AI's Hall of Fame

In 2010, as the part of the celebration of the 25th anniversary of *IEEE Intelligent Systems* magazine, our editorial and advisory boards decided to launch the *IEEE Intelligent Systems* Hall of Fame to express our appreciation and respect for the trailblazers who have made significant contributions to the field of AI and intelligent systems and to honor them for their notable impact and influence on our field and our society.

When we first began our search for candidates, we did not think we would be so overwhelmed. It quickly became clear that there was an immense number of amazing, talented individuals conducting relevant and innovative research in the AI and intelligent systems field across the globe. The task of selecting from such an accomplished list was an extremely difficult process, and we proceeded with great care and consideration. I would like to express my sincere thanks to all the members of our editorial and advisory boards for their great effort in this endeavor.

It is always exciting to see that there are people with such passion in a field, and we hope that our Hall of Fame will be a way to recognize and promote creative work and progress in AI and intelligent systems.

Now, I proudly present the inaugural induction of the *IEEE Intelligent Systems* Hall of Fame. Congratulations to our first ever Hall of Fame recipients!

—*Fei-Yue Wang, Editor in Chief*

1541-1672/11/$26.00 © 2011 IEEE
Published by the IEEE Computer Society

HALL OF FAME

Edward Feigenbaum

Edward Albert Feigenbaum is a professor emeritus of computer science at Stanford University and a cofounder of three start-up firms in applied AI: IntelliCorp, Teknowledge, and Design Power. Often called the "father of expert systems," he founded the Knowledge Systems Laboratory at Stanford University, where the first expert system, DENDRAL, was developed. Feigenbaum completed his BS and a PhD at Carnegie Mellon University. In his 1959 PhD thesis, carried out under the supervision of Herbert Simon, he developed EPAM, one of the first computer models of how people learn. He helped establish Stanford's SUMEX-AIM national computer facility for applications of AI to medicine and biology. He has been a member of the Board of Regents of the National Library of Medicine and was a member of the Board of Directors of Sperry Corporation. He was the second president of the American Association for Artificial Intelligence. Feigenbaum has received the ACM Turing Award and the Feigenbaum Medal of the World Congress on Expert Systems. From 1994 to 1997, he served as US Air Force Chief Scientist, and he won a US Air Force Exceptional Civilian Service Award. He is a member of the National Academy of Engineering and American Academy of Arts and Sciences.

Engineering Knowledge

By Alun Preece

Arguably more than anyone, Ed Feigenbaum was responsible for getting AI out of the lab and into the enterprise. The key insight behind what he termed *expert systems* was to capture valuable problem-solving knowledge in a knowledge base, such that a machine could perform automated inference on it. Typically, the knowledge was captured by means of rules or descriptive frames. As a result, significant advances were made in techniques for rich knowledge and information representation and on efficient methods for large-scale inference. Indeed, one of Feigenbaum's key interests was the use of heuristic methods for solving computationally hard problems. Motivated in part to compete with Japan's "Fifth Generation Computer" project, Feigenbaum worked to alert companies and organizations to the potential benefits of knowledge-based systems.

Feigenbaum coined the term *knowledge engineering* to refer to the incorporation of knowledge into software systems in order to perform complex problem-solving tasks. Essentially, knowledge engineering represented the first widespread industrialization of AI. The AI field was transformed during the years of the expert systems boom, with significant investment both in basic research and application development. The limits of the technology were quickly identified, but nevertheless numerous useful systems were deployed at all levels within enterprises—some of which continue to be used today.

In 1994, at the tail-end of the expert systems boom, Feigenbaum's work was recognized—jointly with that of Raj Reddy—by an ACM Turing Award for "pioneering the design and construction of large-scale artificial intelligence systems, demonstrating the practical importance and potential commercial impact of artificial intelligence technology." As professor emeritus at Stanford, Feigenbaum has focused interest, as a Board of Trustees member of the Computer History Museum, on preserving the history of computer science, and with the Stanford Libraries on software for building and using digital archives.

Feigenbaum's legacy is broad. We can recognize patterns and characteristics of the technology in intelligent assistants and software "wizards" that step users through complex tasks. In the 1990s, there was a great deal of interest in technical aspects of knowledge sharing, by means of federated knowledge bases connected to the Internet. Organizations remained committed to the idea of capturing and leveraging their valuable knowledge, and expert systems technology became one element of the emerging sociotechnical knowledge management field. Effective heuristic problem-solving techniques remain an important tool within many enterprises, for example, for scheduling and resource-allocation tasks. We can also see components of knowledge-based systems technology and knowledge engineering methods in diverse areas such as business rules systems, database integrity management, and Semantic Web ontologies. Ideas from the expert systems field are also evident in recent high-profile launches such as IBM's Watson system and the Wolfram Alpha computational knowledge engine. Ed Feigenbaum will always be associated with the 1980s assertion that "knowledge is power." Thirty years later, a great many AI systems developers still agree.

Alun Preece is a professor of intelligent systems in the School of Computer Science and Informatics at Cardiff University, UK.

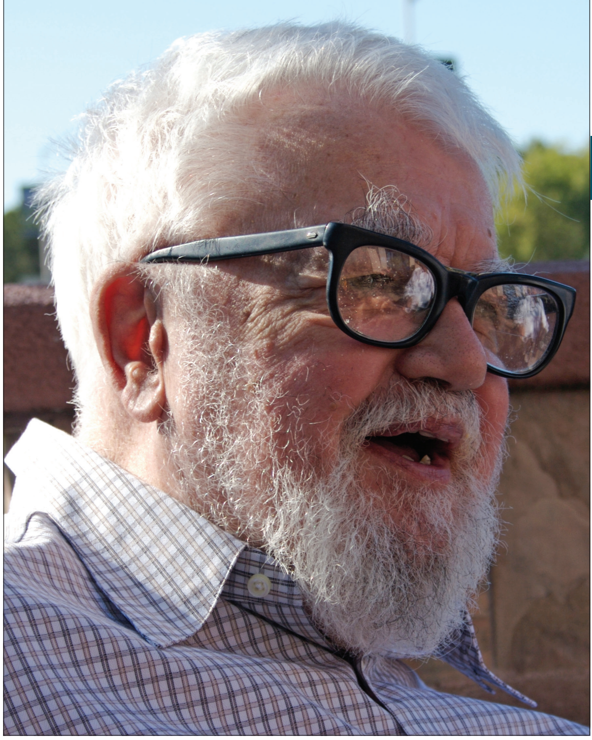

HALL OF FAME

John McCarthy

John McCarthy is an American researcher in computer science and pioneer in the field of AI. Now a professor emeritus at Stanford University, he was previously a professor of computer science, the Charles M. Pigott Professor in the School of Engineering, and director of the Artificial Intelligence Laboratory at Stanford University. McCarthy has a BS in mathematics from the California Institute of Technology and a PhD in mathematics from Princeton University. McCarthy was responsible for coining the term *artificial intelligence* in his 1955 proposal for the 1956 Dartmouth Conference. He founded MIT's AI laboratory (with Marvin Minsky) in 1957 and Stanford's AI laboratory in 1963. In 1958, he proposed the Lisp programming language. He was also one of the first to propose and design time-sharing computer systems, and he pioneered mathematical logic to prove the correctness of computer programs. From 1978 to 1986, he developed the circumscription method of nonmonotonic reasoning. He has received the Turing Award for Computing Machinery, National Medal of Science in Mathematics and Computer Science, Benjamin Franklin Medal in Computer and Cognitive Science, the first Research Excellence Award of IJCAI, and the Kyoto Prize of the Inamori Foundation. He is also a member of the American Academy of Arts and Sciences, National Academy of Engineering, and National Academy of Sciences.

Logic and Common Sense

By Ulrich Furbach

John McCarthy is one of the fathers of our discipline. He organized the famous Dartmouth Conference, where the foundations for the whole field were laid and the term AI was coined. Not only did he define the goals of AI, he also developed and provided a broad spectrum of key tools and methods.

His first seminal paper "Programs with Common Sense" was published before his famous Lisp paper and is still a pleasure to read about logical representation and reasoning from 1959. This was probably the first paper in which logic as a representation formalism was proposed. Other papers followed on this topic, and by 1970, he developed a modal-logic-based semantics of knowledge—an issue that is central in symbolic knowledge representation and in Web Science today.

John McCarthy continued the program he laid down with that paper, and in the 1980s, he came up with circumscription as a logical basis for common-sense reasoning. This work initiated various approaches to nonmonotonic reasoning in AI and logic programming; there are applications in most areas of AI. In his 1969 paper "Some Philosophical Problems from the Standpoint of Artificial Intelligence," together with Pat Hayes he included topics from robot action and planning. It even addressed the topic of free will. This is certainly the first detailed definition of situation calculus, which is used in so many variants for cognitive robotics. This paper also advocated a logic of knowledge by discussing various nonclassical logics. In subsequent papers, McCarthy explored the topics of robot consciousness and introspection, topics challenging for both robotics and knowledge representation.

In the 1950s, at a time when Fortran was just published and Algol was under construction, McCarthy developed the Lisp language based on symbolic expressions and symbolic evaluation. From the beginning, he designed it such that reasoning systems could make deductions about programs. It is a language of great elegance—small and yet powerful and defined with a clear and formal semantics. Lisp did not only become the major AI programming language, it also was the starting point for numerous papers on the mathematical foundation of computing. John McCarthy was probably the first to give a formal mathematical proof of a compiler for arithmetic expressions.

In addition to all these contributions to programming languages, logic, knowledge representation, and planning, John McCarthy published many papers discussing and shaping the field of AI. It is hardly possible for any researcher in AI not to stand on his shoulders!

Ulrich Furbach is a professor of artificial intelligence at the University of Koblenz, Germany.

