AI Agents: Demons or Angels?

Amid the often baffling jargon of artificial intelligence, agents can be the most slithery of concepts to grasp—thanks to conflicting notions of what agents should be. From gifted idiots to "caring" program subroutines to showbiz-inspired smart icons with faces and personalities, agents today fit a wide range of AI precepts. Will one win out? Or, like the high-tech buzzwords *3-D* and *virtual reality*, will they remain dependent on the mind of the beholder?

In his book *Society of Mind*, Marvin Minsky posits agents as the least divisible units of the mind. It's a rather convenient conceit, with Minsky all but admitting as much when defining an agent in the book's glossary as "any part or process of the mind that by itself is simple enough to understand —even though the interactions among groups of such agents may produce phenomena that are much harder to understand." These agents are essentially idiots savants, with incredible capabilities for specialized thinking, yet hardly "smart" enough to be autonomous. