AI Magazine Volume 26 Number 4 (2006) (© AAAI)

Artificial Intelligence: The Next Twenty-Five Years

Edited by Matthew Stone and Haym Hirsh

■ Through this collection of programmatic statements from key figures in the field, we chart the progress of AI and survey current and future directions for AI research and the AI community.

We know—with a title like that, you're expecting something awful. You're on the lookout for fanciful prognostications about technology: Someday computers will fit in a suitcase and have a whole megabyte of memory. And you're wary of lurid Hollywood visions of "the day the robots come": A spiderlike machine pins you to the wall and targets a point four inches behind your eyes with long and disturbing spikes. Your last memory before it uploads you is of it asking, with some alien but unmistakable existential agony, what does this all mean?

We are not here to offer a collection of fiction. Artificial intelligence isn't a nebulous goal that we hope—or fear—humanity one day achieves. AI is an established research practice, with demonstrated successes and a clear trajectory for its immediate future. David Leake, the editor-in-chief of *AI Magazine,* has invited us to take you on a tour of AI practice as it has been playing out across a range of subcommunities, around the whole world, among a diverse community of scientists.

For those of us already involved in AI, we hope this survey will open up a conversation that celebrates the exciting new work that AI researchers are currently engaged in. At the same time, we hope to help to articulate, to readers thinking of taking up AI research, the rewards that practitioners can expect to find in our mature discipline. And even if you're a stranger to AI, we think you'll enjoy this inside peek at the ways AI researchers talk about what they really do.

We asked our contributors to comment on goals that drive current research in AI; on ways our current questions grow out of our past results; on interactions that tie us together as a field; and on principles that can help to represent our profession to society at large. We have excerpted here from the responses we received and edited them together to bring out some of the most distinctive themes.

We have grouped the discussion around five broad topics. We begin in the first section ("Progress in AI") by observing how far AI has come. After decades of research, techniques developed within AI don't just inform our understanding of complex real-world behavior. AI has advanced the boundaries of computer science and shaped users' everyday experiences with computer systems. We continue in the second section ("The AI Journey: Getting Here") by surveying contributors' experience doing AI. This section offers a more intimate glimpse inside the progress that AI has made as our contributors explain how their approaches

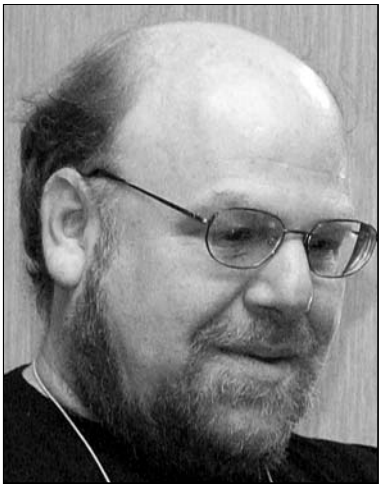

Jim Hendler

Rodney Brooks

to central questions in the field have been formed and transformed over the years by their involvement in research and in the AI community.

The third section ("The AI Journey: Future Challenges") looks forward to characterize some of the new principles, techniques, and opportunities that define the ongoing agenda for AI. By and large, our contributors expect to take up research challenges that integrate and transcend the traditional subfields of AI as part of new, inclusive communities organized around ambitious new projects. They envision big teams solving big problems—designing new fundamental techniques for representation and learning that fit new computational models of perception, human behavior and social interaction, natural language, and intelligent robots' own open-ended activity. The fourth section ("Shaping the Journey") considers the community building that can make such investigations possible. We have a lot to do to facilitate the kind of research we think we need, but a lot of experience to draw on—from creating resources for collaborative research, to establishing meetings and organizations, to training undergraduates and graduate students and administering research institutes. We close, in the fifth section ("Closing Thoughts"), with some reminders of why we can expect takers for this substantial undertaking: the excitement, the satisfaction, and the importance of creating intelligent artifacts.

Across the statements contributors sent us, we see how thoroughly AI has matured into a collaborative process of genuine discovery. AI now seems quite different from what would have been predicted twenty-five years ago: that fact only highlights the progress we have made. AI continues to embrace new perspectives and approaches, and we continue to find new interplay among the complementary principles required for general intelligence.

Our progress itself thus helps tie us together as a community. Whenever new confluences of ideas provide fertile new ground for theory and experiment, they reconnect AI researchers into an evolving network of mutual support, friendship, and fun. Even 50 years after the Dartmouth conference, and 25 years after the founding of AAAI, AI researchers continue to find new points of overlap in one another's interests and new research to explore together. As this research brings results, it cements these intellectual connections in new kinds of computer systems and in a deeper understanding of human experience. We're pleased and honored to present just a slice of a snapshot of this dynamic process. Our only prediction for the next 25 years is that it will continue to bring us unexpected insights and connections to one another.

Every member of AAAI ideally deserves to contribute to this article. But, obviously, space limits do not make that possible, and thus we solicited feedback from far too small a number of the many people who are playing leadership roles in our field. We don't imagine that what we have is definitive; we have set up a website, at www.ai.rutgers.edu/aaai25, to continue the discussion. Your contribution remains welcome.

Progress in AI

AI is thriving. Many decades' efforts of many talented individuals have resulted in the techniques of AI occupying a central place throughout the discipline of computing. The capacity for intelligent behavior is now a central part of people's understanding of and experience with computer technology. AI continues to strengthen and ramify its connections to neighboring disciplines.

The multifaceted nature of AI today is a sign of the range and diversity of its success. We begin with contributions observing the successes we often neglect to mention when we consider AI's progress.

> Last night I saw a prescreening of the movie *Stealth,* the latest science fiction flick in which an artificially intelligent machine (in this case an unmanned combat air vehicle) stars in a major role. The movie has a couple of scenes that pay homage to the 1968 film *2001: A Space Odyssey*, in which HAL, the advanced AI computer, is a key character. What amazed me in the new movie was a realization of just how far AI has come. When I saw *2001,* the idea of the talking computer that understood language was so cool that I decided then and there that I wanted to be an AI scientist someday. In *Stealth,* the AI carries on normal conversation with humans, flies a high-powered airplane, and shows many human-level capabilities without it really raising an eyebrow—the plot revolves around its actions (and emotions), not

around how “cool” it is that a computer can do these things.

From the point of view of the AI vision, we’ve already achieved many of the things the field’s founders used for motivators: for example, a computer beat the world’s chess champ, commercial systems are exploiting continually improving voice and speech capabilities, there are robots running around the surface of Mars, and the word processor I’m using to write this comment helps to correct my grammar mistakes. We’ve grown from a field with one conference to one in which many subareas hold well-attended conferences on a regular basis and in which it is rare to see a university that does not include AI in its undergraduate curriculum. We in the field are sometimes too fast to recognize our own faults and too slow to realize just how amazingly far we’ve come in such a short time.