HALL OF FAME

Marvin Minsky

Marvin Minsky is a pioneer for his work in AI and cognitive science, which includes key forays into mathematics, computational linguistics, robotics, and optics. He is currently the Toshiba Professor of Media Arts and Sciences and a professor of electrical engineering and computer science at the Massachusetts Institute of Technology. After serving in the US Navy from 1944 to 1945, Minsky received a BA from Harvard University and a PhD from Princeton University, both in mathematics. In 1959, he and John McCarthy founded what is now known as the MIT Computer Science and Artificial Intelligence Laboratory. In the early 1970s at the MIT Artificial Intelligence Lab, Minsky and Seymour Papert started developing what came to be called the Society of Mind theory. In 1985, Minsky published *The Society of Mind*, a comprehensive book on the theory written for a general audience. In 2006, Minsky published *The Emotion Machine*, which proposes theories that could account for human higher-level feelings, goals, emotions, and conscious thoughts in terms of multiple levels of processes, some of which can reflect on the others. He has received the ACM Turing Award, Japan Prize, and Benjamin Franklin Medal. He is also a fellow of the American Academy of Arts and Sciences and a member of the National Academy of Engineering and National Academy of Sciences.

Cognitive Science and AI

By Silvia Coradeschi

> We rarely recognize how wonderful it is that a person can traverse an entire lifetime without making a single really serious mistake—like putting a fork in one's eye or using a window instead of a door.
>
> —*The Society of Mind* (1987)

Marvin Minsky wrote this opening quote with his usual wit and insightfulness in one of his most popular books, *The Society of Mind*. In fact, he has always been fascinated by the fact that a five-year-old child can do things that machines are still unable to do and has focused his great talents on the challenging problem of how to give computers the human capacity for common-sense reasoning. His mechanistic view of the human mind inspired him to develop a theory in the 1970s attempting to show how intelligence could be a product of the interaction of many nonintelligent parts. In his owns words, "you mustn't look for one wonderful way to solve all problems. Instead you want to look for 20 or 30 ways to solve different kinds of problems. And to build some kind of higher administrative device that figures out what kind of problem you have and what method to use." Emotions are in his opinion simpler than intellectual responses, even if they are harder to describe. In his most recent book *The Emotion Machine*, he critiques many popular theories of how the human mind works and suggests alternative theories. His two books *The Emotion Machine* and *The Society of Mind* present together his conception of human intellectual structure and function.

He has also been one of the pioneers of artificial neural networks (ANNs) research, and in 1951, he designed the first NN simulator. He established the foundations of NN analysis in *Perceptrons*, a book he coauthored with Seymour Papert. Another great achievement is the formulation of the theory of frames that he presents in his book *A Framework for Representing Knowledge*. Frame theory is fundamental in *Knowledge Representation*, from both a practical and theoretical perspective, and it contributed to the emergence of the object-oriented programming paradigm. He also influenced the popular view of intelligent systems by being the adviser on the movie *2001: A Space Odyssey* and helping to make HAL the ultimate prototype of the intelligent machine.

Beside his great contributions to AI, he has also contributed to cognitive psychology, mathematics, computational linguistics, robotics, and optics. He designed and built early visual scanners, mechanical arms and hands, and other robotic devices; the Confocal Scanning Microscope; and the Muse synthesizer for musical variations (with E. Fredkin). He contributed to the development of the Logic-Oriented Graphic-Oriented (Logo) language and, in 1969, built the first Logo "turtle" device together with Seymour Papert.

His optimistic view that it is possible to build machines as smart as people and that our generation could be the one to realize them has been an inspiration for the entire AI field.

Silvia Coradeschi is a professor at Örebro University's Center for Applied Autonomous Sensor Systems.

