

Remembering Marvin Minsky

*Kenneth D. Forbus, Benjamin Kuipers,
Henry Lieberman*

■ Marvin Minsky, one of the pioneers of artificial intelligence and a renowned mathematical and computer scientist, died on Sunday, 24 January 2016, of a cerebral hemorrhage. He was 88. In this article, AI scientists Kenneth D. Forbus (Northwestern University), Benjamin Kuipers (University of Michigan), and Henry Lieberman (Massachusetts Institute of Technology) recall their interactions with Minsky and briefly recount the impact he had on their lives and their research. A remembrance of Marvin Minsky was held at the AAAI Spring Symposium at Stanford University on March 22. Video remembrances of Minsky by Danny Bobrow, Benjamin Kuipers, Ray Kurzweil, Richard Waldinger, and others can be on the sentient webpage¹ or on youtube.com.

The photographs in this article were taken at the AAAI-05 conference in Pittsburgh, Pennsylvania.

Kenneth D. Forbus

In his 1960 essay, *Steps Toward Artificial Intelligence*, Marvin wrote:

A visitor to our planet might be puzzled about the role of computers in our technology. On the one hand, he would read and hear all about wonderful “mechanical brains” baffling their creators with prodigious intellectual performance. And he (or it) would be warned that these machines must be restrained, lest they overwhelm us by might, persuasion, or even by the revelation of truths too terrible to be borne. On the other hand, our visitor would find the machines being denounced on all sides for their slavish obedience, unimaginative literal interpretations, and incapacity for innovation or initiative; in short, for their inhuman dullness.

Similar conversations are being held today. But the balance has shifted. In the 1960s through the 1980s many philosophers and physicists made public comments that AI was impossible, and its seeming achievements illusory. Today, of course, physicists and philosophers are more likely to issue dire warnings about the dangers of AI. That shows just how far this young scientific enterprise has come since its founding, by Marvin Minsky, John McCarthy, Allen Newell, Herb Simon, and others, only 60 years ago.

Marvin was a cognitive scientist before the term was invented. He looked to psychology, neuroscience, and biology for clues to how minds worked, pulling together disparate ideas through the lens of computation. While all of his work has been influential, his final book, *The Emotion Machine*, provides a grand synthesis that is worthwhile reading for anyone interested in AI.

In popular histories, Marvin has often been caricatured as an enemy of neural models. Nothing could be further from



the truth. His interest in modeling brains as well as minds was clear to all who knew him. For example, it was Marvin who invited a freshly minted Ph.D., David Marr, whose thesis was a mathematical model of the cerebellum, to join the AI Lab as a research scientist. What Marvin could not abide were people who did not clearly understand the limitations of their models. Minsky and Papert's analysis of perceptrons provided needed clarity, delineating important theoretical limitations on them. In his preface to the second edition of *Perceptrons*, Marvin pointed out that Rumelhart and McClelland's own data showed that their networks could not learn exclusive-or without exponential amounts of data, as he predicted. Similarly, Marvin was impatient with the primitive neuroscience techniques such as single-cell recordings available for much of his life. The widespread adoption of imaging techniques by neuroscientists suggests that many modern researchers agree as to the power of analyzing larger-scale neural systems.

I was lucky enough to start working at the MIT AI Lab from the second week of my freshman year at MIT in 1973, through the end of my Ph.D. work in 1984. Staying at the same institution was more common in those days. AI labs were very few in number — most universities didn't have them. Such labs had access to resources, like computers and the ARPANET, that most students didn't have access to, even on the

MIT campus. Marvin had created a vibrant, thriving laboratory. He loved to play with ideas, with anyone who was interested. He could often be found in the Playroom, a large open space, riffing about AI or just about any other aspect of science or engineering (or science fiction).