—Jim Hendler, *University of Maryland*

Artificial intelligence has enjoyed tremendous success over the last twenty five years. Its tools and techniques are now mainstream within computer science and at the core of so many of the systems we use every day. Search algorithms, the backbone of traditional AI, are used throughout operating systems, compilers, and networks. More modern machine-learning techniques are used to adapt these same systems in real-time. Satisfiability of logic formulas has become a central notion in understanding computability questions, and once esoteric notions like semantic ontologies are being used to power the search engines that have become organizers of the world’s knowledge, replacing libraries and encyclopedias and automating business interfaces. And who would have guessed that AI-powered robots in people’s homes would now be counted in the millions? So much accomplishment to bring pride to us all.

—Rodney Brooks, *MIT*

From the engineering perspective, artificial intelligence is a grand success. In education, computer science majors expect to take a subject or two in artificial intelligence, and prospective employers expect it. In practice, big systems all seem to contain elements that have roots in the past half century of research in artificial intelligence.

—Patrick Henry Winston, *MIT*

The AI community has cause for much pride in the progress it has made over the past half century. We have made significant headway in solving fundamental problems in representing knowledge, in reasoning, in machine learning, and more. On the practical side, AI methods now form a key component in a wide variety of real-world applications.

—Daphne Koller, *Stanford University*

Fifty years into AI’s U.S. history and 25 years into AAAI’s history, we’ve come a long way. There are a wide variety of deployed applications based on AI or incorporating AI ideas, especially applications involving machine learning, data mining, vision, natural language processing, planning, and robotics. A large fraction of the most exciting opportunities for research lie on the interdisciplinary boundaries of AI with computer science (systems, graphics, theory, and so on), biology, linguistics, engineering, and science. Vastly increased computing power has made it possible to deal with realistically large though specialized tasks.

—David Waltz, *Columbia University*

The AI success was once composed of a list of offshoot technologies, from time sharing to functional programming. Now it is AI itself that is the contribution.

The AI Journey: Getting Here

The wide-ranging and eclectic sweep of AI as a field is mirrored in the experiences of individual researchers. A career in AI is a license to liberally explore a range of problems, a range of methods, and a range of insights—often across a range of institutions. For many, AI’s successes represent a very personal journey. We are pleased that many of our contributors offered an inside look at their journeys through AI. Thus, Alan Mackworth and Ruzena Bajcsy—in recounting the challenges of bridging discrete, symbolic reasoning with the continuous mathematics of signal processing and control theory—highlight how new concepts of constraint satisfaction and active perception clarified their research programs. Similarly, Bruce Buchanan describes his evolving research into the design of systems that solve complex real-world problems through insights he gained about the interplay of represented knowledge and calculated inference.

In many cases, newly discovered insights lead us to radically change the work we do and the way we talk about it. Consider the range of research domains Aaron Sloman has explored, or watch as

Patrick Winston

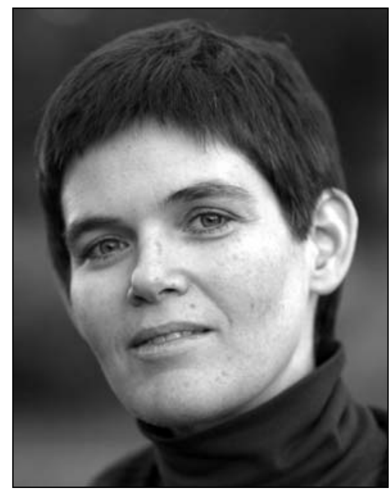

Daphne Koller

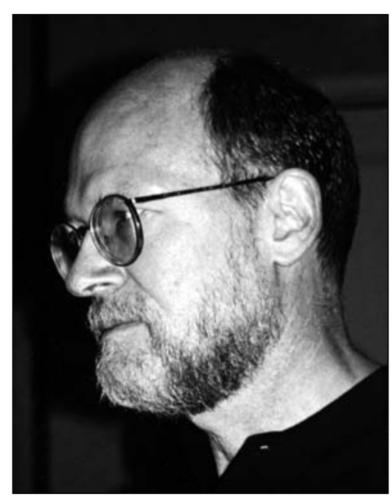

David Waltz

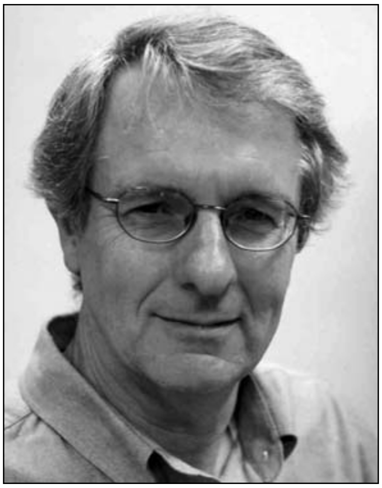

Alan Mackworth

Ruzena Bajcsy

Bruce Buchanan

Michael Kearns and Usama Fayyad trace quite different evolving trajectories through machine learning. As we survey emerging research in AI in the third section ("The AI Journey: Future Challenges"), we'll see that new problems, new discoveries, and new technologies continue to forge new intellectual connections and create new directions for research in AI. This ongoing interplay, in Wolfgang Wahlster's experience, both defines and sustains AI research.

Across the span of a career in AI, the formative mentorship that starts us off has a special place, of course. Ruzena Bajcsy points to the eclectic good taste of John McCarthy, while Aaron Sloman was won over by Max Clowes's passionate advocacy of the fundamental insights in AI and Usama Fayyad was captivated by the romance and energy with which AI was taught at the University of Michigan. Our work is also strongly shaped by our acquaintance with particularly thought-provoking research, as Bruce Buchanan's (and the field's) was by Turing's empiricism of the 1950s, as Alan Mackworth's was by the seminal computer vision research of the 1970s, and as Michael Kearns's was by some of the early mathematics of computational learning theory in the 1980s. We can only hope that our personal and intellectual efforts to sustain AI as a community continue to act so powerfully.

As a young scientist, I found AI's constant ferment exciting, and I still do. I had previously worked in cybernetics, control theory, and pattern recognition, where we modeled intelligence, perception, and action as signal processing. However, that view excluded much of what we know intelligence to require, such as symbolic cognition. Modeling cognition as symbolic computation provided a missing link. But we went too far in modeling intelligence as only symbolic. One of our toughest challenges now is to develop architectures that smoothly combine the symbolic and the subsymbolic. Or, if you like, to synthesize the achievements of logicist AI with those of cybernetics, control theory, neural nets, artificial life, and pattern recognition.