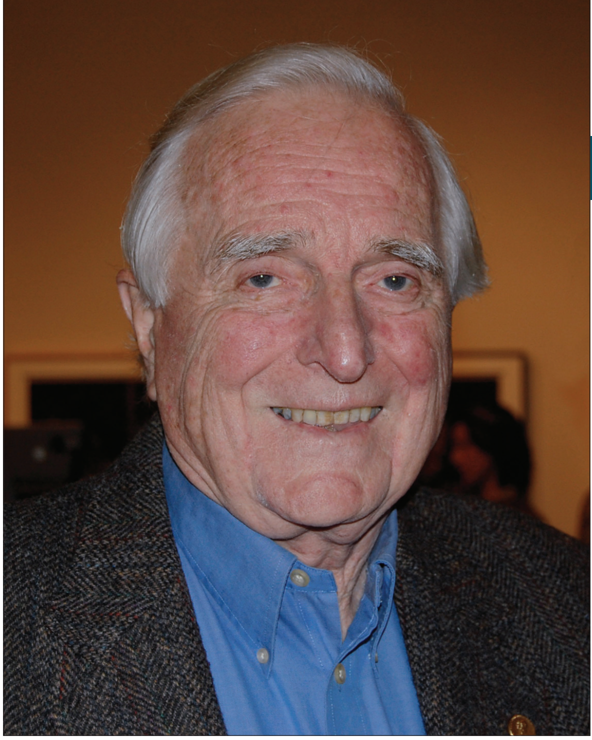

HALL OF FAME

Douglas Engelbart

Douglas Engelbart is an American inventor and early computer and Internet pioneer. He received a BS in electrical engineering from Oregon State University, after which he worked for the National Advisory Committee on Aeronautics (NACA) Ames Laboratory—the precursor to NASA. Subsequently, he earned a PhD in electrical engineering from the University of California, Berkeley. In 1963, Engelbart formed his own research lab at the Stanford Research Institute (SRI) that developed an elaborate hypermedia-groupware system called NLS (Online System). Pioneering firsts in human-computer interaction included a new point and click device (the mouse), natural language commands, and a GUI windowing environment. NLS also integrated word processing, hyper-email, computer-supported software engineering and meeting support, telecommuting, and a shared desktop; it also was first to demonstrate onscreen video teleconferencing. While developing NLS with ARPA funding, Engelbart was involved in the network ARPA was building and started the Network information Center (NIC) to support the first distributed online community. In 1989, Engelbart founded the Bootstrap Institute to foster high-performance organizations. He has received the ACM Turing Award, Lemelson-MIT Prize, the Benjamin Franklin Medal, and the National Medal of Technology. He is also a fellow of the American Academy of Arts and Sciences and a member of the National Academy of Engineering.

AI and Interactive Computing

By Susan B. Barnes

After reading Vannevar Bush's seminal *Atlantic Monthly* article "As We May Think" in 1945, Douglas Engelbart has spent decades creating better methods of human-computer interaction to advance what he described as our "collective IQ." Similar to Bush, Engelbart realized that the post-war world was becoming ever-more complex, and the problems we as a society would face would become increasingly complex and urgent at an alarming rate. He reasoned that radically new tools and strategies would be needed for effectively addressing these challenges and that computers could play an important role. His research began with the 1962 report "Augmenting the Human Intellect: A Conceptual Framework," funded by the Air Force Office of Scientific Research. The report described interactive computing, and he conceptualized tools to help intellectual workers make decisions. Specifically, Engelbart wanted to match human capabilities with computer tools. The point where humans and computers meet is called the graphical user interface.

In 1963, Engelbart set up the Augmentation Research Center (ARC) at the Stanford Research Institute (now SRI International). He was aware that moving from manual (pen, pencils, and paper) to automated symbol manipulation systems (computers) would alter how people think, perceive, conceptualize, assess, create, and deploy solutions collectively. To achieve this goal, a team of people with backgrounds in psychology, computer programming, computer engineering, and organizational development were assembled. In addition to developing the tools, the team also experimented and used the tools themselves, in the first major achievements in interactive knowledge work.

Throughout the 1960s and 1970s, Engelbart's ARC evolved the Online System (NLS) computer system, which facilitated the storage and retrieval of electronic documents and the development of digital libraries. The documents could be accessed from any connected terminal in the lab. A shared electronic space was created for software journal entries. GUIs were designed to help users interact with information displayed on the screen. Two new methods of interaction included a five-key, one-hand chord keyset and the mouse.

In 1968, Engelbart presented NLS in a live demonstration at the AFIPS Fall Joint Computer Conference in San Francisco. Even today, this demonstration is known at the "the mother of all demos" because it was so far ahead of its time. Engelbart showed interactive computing, the computer mouse, key chord, hypertext links, computer-supported cooperative work, videoconferencing, computer graphics, and word processing—all technologies that have become ubiquitous today.

The following year, SRI became the second site connected to the Arpanet. Due to its work on computer-supported cooperative work and hypertext, the ARC became an online clearing-house for Arpanet resources. In 1989, with his daughter Christina, Engelbart founded the Bootstrap Institute (now the Doug Engelbart Institute), which continues to support the development of enabling technologies for promoting collaboration.

Engelbart's interest in the synergistic relationship between humans and computers has played an important role in inspiring a generation of computer scientists. His 1968 demonstration paved the way for new interactive technologies that had not previously been imagined.

Susan B. Barnes is a professor and the associate director of the Lab for Social Computing at the Rochester Institute of Technology.

HALL OF FAME

Tim Berners-Lee

Tim Berners-Lee is the 3Com Founders Professor of Engineering in the School of Engineering with a joint appointment in the Department of Electrical Engineering and Computer Science at the Laboratory for Computer Science and Artificial Intelligence at the Massachusetts Institute of Technology, where he also heads the Decentralized Information Group (DIG). He is also a professor in the Electronics and Computer Science Department at the University of Southampton, UK, and the director of the World Wide Web Consortium (W3C). He is a founding director of the Web Science Trust (WST), which was launched in 2009 to promote research and education in Web Science, the multidisciplinary study of humanity connected by technology. He is also a director of the World Wide Web Foundation, launched in 2009 to fund and coordinate efforts to further the potential of the Web to benefit humanity. A graduate of Oxford University, Berners-Lee invented the World Wide Web at CERN, the European Particle Physics Laboratory, in 1989. He wrote the first Web client and server in 1990. His specifications of uniform resource identifiers (URIs), HTTP, and HTML were refined as Web technology spread. In 2009, Berners-Lee also advised the UK government's "Making Public Data Public" initiative. He is the author of *Weaving the Web*. He has received the Japan Prize. He is also a foreign associate of the National Academy of Science and the National Academy of Engineering and a member of American Academy of Arts and Sciences. In 2004 he was knighted by H.M. Queen Elizabeth, and in 2007, he was awarded the Order of Merit.