Like many brilliant people, Marvin could be eccentric. For example, his basement was legendary for being a kind of Sargasso Sea of papers and artifacts. In fact, I saw a talk by a historian in the Playroom on the history of Lisp that used it as a resource. The historian dated events in Lisp's creation by treating Marvin's basement as an archeological dig site, dating documents based on their position within layers of other papers. Marvin's desk was similarly cluttered. This was a challenge for students who needed to get a copy of their thesis to him to read, since the document would quickly vanish into the morass. (People didn't read on screen back then, there weren't enough computers around.) Students came up with various ways to ensure that their document attracted Marvin's attention. My technique was to use a telephone cord to hang the thesis from the ceiling so that it dangled over the middle of his desk, at eye level, bobbing gently. But his comments, once you got his attention, were always insightful.

Artificial intelligence has come a long way in just 60 years. The scientific and engineering triumphs



have already changed our lives in many ways, and this process will continue for the foreseeable future. And in doing so, we will continue that voyage that Marvin and others started, to understand minds by building them.

Benjamin Kuipers

Marvin Minsky was my advisor, and he was very important in my life.

My story starts back when I was in high school, and my dad and I would take long walks, talking about math, science, and the nature of the mind. We agreed that this was one of the great scientific problems of all time. I couldn't wait to get to college, so I could take a psychology course, and learn about the science of the mind.

Well, I got to college. And I took a psychology course. And it was a crashing disappointment. The interesting parts weren't rigorous, and the rigorous parts weren't interesting. So I decided I had guessed wrong about psychology, and majored in mathematics. After college, I went to MIT for graduate school in pure math. I was planning to be pure as the driven snow, committed to crystalline mathematical beauty.

But in the spring of 1973, just as a lark, I took the graduate Introduction to Artificial Intelligence course

offered by Marvin Minsky and Seymour Papert. The skies opened! I realized that this was what I had been looking for all along! Their computational methods demonstrated ways to model interesting and important properties of the mind. Like calculus and differential equations had transformed physics more than three centuries before, symbolic knowledge representation and inference methods were beginning a historical transformation of the science of the mind. Like calculus and differential equations had grown and evolved over centuries to become more and more powerful tools for doing physics (and so much else), I knew that these methods would grow and evolve over the decades, and perhaps the centuries, to become powerful enough to describe the mind. I knew that this was what I wanted to spend my life on.

I became a phantom grad student in the Math Department, with a picture on the bulletin board in Building 2, but physically at home in the AI Lab in Tech Square. I was inspired by Marvin's essays, by his students' theses in *Semantic Information Processing*, by class meetings in the iconic Minsky living room, and by discussions in the AI Lab Playroom. Most of all, I was inspired by circulating drafts of the famous "frames paper." There, in a precursor to *The Society of Mind*, Marvin laid out his vision for the organization of knowledge as rich descriptions of complex objects,



© 2005, AAAI

situations, and events, rather than at the finer granularity of logical sentences.

Marvin agreed to be my thesis advisor, as I undertook a project to understand knowledge of space, specifically the large-scale space of the “cognitive map,” describing the structure of buildings and cities at a larger scale than can be observed all at once. He was a wonderful doctoral advisor for me, respecting my ideas, and giving me full freedom to pursue them. My priorities and my direction were deeply inspired by his thinking, but he never pushed me to do things his way. Someone quoted his advice on advising students: “Make sure the students believe that all the good ideas are their own.”

Every month or so, as I worked on my thesis research, I would sit down in his office and tell him, “I think I need some advising.” He would ask what I was up to, and I would explain my progress. Then he would tell me what seemed like a completely random story, like a time that he and Seymour got lost in Buenos Aires, while they were there for a conference. I would leave thinking, “What was that all about?” Then I would think about it, and think about it, and eventually, I would realize that there was a really great idea hidden in there, and it would contribute to my thesis. But I never knew where the idea came from.

At a rough time in my career, he was an enormous

help, and the situation resolved very much for the better. I am very grateful, and I try to pay it forward. He also gave brief and pointed praise, that I still treasure, for coming up with an idea that he really liked.