Inspired initially by David Waltz, Ugo Montanari, David Huffman, Max Clowes, and David Marr, I've advocated constraint satisfaction as the unifying model. At both the symbolic and subsymbolic levels we can specify the internal, external, and coupled constraints that agents must satisfy. Those constraints can be static or dynamic. Our development of the robot soccer challenge has forced all of us to develop architectures supporting both proactive and reactive behaviors.

—Alan Mackworth, *University of British Columbia*

I came from Czechoslovakia to the Stanford AI Laboratory in October 1967. This laboratory was one of the three AI labs in the USA and was under the leadership of John McCarthy. The basic philosophy of Professor McCarthy was that AI was about representation of knowledge and that this representation was symbolic. The language we used for the representation was Lisp. To his credit, McCarthy recognized that perception and robotic interaction with the environment was equally important as reasoning strictly on symbolic information. Hence we faced the problem of how to systematically convert the measurements or observations into symbols. What is an edge, straight line, circle, cube, and so on? This is still an open problem.

The tradition that was set at that time (and it has prevailed) is the foundation of a good engineering science: every good theory needs experimental verification. As we go on and understand more complex phenomena, the experiments reflect this complexity.

I implemented this tradition in the GRASP laboratory during my thirty years at the University of Pennsylvania in Philadelphia. Furthermore, coming from a background of control engineering, we recognized the need in building intelligent systems, the importance of controlling the data acquisition, and introduced an new paradigm: *active perception*. We stated that we not just see but we also *look,* and we not only touch but we also *feel*.

—Ruzena Bajcsy, *University of California at Berkeley*

Every empirical science needs both theoreticians and experimenters. Turing saw that operational tests of behavior would be more informative than arguing in the abstract about the nature of intelligence, which established the experimental nature of AI.

The two major research themes for both theoretical and experimental AI have always been knowledge representation (KR) and inference. Clearly an intelligent person or program needs a store of knowledge and needs inferential capabilities to arrive at answers to the problem he/she/it faces in the world. Other big issues, like learning and planning, can be seen as secondary to KR and inference.

Ed Feigenbaum and I were early players in two major controversies: (1) what are the relative contributions of knowledge and inference, and (2) what representation methods are both simple enough to work with and sophisticated enough to capture the kinds of knowledge that experts use? The DENDRAL and MYCIN programs provide experimental evidence on the side of more knowledge, represented simply.

—Bruce Buchanan,
University of Pittsburgh

AI is today routinely employed in so many areas of advanced information technology that it is fair to say that AI stands also for avant-garde informatics, since it is always pushing informatics to its limits. For the steady growth of AI in Germany, it was imperative for AI researchers to stay integrated with mainstream informatics and to collaborate intensively with colleagues from all subareas of computer science. The attempts of AI researchers in some other countries to establish AI as another metascience like cybernetics outside of informatics were unsuccessful.

—Wolfgang Wahlster, *German Research Center for Artificial Intelligence (DFKI)*

I met AI through Max Clowes in 1969 when I was still a lecturer in philosophy at Sussex University[1] and soon became deeply involved, through a paper at IJCAI 1971 criticizing the 1969 logicist manifesto by John McCarthy and Pat Hayes, followed by a fellowship in Edinburgh 1972–1973. Since then I've worked on forms of representation, vision, architectures, emotions, ontology for architectures, tools for AI research, and teaching, links with psychology, biology and philosophy, and most recently robotics, and I have helped to build up two major AI centers for teaching and research (at Sussex and Birmingham).

I believe philosophy needs AI and AI needs philosophy. Much of what philosophers write about consciousness and the mind-body problem shows their ignorance of AI, and many silly debates between factions in AI (for example, about representations, use of symbols, GOFAI) and some fashions (for example, recent enthusiasm for "emotions") result from doing poor philosophical analysis.

I always thought progress in AI would be slow and difficult and that people who predicted rapid results had simply failed to understand the problems, as sketched in my 1978 book.[2]

—Aaron Sloman,
University of Birmingham

Twenty years ago, I arrived at Harvard to work with Les Valiant on the field that would shortly become known as computational learning theory but that at the time consisted exclusively of two algorithmically focused papers by Valiant, and an early draft of the rather mind-bending (to a first-year graduate student, at least) "four Germans" paper on the exotic and powerful Vapnik-Chervonenkis dimension.[3] It was a great time to enter the field, as virtually any reasonable problem or model one might consider was untouched territory.

Now that the field is highly developed (with even many unreasonable problems sporting hefty literatures), I think that the greatest sources of innovation within computational learning theory come from the interaction with the experimental machine learning and AI communities. In a 2003 International Conference on Machine Learning (ICML) talk, I recalled how my first paper was published in ICML 1987, then an invitation-only workshop. To our amusement, the program committee strongly advised us not to use abstract symbols like x_1 for feature names, but warmer and fuzzier terminology like *can_fly* and *has_wings*.

Perhaps we smirked a bit, but we understood the sentiment and complied. Both sides have come a long way since then, to their mutual benefit. The richness of the theory that has been either directly or indirectly driven by the concerns and findings of empirical machine learning and AI work is staggering to me, and it has been a great pleasure to be a theoretician working in a field in such a close dialogue with practitioners. I am hard-pressed to think of other branches of computer science that enjoy comparable marriages. May the next twenty years bring even more of the same; I cannot predict the results but know they will be interesting.

—Michael Kearns,
University of Pennsylvania

I have fond recollections of the early years of my "discovering" the field of AI. Coming across it in a graduate course in Michigan back in 1985, I was fascinated and inspired by a field that had the bold vision of nothing short of modeling human intelligence. The field in its early days consisted of a collection of works that spanned everything from computer science theory to biology, to psychology, to neural sciences, to machine vision, to more classical computer science tricks and techniques. The excitement was very high and the expectations even higher. As I decided to start working in the subarea of machine learning, I started to re-