Collective Knowledge

By Steffen Staab

At the foundation of AI there is the need to have knowledge and to work with this knowledge in order to solve our problems or accomplish something intelligent and exciting. When I started working in AI in the early 1990s, one of the most pressing issues was the bottleneck of procuring this knowledge from experts or from a limited set of documents. Sir Tim Berners-Lee changed this situation completely.

First, he recognized that people are willing to provide knowledge for free when they are given the right incentives and the right tooling. He developed formats (HTML), protocols (HTTP), and services (Web server) that led to the World Wide Web, a base of collected knowledge that is bigger than I could have ever dreamed of. By its very principles, it not only constitutes a completely unstructured aggregation of text and media, but through the participation of and recommendation by all people with access to computing networks, it became a truly collective source of knowledge long before the term Web 2.0 was coined.

Second, he recognized that all of us not only have knowledge implicitly available in text and media sources, but we all have structured data and knowledge en masse, though often we are not even aware of it. By putting data on the Semantic Web, he conceived the drastic next step that is necessary to both find implicit knowledge and gain insight into the knowledge collected by the many.

The World Wide Web and the Semantic Web are two, of his many, technical contributions that Berners-Lee has given to the AI community that have changed its playing field forever. It is important to note that he did not just conceive the ideas, but he was a promoter, developer, and nurse to his brain children. In fact, if it would have been for the technological development alone, we might not see either a Web or Semantic Web today. Berners-Lee was also a creator of the social machinery that backed his technical innovations. He built it up the with a truly democratic momentum that favored the effective small step understood and appreciated by the many over the fortune or wisdom of the few, leading to an unprecedented global uptake of his technological innovations. Thus, the World Wide Web Consortium (W3C), which he founded, has achieved its immense influence not by any regulatory decision power, but simply by the trust that so many people put into it because of the way it was conceived.

The Web is now an unprecedented social machine that is producing both implicit and explicit collective knowledge. Immense progress in AI, like the recent success of IBM's Watson on the game show *Jeopardy* (http://www.ibm.com/innovation/us/watson), has been achieved by building on it.

To the luck of all of us, Berners-Lee is not resting; he is working on Web Science to better understand the socio-technological machine that has built and is building the Web, further changing what it means to produce and harness intelligence.

Steffen Staab is a professor for databases and information systems in the Department of Informatics and director of the Institute for Web Science and Technologies (WeST) at the University Koblenz-Landau.

Lotfi Zadeh

Lotfi Zadeh is a mathematician, electrical engineer, computer scientist, professor of computer science, and the director of the Berkeley Initiative in Soft Computing (BISC) at the University of California, Berkeley. He received an MS in electrical engineering from the Massachusetts Institute of Technology and a PhD in electrical engineering from Columbia University. He published his seminal work on fuzzy sets in 1965, in which he detailed the mathematics of fuzzy set theory. In 1973, he proposed his theory of fuzzy logic. Zadeh is also credited, along with John R. Ragazzini, with having pioneered the development of the z-transform method in discrete time signal processing and analysis in 1952. These methods are now standard in digital signal processing, digital control, and other discrete-time systems used in industry and research. In 1991, Zadeh introduced the concept of soft computing, which highlights the emergence of computing methodologies in which the accent is on exploiting the tolerance for imprecision and uncertainty to achieve tractability, robustness, and low solution cost. He has received the Benjamin Franklin Medal, IEEE Richard W. Hamming Medal, ACM Allen Newell Award, and AIM Information Science Award. He is also a member of the National Academy of Engineering and a foreign member of the Russian Academy of Natural Sciences.

Fuzzy Logic and Computational Intelligence

By Derong Liu

As an indispensable constituent of AI, fuzzy logic is a superset of conventional (Boolean) logic that has been extended to handle the concept of partial truth, where the truth value can range between completely true and completely false. As the creator of a new field of mathematics—fuzzy set theory and fuzzy logic—Lotfi Zadeh's intellectual contributions are myriad. He is also known for his research in system theory, information processing, AI, expert systems, natural language understanding, and the theory of evidence. His current research is focused on fuzzy logic, computing with words, and soft computing, which is a coalition of fuzzy logic, neurocomputing, evolutionary computing, probabilistic computing, and parts of machine learning.

In 1965, Zadeh conceived of the idea that developed into what is now known as fuzzy logic, a model for human reasoning in which everything, including truth, is a matter of degree. The theory challenges classical logic's belief in absolute true or false. Although initially met with disdain, fuzzy logic is widely accepted today, with applications for everything from consumer products, industrial systems, and operations research to medicine, geology, and physics.

Because of the importance of the relaxation of Aristotelian logic, which opens up applicability of rational methods to the majority of practical situations without dichotomous truth values, Zadeh is one of the most referenced authors in the applied mathematics and computer science fields. In the theory of fuzzy sets, he proposed using a membership function (with a range covering the interval [0, 1]) operating on the domain of all possible values. He proposed new operations for the calculus of logic and showed that fuzzy logic was a generalization of classical and Boolean logic. He also proposed fuzzy numbers as a special case of fuzzy sets, as well as the corresponding rules for consistent mathematical operations (fuzzy arithmetic). In addition, Zadeh is credited, along with John R. Ragazzini, in 1952, with having pioneered the development of the z-transform method in discrete time signal processing and analysis. These methods are now standard in digital signal processing, digital control, and other discrete-time systems used in industry and research.