Marvin had many different accomplishments, but chief among them was being a founder of the field of AI. He and his colleagues started the process of creating the tools and ideas that are revolutionizing the science of the mind. He was disappointed that we haven’t come closer to achieving the goals of AI during his lifetime.

Centuries ago, Newton and Leibniz created tools and ideas that revolutionized the problem of physics, and the work of understanding physics is still not done. Likewise, the problem of the mind is a problem for the centuries, not merely for the decades. Marvin’s contributions have made lasting changes to our understanding of the problem of the mind.

Henry Lieberman

Whenever we encountered each other, Marvin’s eyes would light up, he’d smile, and, instead of “Hello” or “How are you?”, his greeting was always: “What’s new?” I’d tell him what was new with me, or some topic I was thinking about. He’d always react in a way that was surprising, amusing, and profound.

Marvin would plant time bombs in your head.



He'd say something, and you would have a hard time deciding, was that serious? What did that mean? Did he really believe that? What kind of point was he trying to make? Three months later, you'd be walking down the street, and then suddenly it would hit you. "Wait a second! Why *do* we like fun?"

Marvin was a no-nonsense person. He didn't care about money, power, or status. Many times, we would visit some university where he'd be fawned over by the tenured professors. He'd walk right past them, find a grad student in front of a computer and cold coffee and cold pizza. He'd tap the student on the shoulder and say, "What are you thinking about?"

He didn't even really care about computers. He cared about understanding how the mind works. If you wanted to think about it with him, he always had time for you. The human mind is the most complicated thing in existence. What could be more fun than trying to figure out how it works?

But what's the right level to describe it? I think the answer to that question was Marvin's greatest contribution to computer science — and to psychology.

We're just bags of chemicals. Can intelligence be explained at the chemical level? We're just strings of neurons. We can map the 302 neurons of *C. elegans*, but does that tell us how the worm thinks? The mind

must work by electrical impulses flowing around the brain. Will looking at the signals tell us how it works? Different parts of the brain do different things. Does the geography of the brain explain thinking? Other people will tell you the mind is all math, is all biology, is all social interaction. And they'd be right each time.

The brain is hardware. The mind is the software it runs. Marvin explained the workings of the mind as components of software in a high-level programming language. What we think of as concepts can be reified as knowledge representations, as Marvin taught us in the Frame paper. The activity of thinking can be represented as heuristic procedures—trans-frames, micronemes, ways to think, K-lines, A-brains, and B-brains.

Marvin wasn't afraid to push people's hot buttons, especially concerning those aspects of human thought that people cherished, but obstinately refused to try to explain: consciousness, emotion, religion.

Marvin worked in every aspect of AI, from neural nets to cognitive science. The controversies that ravage the field are often just silly arguments between top-down and bottom-up approaches. Marvin advocated coexistence among reactive, reflective, and deliberate layers. What Marvin thought was important was the architecture for managing what he called the Society of Mind, showing how components of the mind both cooperate and compete. His Causal-Diversity Matrix classified the diversity of methods according to what they were good for.

I was once at a graduate seminar with Marvin, and everybody went around the room introducing themselves. "Hi, I'm Robert, I'm a first-year grad student," "Hi, I'm Rebecca, I'm a second year grad student," and so on. Then Marvin's turn: "Hi, I'm Marvin. I'm a 60th year grad student!"

I can only claim to be a 40th year grad student myself. Right now, I'm trying to help launch the Minsky Center for Artificial Intelligence at MIT's CSAIL, the modern incarnation of the Lab he founded. My fondest hope, and the best way to honor his memory, is to try to help create more 60-year graduate students.

Note

1. www.sentient.ai/minsky.

Kenneth D. Forbus is the Walter P. Murphy professor of computer science and a professor of education at Northwestern University.

Benjamin Kuipers is a professor of computer science and engineering at the University of Michigan.

Henry Lieberman is a research scientist at the Massachusetts Institute of Technology's Computer Science and Artificial Intelligence Lab (CSAIL).