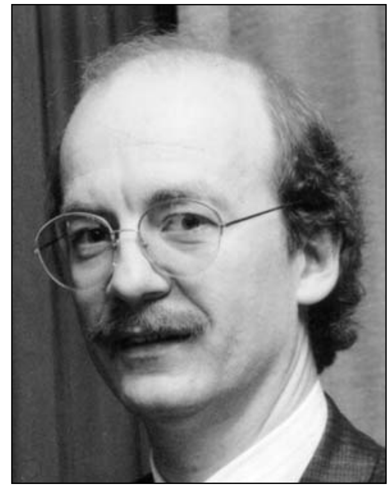

Wolfgang Wahlster

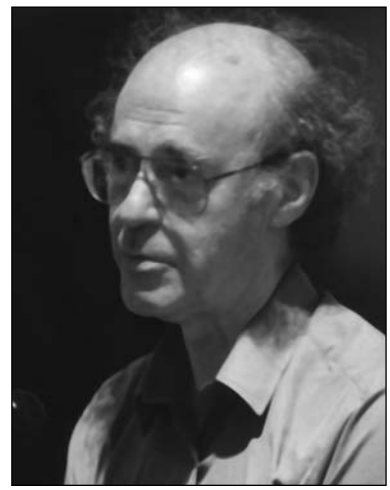

Aaron Sloman

Michael Kearns

Usama Fayyad

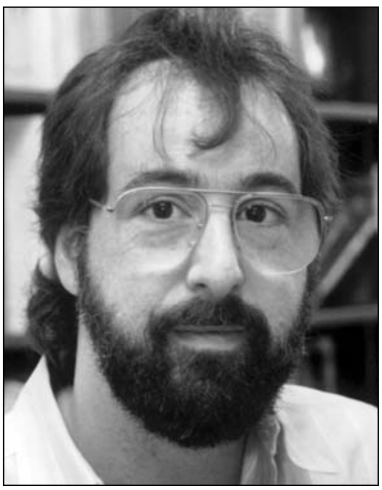

Paul Cohen

alize how difficult the problems are and how far we truly are from realizing the ultimate dream of a thinking machine. What I also realized at the time was that specialization with some deep technical approaches and mathematical rigor were a necessity to make progress.

In reflecting back on those days of romantic excitement, I am very pleased at what they drove in terms of engineering achievements and new fields of study. In my own area of machine learning, while the vision of pursuing general algorithms that "learn from experience" morphed itself into highly specialized algorithms that solve complex problems at a large scale, the result was the birth of several new subfields of specialization. Combining learning algorithms with database techniques and algorithms from computational statistics resulted in data-mining algorithms that work on very large scales. The resulting field of data mining is now a vibrant field with many commercial applications and significant economic value. This journey has also taken me personally from the world of basic scientific research to the business side of realizing economic value from applying these algorithms to commercial problems and finally to working at the "strategy" level on the senior executive team of the largest Internet company in the world, Yahoo!, where data drives many products and strategies.

In looking back at it, I can only say in wonder: what a ride!

—Usama Fayyad, *Yahoo!*

The AI Journey: Future Challenges

As a field, AI researchers have always looked for generality in the intelligent behavior our artifacts exhibit, and generality remains a central challenge. Rod Brooks, Patrick Winston, and Paul Cohen offer us a call to arms.

Artificial intelligence has not yet succeeded in its most fundamental ambitions. Our systems are still fragile when outside their carefully circumscribed domains. The best poker-playing program can't even understand the notion of a chess move, let alone the conceptual idea of animate versus inanimate. A six-year-old child can discuss all three domains but may not be very good at any of them compared to our specialized systems. The challenge for AI, still, is to capture the fundamental nature of generalized perception, intelligence, and action. Worthy challenges for AI that would have tremendous practical impact are, in my opinion: (1) the generic visual object-recognition capabilities of a two-year-old child; (2) the manual dexterity of a six-year-old child; (3) the social interaction and language capabilities of a ten-year-old child. So much work for all of us to be challenged by.

—Rod Brooks, *MIT*

From the scientific perspective, not so much has been accomplished, and the goal of understanding intelligence, from a computational point of view, remains elusive. Reasoning programs still exhibit little common sense. Language programs still have trouble with idioms, metaphors, convoluted syntax, and ungrammatical expressions. Vision programs still stumble when asked to describe an office environment.

—Patrick Henry Winston, *MIT*

Our choice of problems is telling: Small and technical, not large and important. A large, important problem is to work out the semantics of natural language—including all the required commonsense knowledge—so machines can read and understand the web. Another is to develop robots that understand what they see and hear.

Understanding is hard, so AI approximates it with increasingly sophisticated mappings from stimuli to responses: feature vectors to class labels, strings in one language to strings in another, states to states. I once had a robot that learned to map positive and negative translational velocity to the words "forward" and "backward," but never learned that forward and backward are antonyms. It understood the words superficially. It had the kind of understanding we can measure with ROC curves. Every child does better.

—Paul Cohen, *University of Southern California*

Such broad capabilities need not originate in a single fundamental principle or algorithm that applies across the board. Instead, they may be the product of a range of different models, representations, and experience, appropriately combined. Building a general system may hinge on principled, flexible, and extensible ways of putting the pieces together. If that's right, we'll have to start with a grounded understanding of the meanings of representations, as Sebastian Thrun argues, but we'll also need ways of scaffolding sophisticated intelligent behavior over underlying abilities to perceive and act in the world, and we'll need operational ways of weaving together restricted solutions into systems that exhibit more robust behavior. In such architectures, David Waltz, Manuela Veloso, and Patrick Winston find parallels to human intelligence. Despite its flexibility, our own intelligence is an evolved capacity with clear limitations. We manage to act so successfully in part because we bring the right sets of cognitive skills, including our notable strengths in domains of vision, language, and action.

One of the big dreams of AI has been to build an artificially "intelligent" robot—a robot capable of interacting with people and performing many different tasks. We have seen remarkable progress on many of the component technologies necessary to build AI robots. All these tremendous advances beg the obvious question: Why don't we have a single example of a truly multipurpose robot that would, even marginally, deserve to be called artificially intelligent?

I believe the key missing component is representation. While we have succeeded in building special-purpose representations for specialized robot applications, we understand very little about what it takes to build a lifelong learning robot that can accumulate diverse knowledge over long periods of time and that can use such knowledge effectively when deciding what to do. It is time to bring knowledge representation and reasoning back into robotics. But not of the old kind, where our only language to represent knowledge was binary statements of (nearly) universal truth, deprived of any meaningful grounding in the physical world.

We need more powerful means of representing knowledge. Robotics knowledge must be grounded in the physical world, hence knowledge acquisition equals learning. Because data-driven learning is prone to error, reasoning with such knowledge must obey the uncertainties that exist in the learned knowledge bases. Our representation languages must be expressive enough to represent the complex connections between objects in the world, places, actions, people, time, and causation, and the uncertainty among them. In short, we need to reinvent the decades-old field knowledge representation and reasoning if we are to succeed in robotics.

