studies* (AM-34, Annals of Mathematics Studies). Princeton: Princeton University Press, Apr. 1, 1956.
6. Transcript of “Oral History of John McCarthy,” op. cit.
7. From <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>. Portions of the proposal have been reprinted in J. McCarthy, M. L. Minsky, N. Rochester, and C. E. Shannon, A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, *AI Mag* 27(4)(2006):12.
8. Transcript of “Oral History of John McCarthy,” op. cit.
9. P. J. Hayes and L. Morgenstern. On John McCarthy’s 80th birthday, in honor of his contributions. *AI Mag* 28(4)(2007):93-102.
10. J. McCarthy. Recursive functions of symbolic expressions and their computation by machine. I. *Commun. ACM* 3(4)(1960):184-195.
11. McCarthy’s recollections about his suggestions for timesharing can be found at <http://www-formal.stanford.edu/jmc/history/timesharing/timesharing.html>.
12. A. Kotok. A Chess Playing Program for the IBM 7090 Computer. MIT bachelor’s thesis in electrical engineering, Jun. 1962.
13. Transcript of “Oral History of John McCarthy,” op. cit.

NOTES

14. J. McCarthy. Programs with commonsense. In *Proceedings of the Symposium on Mechanisation of Thought Processes*, vols. 1 and 2, eds. D. V. Blake and A. M. Uttley. London: Her Majesty’s Stationary Office, 1959.
15. P. J. Hayes and L. Morgenstern. On John McCarthy’s 80th birthday, in honor of his contributions, op. cit.
16. J. McCarthy. A basis for a mathematical theory of computation. In *Computer Programming and Formal Systems*, eds. P. Braffort and D. Hirschenberg. Amsterdam: North-Holland, 1963.
17. J. McCarthy. Towards a mathematical science of computation. In *Information Processing 62: Proceedings of IFIP Congress 1962*, pp. 21-28. (Also at <http://www-formal.stanford.edu/jmc/towards.pdf>)
18. J. McCarthy and J. Painter. Correctness of a compiler for arithmetic expressions. Mathematical aspects of computer science. 1. *Proc. Symp. Appl. Math.*, vol. 19. American Mathematical Society, 1967.
19. Transcript of “Oral History of John McCarthy,” op. cit.
20. J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, vol. 4, eds. B. Meltzer and D. Michie, pp. 463-502. Edinburgh: Edinburgh University Press, 1969.
21. J. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artif. Intell.* 13(1980):27-39.
22. J. McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artif. Intell.* 28(1986):89-116.
23. P. J. Hayes and L. Morgenstern. On John McCarthy’s 80th birthday, in honor of his contributions, op. cit.
24. Transcript of “Oral History of John McCarthy,” op. cit.

NOTES

-
25. R. Bolles and R. Paul. The use of sensory feedback in a programmable assembly system. Stanford AI Laboratory Memo AIM-220, Stanford Computer Science Department Report STAN-CS-396, October 1973.
26. H. P. Moravec. The Stanford cart and the CMU rover. *Proc. IEEE* 71(7) (1983):872-884.
27. J. McCarthy. Review of “Artificial intelligence: A general survey.” In *Formalizing Common Sense: Papers by John McCarthy*, ed. V. Lifschitz. Ablex Publishing, 1990. Preprint available online at <http://www-formal.stanford.edu/jmc/reviews/lighthill/lighthill.html>.
28. <http://groups.google.com/>.
29. J. McCarthy. What has AI in common with philosophy. Unpublished memo, Apr. 23, 1996.
30. See: J. McCarthy. Free will—even for robots, unpublished memo, Feb. 14, 2000. Simple deterministic free will, unpublished memo, May 16, 2002. Published online at www-formal.stanford.edu/jmc/freewill.html.
31. J. McCarthy. Todd Moody’s zombies. *J. Consciousness Stud.* 2(4)(1995).
32. J. McCarthy. Making robots conscious of their mental states. (Paper presented at Machine Intelligence 15 workshop, Oxford University, UK, July 24, 1995.) Published online at www-formal.stanford.edu/jmc/consciousness/consciousness.html.
33. P. J. Hayes and L. Morgenstern. On John McCarthy’s 80th birthday, in honor of his contributions, op cit.
34. About John McCarthy. Published online at www-formal.stanford.edu/jmc/personal.html.

SELECTED BIBLIOGRAPHY

-
- 1958
Programs with common sense. In “Mechanisation of Thought Processes.” *Proc. Symp. Natl. Phys. Lab.* 1:77-84 (London: Her Majesty’s Stationery Office).
- 1960
Recursive functions of symbolic expressions and their computation by machine. *Commun. ACM* 3(4):184-195.
- 1963
Situations, actions and causal laws. Memo 2. Stanford University Artificial Intelligence Project. (Reprinted in M. Minsky, ed., *Semantic Information Processing*, pp. 410-418. Cambridge, Mass.: MIT Press, 1968.)
- 1969
With P. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, vol. 4, eds. B. Meltzer and D. Michie, pp. 463-502. Edinburgh: Edinburgh University Press.
- 1977
Epistemological problems of artificial intelligence. *Proc. Fifth Int. Joint Conf. Artif. Intell.* (IJCAI-77). Los Altos, Calif.: William Kaufmann Inc.
- 1980
Circumscription—a form of non-monotonic reasoning. *Artif. Intell.* 13(1,2):27-39.
- 1986
Applications of circumscription to formalizing commonsense knowledge. *Artif. Intell.* 28(1):89-116.

SELECTED BIBLIOGRAPHY

1997

With S. Buvac. Formalizing context: Expanded notes. In *Computing Natural Language*, vol. 81, eds. A. Aliseda, R. J. van Glabbeek, and D. Westerstah, pp. 13-50, CSLI Lecture Notes, Stanford, Calif.: Center for the Study of Language and Information.

2002

Actions and other events in situation calculus. In *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR-02)*. San Francisco: Morgan Kaufmann Publishers.

Published since 1877, *Biographical Memoirs* are brief biographies of deceased National Academy of Sciences members, written by those who knew them or their work. These biographies provide personal and scholarly views of America's most distinguished researchers and a biographical history of U.S. science. *Biographical Memoirs* are freely available online at www.nasonline.org/memoirs.