Lotfi Zadeh belongs to a world where there are no boundaries limited to time or place. He really is best characterized as an internationalist. He is a fellow of IEEE, American Academy of Arts and Sciences, ACM, AAAI, and International Fuzzy Systems Association (IFSA). He has published extensively on a wide variety of subjects relating to the conception, design, and analysis of information and intelligent systems and serves on the editorial boards of more than 60 journals.

Derong Liu is a professor in the Institute of Automation, Chinese Academy of Sciences, and a professor in the Department of Electrical and Computer Engineering at the University of Illinois at Chicago. He is also the editor in chief of the *IEEE Transactions on Neural Networks*.

HALL OF FAME

Noam Chomsky

Noam Chomsky is an American linguist, philosopher, cognitive scientist, and activist. After receiving a BA, MA, and PhD in linguistics from the University of Pennsylvania, Chomsky remained on the staff of the Massachusetts Institute of Technology for approximately 19 years. He is now an institute professor and professor emeritus of linguistics at MIT. He is well known in the academic and scientific community as one of the fathers of modern linguistics and a major figure of analytic philosophy. Chomsky developed a theory of transformational (sometimes called generative or transformational-generative) grammar that revolutionized the scientific study of language. He first set out his abstract analysis of language in his doctoral dissertation (1955) and his book *Syntactic Structures* (1957). His work on formal languages (as used in mathematics and logic) and on the acquisition and processing of language has greatly affected AI research as well as computerized language translation, evolutionary psychology, and even immunology. He has received the Kyoto Prize, Helmholtz Medal, Dorothy Eldridge Peacemaker Award, and Ben Franklin Medal in Computer and Cognitive Science. He is also a member of the American Academy of Arts and Sciences, National Academy of Sciences, and American Philosophical Society.

Computational Linguistics and Cognitive Science

By Ee-Peng Lim

Noam Chomsky is a distinguished linguistics and cognitive scientist. He developed grammar rules for natural languages in the 1950s, laying the foundation work for language studies. The work also led to a large body of interesting research on natural language processing and subsequently computational linguistics. Chomsky proposed the well-known transformational grammar and authored the highly influential and cited book *Syntactic Structures*, which brought the computer science and linguistics fields closer to each other. He also created the well-known Chomsky hierarchy of formal languages that classifies formal languages by their different expressive power. The Chomsky hierarchy has been used in linguistics to study languages as well as by computer science researchers to study compiler and automata theory.

Chomsky is the pioneer linking linguistics with cognitive psychology. His review of B.F. Skinner's book *Verbal Behavior* pointed out that language cannot be acquired by a stimulus response process. He instead proposed the view that humans have an innate ability to construct language grammar and generate sentence phrases. His view was considered highly controversial at the time because it changed the way language was studied and shaped the subsequent development of cognitive science.

Chomsky has authored more than 70 books and 1,000 articles. At MIT, institute professorship is awarded to the most distinguished faculty. He became an endowed chair professor at MIT at the age of 37 and an institute professor at 47.

Ee-Peng Lim is a professor in the School of Information Systems and director of the Knowledge Discovery Systems Lab at Singapore Management University.

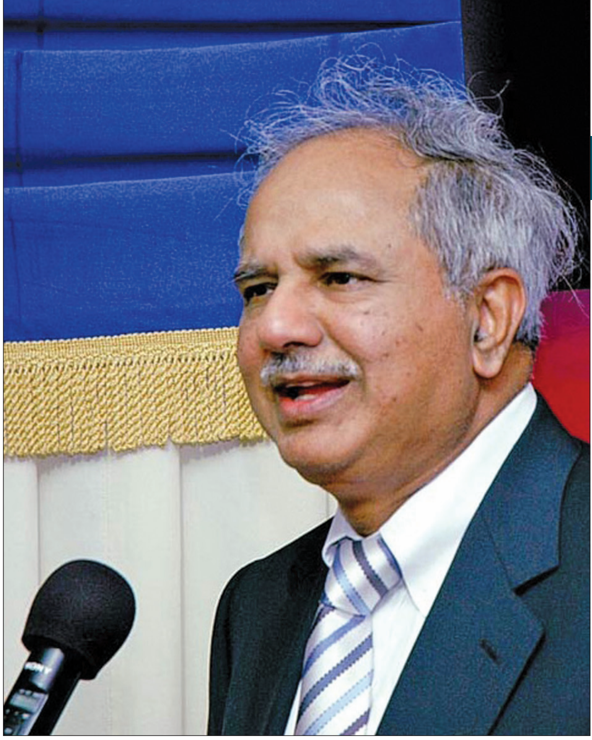

HALL OF FAME

Raj Reddy

Raj Reddy is the Mozah Bint Nasser University Professor of Computer Science and Robotics in the School of Computer Science at Carnegie Mellon University. After receiving a BE from the Guindy Engineering College of the University of Madras, India, an MTech from the University of New South Wales, Australia, and a PhD in computer science from Stanford University, Reddy served on the faculties of Stanford University and Carnegie Mellon University for more than 40 years. He was the founding director of the Robotics Institute and the dean of the School of Computer Science at CMU. His AI research has concentrated on perceptual and motor aspects of intelligence such as speech, language, vision, and robotics. In speech, he and his colleagues developed large vocabulary connected speech recognition systems (including Hearsay and Harpy) as part of the DARPA speech understanding research program. Reddy served as the chief scientist of the Centre Mondial Informatique et Ressource Humaines and was awarded the Legion d'Honneur by French President Mitterrand in 1984. He has received the ACM Turing Award, Okawa Prize, Honda Prize, and Vannevar Bush Award. He is also a member of the National Academy of Engineering and American Academy of Arts and Sciences.