—Sebastian Thrun, *Stanford University*

We are still far short of truly intelligent systems in the sense that people are intelligent—able to display "common sense," deal robustly with surprises, learn from anything that can be expressed in natural language, understand natural scenes and situations, and so on. At the same time AI has tended to splinter into specialized areas that have their own conferences and journals and that no longer have the goal of understanding or building truly intelligent systems.

My own sense is that the AI research program needs to be rethought in order to have a realistic hope of building truly intelligent systems, whether these are autonomously intelligent or "cognitive prostheses" for human-centered systems. Early AI focused on the aspects of human thought that were not shared with other creatures—for example, reasoning, planning, symbolic learning—and minimized aspects of intelligence that are shared with other creatures—such as vision, learning, adaptation, memory, navigation, manipulation of the physical world. A truly intelligent system will require an architecture that layers specifically humanlike abilities on top of abilities shared with other creatures. Some recent programs at the Defense Advanced Research Projects Agency (DARPA) and the National Science Foundation (NSF) are setting ambitious goals that will require integrated generally intelligent systems, a very promising trend. The best news about the neglect of integrated intelligent systems is that researchers going into this area are likely to encounter a good deal of "low hanging fruit."

—David Waltz, *Columbia University*

Creating autonomous intelligent robots with perception, cognition, and action and that are able to coexist with humans can be viewed as the ultimate challenging goal of artificial intelligence. Approaches to achieve such a goal that depend on rigid task and world models that drive precise mathematical algorithms, even if probabilistic, are doomed to be too restrictive, as heuristics are clearly needed to handle the uncertainty that inevitably surrounds autonomy within human environments. Instead we need to investigate rich approaches capable of using heuristics and flexible experience-built webs of knowledge to continuously question and revise models while acting in an environment. Significant progress depends upon a seamless integration of perception, cognition, and action to provide AI creatures with purposeful perception and action, combined with the ability to handle surprise, to recognize and adapt past similar experience, and to learn from observation. Hence and interestingly, I find that the achievement of the ultimate goal of the field requires us, researchers, to accept that AI creatures

Sebastian Thrun

Manuela Veloso

Tuomas Sandholm

Peter Norvig

Daniel Bobrow

are evolving artifacts most probably always with limitations, similarly to humans. Equipped with an initial perceptual, cognitive, and execution architecture, robots will accumulate experience, refine their knowledge, and adapt the parameters of their algorithms as a function of their interactions with humans, other robots, and their environments.

—Manuela Veloso, *Carnegie Mellon University*

Since the field of artificial intelligence was born in the 1960s, most of its practitioners have believed—or at least acted as if they have believed—that language, vision, and motor faculties are the I/O channels of human intelligence. Over the years I have heard distinguished leaders in the field suggest that people interested in language, vision, and motor issues should attend their own conferences, lest the value of artificial intelligence conferences be diminished by irrelevant distractions.

To me, ignoring the I/O is wrongheaded, because I believe that most of our intelligence is in our I/O, not behind it, and if we are to understand intelligence, we must understand the contributions of language, vision, and motor faculties. Further, we must understand how these faculties, which must have evolved to support survival in the physical world, enable abstract thought and the reuse of both concrete and abstract experience. We must also understand how imagination arises from the concert of communication among our putative I/O faculties, and we must learn how language's symbols ground out in visual and other perceptions.

—Patrick Henry Winston, *MIT*

Intelligence is key not only for our physical robots but also for the intermediaries we recruit for competition, communication, and collaboration in virtual worlds. For example, Tuomas Sandholm sees electronic marketplaces as an area where AI can change the world. Peter Norvig, meanwhile, delights us with the potential of AI to forge relationships among the billions of readers and writers on the web. Though it's easy to forget, much of our own intelligence is specifically social. We count on our abilities to explain, coordinate, and adapt our actions to one another, and as Danny Bobrow and Joe Marks, Chuck Rich, and Candy Sidner argue, duplicating these abilities offers an exciting bridge between AI and human–computer collaboration.

A significant portion of trade is already conducted electronically, and markets are increasingly going electronic. This presents a wonderful opportunity for AI: electronic marketplaces provide new exciting research questions, and AI can significantly help generate more efficient market outcomes and processes.

One example is *expressive competition,* a generalization of combinatorial auctions. The idea is to let buyers and sellers (human or software) express demand and supply at a drastically finer granularity than in traditional markets—much like having the expressiveness of human-to-human negotiation, but in a structured electronic setting where demand and supply are algorithmically matched. A combination of AI and operations research techniques have recently made expressive competition possible, and today almost all expressive competition markets are cleared using sophisticated tree search algorithms. Tens of billions of dollars of trade have been cleared with this technology, generating billions of dollars of additional value into the world by better matching of supply and demand.

Less mature, but promising, roles of AI include the following: (1) Automatically designing the market mechanism for the specific setting at hand. This can circumvent seminal economic impossibility results. (2) Designing markets where finding an insincere (socially undesirable) strategy that increases the participant's own utility is provably hard computationally. (3) Supplementing the market with software that elicits the participants' preferences incrementally so that they do not have to determine their preferences completely when that is unnecessary for reaching the right market outcome. (4) Taking into account the issues that arise when the participants incur costs in determining their preferences and can selectively refine them.

There are numerous other roles for AI, undoubtedly including many that have not even been envisioned. With this brief note I would like to encourage bright young (and old) AI researchers to get involved in this intellectually exciting area that is making the world a better place.

—Tuomas Sandholm, *Carnegie Mellon University*

So far we know of exactly one system in which trillions of pieces of information can be intelligently transmitted to billions of learners: the system of publishing the written word. No other system, artificial or otherwise, can come within a factor of a million for successful communication of information. This is despite the fact that the written word is notoriously ambiguous, ill-structured, and prone to logical inconsistency or even fallacy.

In the early days of AI, most work was on creating a new system of transmission—a new representation language, and/or a new axiomatization of a domain. This was necessary because of the difficulty of extracting useful information from available written words. One problem was that people tend not to write down the obvious; as Doug Lenat and Ed Feigenbaum put it, "each of us has a vast storehouse of general knowledge, though we rarely talk about any of it ... Some examples are: 'water flows downhill' ...."[4] This is undeniably true; if you look at a 1,000-page encyclopedia, there is no mention of "water flows downhill." But if you look at an 8 billion page web corpus, you find that about one in a million pages mentions the phrase, including some quite good kindergarten lesson plans.

This suggests that one future for AI is "in the middle" between author and reader. It will remain expensive to create knowledge in any formal language (Project Halo suggests $10,000/page) but AI can leverage the work of millions of authors of the written word by understanding, classifying, prioritizing, translating, summarizing, and presenting the written word in an intelligent just-in-time basis to billions of potential readers.

—Peter Norvig, *Google*

People build their own models of systems they don't understand and may make unwarranted extrapolations of their capabilities—which can lead to disappointment and lack of trust. An effective intelligent system should be transparent, able to explain its own behavior in a way that connects to its users' background and knowledge. An explanation is not a full trace of the process by which the system came to a conclusion. It must highlight important/surprising points of its process and indicate provenance and dependencies of resources used. Systems that evolve through statistical learning must explain (and exemplify) categories it uses and clarify for a user what properties make a difference in a particular case. Such systems must not be single minded, hence should be interruptible and able to explain current goals and status. They should be able to take guidance in terms of the explanation they have given. Artificial intelligence systems must not only understand the world, and the tasks they face, but understand their users and— most important—make themselves understandable, correctable, and responsible.

—Daniel G. Bobrow, *PARC*

In his prescient 1960 article titled "Man-Computer Symbiosis," J. C. R. Licklider wrote: "[Compare] instructions ordinarily addressed to intelligent human beings with instructions ordinarily used with computers. The latter specify precisely the individual steps to take and the sequence in which to take them. The former present or imply something about incentive or motivation, and they supply a criterion by which the human executor of the instructions will know when he has accomplished his task. In short: instructions directed to computers specify courses; instructions directed to human beings specify goals."[5]

Licklider goes on to argue that instructions directed to computers should be more like instructions to human beings. Even today, this is a radical idea outside of AI circles. Most research on human-computer interaction (HCI) has focused on making interaction with computers more efficient by adding new input and output mechanisms. AI researchers, however, are working to fundamentally change the level of HCI from command-oriented interaction to goal-oriented collaboration. Furthermore, since Licklider wrote the words above, researchers in AI and the neighboring fields of linguistics and cognitive science have accumulated a large body of empirical knowledge and computational theory regarding how human beings collaborate with one another, which is helping us to realize Licklider's vision of human-computer symbiosis.

—Joe Marks, Charles Rich, and Candace Sidner, *Mitsubishi Electric Research Labs*

As we aim for broader capabilities, our systems must draw on a more diverse range of ideas. This brings new challenges for research in the field. As Daphne Koller explains, our problems now require us to take ideas from across subfields of AI and make them work together. Several of our contributors suggest that well-defined, long-range projects can inspire researchers with different backgrounds and specializations to bridge their ideas and results. Henry Kautz will be working on understanding human activities in instrumented environments. Tom Mitchell will be working on understanding the language of the web.

A solution to the AI problem—achieving a truly intelligent system—remains elusive.

The capabilities that appear the hardest to achieve are those that require interaction with an unconstrained environment: machine perception, natural language understanding, or commonsense reasoning. To build systems that address these tasks, we need to draw upon

Joe Marks

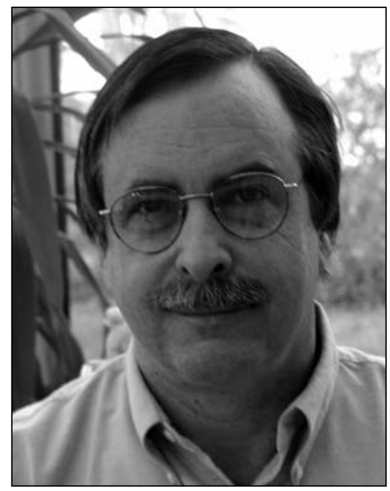

Charles Rich

Candace Sidner

Henry Kautz

Tom Mitchell

the expertise developed in many subfields of AI. Of course, we need expertise in perception and in natural language models. But we also need expressive representations that encode information about different types of objects, their properties, and the relationships between them. We need algorithms that can robustly and effectively answer questions about the world using this representation, given only partial information. Finally, as these systems will need to know an essentially unbounded number of things about the world, our framework must allow new knowledge to be acquired by learning from data. Note that this is not just "machine learning" in its most traditional sense, but a broad spectrum of capabilities that allow the system to learn continuously and adaptively.

Therefore, in addition to making progress in individual subfields of AI, we must also keep in mind the broader goal of building frameworks that integrate representation, reasoning, and learning into a unified whole.

—Daphne Koller, *Stanford University*

One of the earliest goals of research in artificial intelligence was to create systems that can interpret and understand day to day human experience.

Early work in AI, in areas such as story understanding and commonsense reasoning, tried to tackle the problem head on but ultimately failed for three main reasons. First, methods for representing and reasoning with uncertain information were not well understood; second, systems could not be grounded in real experience without first solving AI-complete problems of vision or language understanding; and third, there were no well-defined, meaningful tasks against which to measure progress.

After decades of work on the "bits and pieces" of artificial intelligence, we are now at a time when we are well-poised to make serious progress on the goal of building systems that understand human experience. Each of the previous barriers is weakened:

First, we now have a variety of expressive and scalable methods for dealing with information that is both relational and statistical in nature. Second, the development and rapid deployment of low-cost ubiquitous sensing devices—including RFID tags and readers, global positioning systems, wireless motes, and a wide variety of wearable sensors—make it possible to immediately create AI systems that are robustly grounded in direct experience of the world. Third, there are a growing number of vital practical applications of behavior understanding, including assistive technology for the disabled, aging in place, security and surveillance, and data collection for the social sciences.

—Henry Kautz, *University of Washington*

I believe AI has an opportunity to achieve a true breakthrough over the coming decade by at last solving the problem of reading natural language text to extract its factual content. In fact, I hereby offer to bet anyone a lobster dinner that by 2015 we will have a computer program capable of automatically reading at least 80 percent of the factual content across the entire English-speaking web, and placing those facts in a structured knowledge base.

Why do I believe this breakthrough will occur in the coming decade? Because of the fortunate confluence of three trends. First, there has been substantial progress over the past several years in natural language processing for automatically extracting named entities (such as person names, locations, dates, products, and so on) and facts relating these entities (for example, WorksFor[Bill, Microsoft]). Much of this progress has come from new natural language processing approaches, many based on machine learning algorithms, and progress here shows no sign of slowing. Second, there has been substantial progress in machine learning over the past decade, most significantly on "bootstrap learning" algorithms that learn from a small volume of labeled data, and huge volumes of unlabeled data, so long as there is a certain kind of redundancy in the facts expressed in this data. The third important trend is that the data needed for learning to read factual statements is finally available: for the first time in history every computer has access to a virtually limitless and growing text corpus (such as the web), and this corpus happens to contain just the kind of factual redundancies needed. These three trends, progress in natural language analysis, progress in machine learning, and availability of a sufficiently rich text corpus with tremendous redundancy, together make this the right time for AI researchers to go back to one of the key problems of AI—natural language understanding—and solve it (at least for the factual content of language).