AI and Societal Impact

By Harry Shum

Raj Reddy is one of the early and most influential pioneers in AI and computer science. He has been instrumental in broadening the practical real-world impact of AI technology through a relentless pursuit of designing and constructing various practical, large-scale AI systems. His name is associated with many founding and critical breakthroughs in the perceptual and motor aspects of intelligence such as speech, language, vision, and robotics. In speech, he and his colleagues developed large vocabulary continuous speech recognition systems, such as Hearsay, Dragon, Harpy, and Sphinx. Since then, the CMU team has been responsible for many advances in speech research and the graduates from CMU have gone on to become leaders in the field. He was the founding director of the Robotics Institute, the first such institute or department in the US.

Reddy's contribution in AI, however, is far broader than speech and beyond improving a machine's intelligence. His distinct visionary work of using technologies for society and humanity set him apart from other outstanding technologists. He is a true pioneer in exploring the role of "technology in service of society." He not only created the first universal digital library, a free, online digital library that includes millions of books, monuments, and archives that can be accessed anytime, anywhere, but also continues advocating the need to provide anytime, anywhere access to creative works of the human race. To help bridge the "digital divide," he was instrumental in establishing the Rajiv Gandhi University of Knowledge Technologies in India, to cater to the educational needs of the thousands of low income, gifted rural youth.

Harry Shum is corporate vice president at Microsoft.

Judea Pearl

Judea Pearl is a computer scientist and philosopher, best known for introducing the probabilistic approach to AI and developing Bayesian networks as a tool of inference. He is also credited with developing a computational theory of causal and counterfactual reasoning applicable in several sciences. Pearl has a BS in electrical engineering from the Technion, Israel, an MS in physics from Rutgers University, and a PhD in electrical engineering from the Polytechnic Institute of Brooklyn. After graduation, he worked at RCA Research Laboratories on superconductive phenomena and on advanced memory systems at Electronic Memories. In 1970, he joined the University of California, Los Angeles, where he is currently a professor of computer science and statistics and the director of the Cognitive Systems Laboratory. Pearl was one of the first to mathematize causal modeling in the empirical sciences. His work is also intended as a high-level cognitive model and has contributed to knowledge representation, the philosophy of science, nonstandard logics, and machine learning. He received the ACM Allen Newell Award, IJCAI Award for Research Excellence in Artificial Intelligence, LSE Lakatos Prize in Philosophy of Science, Benjamin Franklin Medal in Computers and Cognitive Science, and David Rumelhart Prize in Cognitive Science. Pearl is a fellow of IEEE and a member of the National Academy of Engineering. He is also president of the Daniel Pearl Foundation (named after his son) and writes frequently on the Middle East peace process.

HALL OF FAME

Probability, Causality, and Intelligence

By Eric Horvitz

Methods for machine learning, reasoning, and decision making under uncertainty with probabilistic graphical models lay at the heart of a 25-year rolling revolution in AI—and Judea Pearl has been a scholar and visionary at the forefront of this revolution. Pearl has pursued principles of intelligent reasoning by elucidating foundational representations and inferential machinery for reasoning under uncertainty. His research played a critical role in catalyzing a paradigm shift in computer science, resulting in an effective and fruitful coupling of Bayesian statistics and computer science.

In the mid-1980s, Pearl and his students were a fount of results on probabilistic graphical models. Pearl and colleagues specified with clarity important proofs and procedures on the representation and manipulation of probabilistic independence within Bayesian networks. Pearl also developed useful algorithms for performing probabilistic inference within these graphical representations. Such inference includes diagnostic reasoning, where certain variables are specified (observations such as patient symptoms) and probabilistic updates are desired for sets of unobserved or hidden variables of interest (hypotheses such as the potential diseases afflicting a patient). Beyond quantitative methods, Pearl's work included efforts to understand qualitative patterns of inference such as intercausal reasoning, capturing how beliefs may change in the face of evidence favoring alternative hypotheses.

Pearl provided an influential synthesis of important results and methods in his 1988 book *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. The book was a revolutionary compendium of ideas and vision when it was published and it remains a classic treatise on representation and reasoning under uncertainty.

Building on his studies of inference and learning from data, in the 1990s, Pearl turned his attention to the foundations of science, exploring standing challenges with understanding causality and causal inference. His work at the intersection of computer science, statistics, and the philosophy of science shed new light on causality and statistical inference. In a series of papers, Pearl introduced clarity on topics that had long been rife with challenge and confusion. His 2000 book *Causality: Models, Reasoning, and Inference* has grown to become yet another classic treatise.