—Tom Mitchell,
Carnegie Mellon University

Shaping the Journey

To pursue the kinds of goals and projects sketched in the previous section, we need the right institutions as well as the right ideas. If we share programs and data in more open collaborations, we can make it easier for individual researchers to make meaningful contributions to big new problems. If we slant major conferences to emphasize integrative research, we can help these contributions find their audiences. These changes are under way: talk to Tom Mitchell or Paul Cohen about shared resources; talk to John Laird or Paul Cohen about integrative research meetings. Still, to tackle the really big problems, our institutions may have to become bigger and broader, too. Perhaps we will see more large research centers, like DFKI, specifically dedicated to integrative research in artificial intelligence.

> I can think of nothing more exciting than working on integrated human-level intelligent systems, and after 25 years it is still captivating and challenging. One challenge is that it requires interdisplinarity in the small (across subfields of AI) and large scale (with other disciplines outside of AI). A second challenge is that teams are needed to attack the large-scale problems that can challenge integrated AI systems (as evident in many of the recent DARPA programs). This type of research isn't for the faint of heart or those who enjoy solitary work. A third challenge is to communicate research results—if the work is truly interdisciplinary, which field or subfield should it be published in? Moreover, how (and where) can I talk about what is learned about integration, which itself is not native to any specific field? AI has done a marvelous job, leveraging specialization with the inexhaustible growth of new subfields, new applications, and new conferences; but we also need to fiercely support integration. What better way to do this than by making integrated cognitive systems a major emphasis of the AAAI conference?
>
> —John Laird, *University of Michigan*

> What to do? Form societies defined by and dedicated to solving large, important problems. Hold conferences where the criteria for publication are theoretical and empirical progress on these problems. Discourage sophistication-on-steroids; encourage integrations of simple (preferably extant) algorithms that solve big chunks of important problems. Work together; share knowledge bases, ontologies, algorithms, hardware, and test suites; and don't fight over standards. Don't wait for government agencies to lead; do it ourselves. If anyone wants to join a society dedicated to the cognitive, perceptual, and social development of robot babies, or to learning language sufficient to understanding children's books at a 5-year-old's level of competence, please drop me a line.
>
> —Paul Cohen, *University of Southern California*

> Some have mentioned to me that this [natural language understanding for the web] is a large goal. I agree and propose we approach it by forming a shared web repository where facts that are extracted from the web by different researchers' efforts are accumulated and made accessible to all. This open-source shared repository should also accumulate and share learned rules that extract content from different linguistic forms. Working as a research community in this fashion seems the best way to achieve this ambitious goal. And I'd hate to have to leave the field to open a lobster fishery.
>
> —Tom Mitchell, *Carnegie Mellon University*

> DFKI, the German Research Center for AI, employs today more than 200 full-time researchers, many of them holding a Ph.D. degree in AI. With yearly revenues of more than US$23 million, it is probably the world's largest contract research center for AI. It has created 39 fast-growing spin-off companies in many fields of AI. DFKI views itself as a software technology innovator for government and commercial clients. DFKI is a joint venture including Bertelsmann, DaimlerChrysler, Deutsche Telekom, Microsoft, SAP, and the German federal government. Its mission is to perform "innovation pure" application-oriented basic AI research. Although we have always tried to contribute to the grand challenges of AI, we have never experienced an AI winter at DFKI, since we have always been quite cautious with promises to our sponsors and clients, trying to deliver down-to-earth, practical AI solutions. Since its foundation in 1988, many maturing AI technologies have left DFKI's labs and become ubiquitous to the point where they are almost invisible in embedded software solutions.
>
> —Wolfgang Wahlster, *DFKI*

Integrative research will be particularly challenging for research students. To do it, they must master a wide range of formal techniques and understand not just the mathematical details but also their place in overall accounts of intelligent behavior. At the same time, to launch productive careers, they must

John Laird

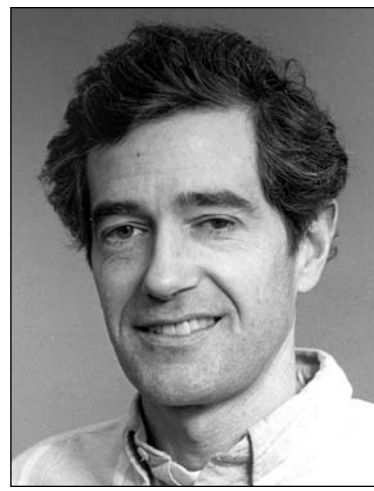

Tom Dean

make a name for themselves with important new ideas of their own. It may take longer than we are used to and require us to think differently about how we nurture new scientists.

So, if AI as a community is to tackle the big problems and continue its progress, we have a lot of work to do besides our own research. But we mustn't think of this work as painful. We'll be doing it with friends and colleagues, as Tom Dean and Bruce Buchanan remind us, and we can expect a unique kind of satisfaction in seeing AI's collaborative process of discovery strengthened and energized by the new communities we foster.

Surely, with such intriguing problems to work on, and with allied fields on the march, this should be a time for universal optimism and expectation; yet many of today's young, emerging practitioners seem to have abandoned the grand original goals of the field, working instead on applied, incremental, and fundamentally boring problems. Too bad. Pessimists will miss the thrill of discovery on the Watson-and-Crick level.

—Patrick Henry Winston, *MIT*

Another reason for slow progress is the fragmentation of AI: people learn about tiny fragments of a whole system and build solutions that could not form part of an integrated humanlike robot. One explanation is that we do not have full-length undergraduate degrees in AI and most researchers have to do a rapid switch from another discipline, so they learn just enough for their Ph.D. topic, and they and their students suffer thereafter from the resulting blinkered vision.

I've proposed some solutions to this problem in an introduction to a multidisciplinary tutorial at IJCAI'05, including use of multiple partially ordered scenarios to drive research.

It requires a lot more people to step back and think about the hard problems of combining diverse AI techniques in fully functional humanlike robots, though some room for specialists remains.

—Aaron Sloman, *University of Birmingham*

The students of AI are sophisticated in both discrete and continuous mathematics, including a recognition of the role of uncertainty. This is necessary because of the increased complexity of problems that we need to attack.

—Ruzena Bajcsy, *University of California at Berkeley*

As teachers, we must challenge students to work on problems in which integration is central and not an afterthought: problems that require large bodies of different types of knowledge, problems that involve interaction with dynamic environments, problems that change over time, and problems in which learning is central (and sometimes problems in which determining the appropriate metrics is part of the research). But given the dynamics in the field of AI, a Ph.D. student must forge an association to an identifiable subfield of AI—some community in which to publish and build a reputation—and as of today that is not "human-level intelligence," "integrated cognitive systems," or even my favorite, "cognitive architecture." So even more than finding a home for publishing, we must grow a community of researchers, teachers, and students in which the integration is the core and not the periphery.

—John Laird, *University of Michigan*

In more than twenty years in this field, the most satisfying moments by far have come from working with people who have set aside their individual interests and biases to inspire students, nurture young scientists, and create community and esprit de corps. And, while I truly enjoyed collaborating with Kathy McKeown on AAAI-91 and Gigina Aiello on IJCAI-99, helping to create the AAAI Robot Competition and Exhibition series with Pete Bonasso, Jim Firby, Dave Kortenkamp, David Miller, Reed Simmons, Holly Yanco, and a host of others was actually a lot of fun. The exercise was certainly not without its aggravations, as getting a sizable group of researchers to agree on any issue is not easy. But most of the effort was spent thinking about how to create useful technology, advance the science and art of robotics, and make the entire experience both educational and inspirational to participants and spectators alike.

It was particularly gratifying to see the buzz of activity around this year's event in Pittsburgh and learn about some of the new ideas involving social robots, assistive technologies, and, of course, cool hardware hacks. I don't know what direction the field should take, and at this particular moment in my career as I return to research after several years in senior administration at Brown University, I'm content to pursue my own personal research interests in the intersection of robotics, machine learning. and computational neuroscience. But I am thinking about how to get students interested in my research area, and in due course, I hope to work with the AI community to run workshops, develop tutorials, sponsor undergraduate research, and pursue all the other avenues open to us to nurture and sustain community, both scientific and social, in our rapidly evolving and increasingly central field.

—Tom Dean, *Brown University*

The sense of collegiality in the AI community has always made AI more fun. Most of the time, the statesmanlike conduct of senior people like Al Newell set an example for debate without rancor. The common goal of understanding the nature of intelligence makes everyone's contribution interesting.

—Bruce Buchanan, *University of Pittsburgh*

Closing Thoughts

As we find new discoveries in AI, as new communities form, as new sets of ideas come together, as new problems emerge and we find new ways of working on them, we will continue to frame new accounts about what our work is. The best stories will resonate not only with what we are doing now but also with how we got here—they will transcend today's fashions and motivate our present activities as a synthesis of our earlier goals and experiences. We close with some brief contributions that take up this discipline of self-examination and, in different ways, distill something important about the landscape of our field. We respond deeply to Bruce Buchanan's characterization of the enduring significance of AI as a goal and the enduring bottom line of AI methodology and to Henry Kautz's acknowledgment of the human meaning of AI results. We endorse Usama Fayyad's prediction that an exciting ride through intellectual space will continue to define the careers of AI researchers. And, with Alan Mackworth, we remind you to keep it real. The discussion will doubtless continue.

> Because there is not enough intelligence in the world and humans often ignore relevant consequences of their decisions, AI can provide the means by which decision makers avoid global catastrophes. I believe we can realize this vision by formulating and testing ideas in the context of writing and experimenting with programs.
>
> —Bruce Buchanan, *University of Pittburgh*

> I believe that understanding human experience will be a driving challenge for work in AI in the years to come and that the work that will result will profoundly impact our knowledge of how we live and interact with the world and with each other.
>
> —Henry Kautz, *University of Washington*

> In looking at the future, we have much to do, and I hope we make some serious progress on that original romantic dream of building machines that truly exhibit learning and thought in general. The dream is still worth pursuing, especially after all we have learned over the past decades. The ride will be a lot more exciting for the new researchers entering the AI field.
>
> —Usama Fayyad, *Yahoo!*

> Think of AI itself as an agent. We need a clear understanding of our own goals, but we must also be willing to seize opportunistically on new developments in technology and related sciences. This anniversary is a lovely opportunity to take stock, to remind ourselves to state our claims realistically, and to consider carefully the consequences of our work. Above all, have fun.
>
> —Alan Mackworth,
> *University of British Columbia*

Notes

1. Sloman, A. 1981. Experiencing Computation: A Tribute to Max Clowes. *Computing in Schools*. (www.cs.bham.ac.uk/research/cogaff/sloman-clowestribute.html.)

2. Sloman, A. 1978. T*he Computer Revolution in Philosophy: Philosophy, Science and Models of Mind*. Brighton, U.K.: Harvester Press. (http://www.cs.bham.ac.uk/research/cogaff/crp.)

3. Valiant, L. G. 1984. A Theory of the Learnable. *Communications of the ACM* 27(11):1134–1142; Valiant, L. G. 1985. Learning Disjunctions of Conjunctions. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, vol. 1, 560–566. Los Altos, CA: William Kaufmann, Inc.; and Blumer, A.; Ehrenfeucht, A.; Haussler, D.; and Warmuth, M. 1985. Classifying Learnable Geometric Concepts with Vapnik-Chervonenkis Dimension. In *Proceedings of the Twenty-Fourth Symposium on Theory of Computation*, 175–85. New York: Association for Computing Machinery.

4. Lenat, D., and Feigenbaum, E. 1991. On the Thresholds of Knowledge. *Artificial Intelligence* 47 (1–3): 185–250.

5. Licklider, J. C. R. 1960. Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics,* Volume HFE-1: 4–11.

Haym Hirsh is a professor and chair of computer science at Rutgers University. His research explores applications of machine learning, data mining, and information retrieval in intelligent information access and human-computer interaction. He received his B.S. in 1983 in mathematics and computer science from UCLA, and his M.S. and Ph.D. in 1985 and 1989, respectively, in computer science from Stanford University. He has been on the faculty at Rutgers University since 1989 and has also held visiting faculty positions at Carnegie Mellon University, MIT, and NYU's Stern School of Business.

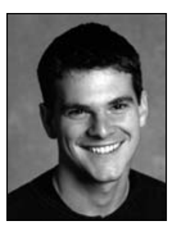

Matthew Stone is associate professor in the department of computer science at Rutgers and has a joint appointment in the Center for Cognitive Science; for the 2005–2006 academic year, he is also Leverhulme Trust Visiting Fellow at the School of Informatics, University of Edinburgh. He received his Ph.D. in computer science from the University of Pennsylvania in 1998. His research interests span computational models of communicative action in conversational agents and include work on generating coordinated verbal and nonverbal behavior (presented most recently at SIGGRAPH 2004 and ESSLLI 2005). He served on the program committee of IJCAI 2003 and as tutorial chair of AAAI 2004. His e-mail address is Matthew.Stone@Rutgers.edu.