Pearl has been a veritable pioneer in pursuing a scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in computational representations and procedures. He has provided the computer science and statistics communities with a precious cache of results, algorithms, and insights, and has mentored several generations of students along the way.

Eric Horvitz is a distinguished scientist at Microsoft Research, past president of AAAI, and a fellow of the AAAI and of the American Academy of Arts and Sciences.

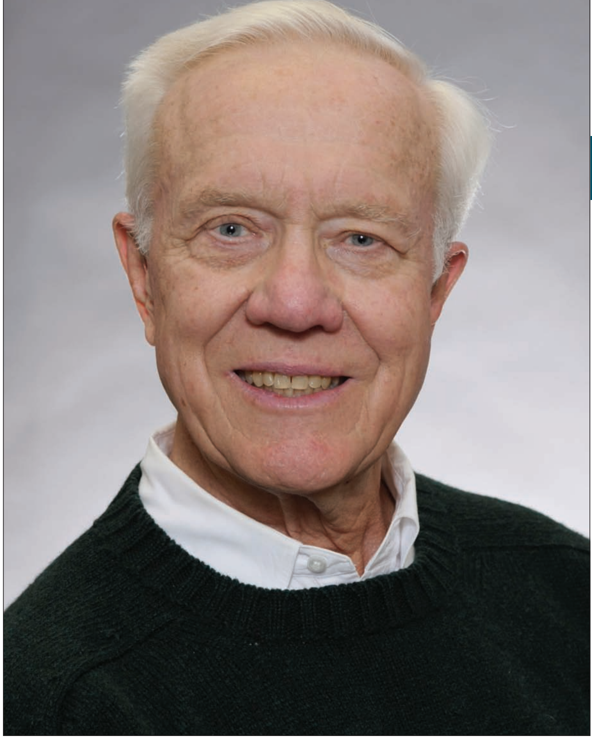

Nils J. Nilsson

Nils J. Nilsson is the Kumagai Professor of Engineering (Emeritus) in the Department of Computer Science at Stanford University. He received a PhD in electrical engineering from Stanford University. Nilsson spent 23 years at the Artificial Intelligence Center of SRI International working on statistical and neural-network approaches to pattern recognition, coinventing the A* heuristic search algorithm and the Strips automatic planning system, directing work on the integrated mobile robot Shakey, and collaborating in the development of the Prospector expert system. He returned to Stanford in 1985 as the chairman of the Department of Computer Science, a position he held until August 1990. He has served on the editorial boards of the *Artificial Intelligence* journal and the *Journal of Artificial Intelligence Research* and was an area editor for the *Journal of the ACM*. He is a past-president of the American Association for Artificial Intelligence (AAAI) and a cofounder of Morgan Kaufmann Publishers. He received the IJCAI Research Excellence Award and IEEE Neural-Network Pioneer Award. He is also a foreign member of the Royal Swedish Academy of Engineering Sciences and a fellow of the AAAI and the American Association for the Advancement of Science.

Problem Solving and Planning

By Daniel Zeng

Much of early AI research focused on heuristic search-based problem solving and AI planning in constrained environments such as microworlds. Nils Nilsson was a major contributor to these studies. He coinvented (with colleagues Peter Hart and Bertram Raphael) one of the best-known AI algorithms, A*, which delivers high performance in graph traversal by using heuristics. A* possesses many desirable properties. It will always find a solution if one exists. When given an admissible heuristic function, one that will never overestimate the cost of reaching from the current state to the goal state, A* is optimal and efficient. A* and its more space-efficient variations have enjoyed success in a range of applications going beyond AI.

Jointly with his colleague at the Stanford Research Institute, Richard Fikes, Nilsson also developed Strips, a formal language to describe AI planning problems, and related automated planning techniques. Strips quickly became the de facto representational formalism for AI planning research and laid the foundation for new ways of applying theorem-proving techniques to problem solving. Strips has been extended in several important dimensions (such as to enable learning, incorporate conditionals, and handle probabilistic state information and actions) to support complex problem solving in a variety of application domains. Furthermore, it makes it possible to formally analyze a range of AI problems from the point of view of computational complexity. One of the early applications of Strips was in Shakey, a mobile robot developed by Nilsson and his colleagues at the Stanford Research Institute. Shakey was capable of interacting with a simple physical environment and reasoning about its own actions.

To attest to the influence of Nilsson' work in AI research and education, most first AI courses include A* and Strips at the beginning of the syllabus. Nilsson is also known for his research in machine learning, pattern recognition, and computational intelligence. His long, distinguished research career has been spent between SRI International's Artificial Intelligence Center and Stanford University.

In addition to conducting cutting-edge AI research, Nilsson has published a number of influential AI textbooks and a definitive history of the field. He has made important contributions to the AI research community as the AAAI's fourth president and as a cofounder of Morgan Kaufmann Publishers, which has published many important AI works.

Daniel Zeng is a professor and Eller Fellow in the Eller College of Management at the University of Arizona and a research professor at the Institute of Automation in the Chinese Academy of Sciences.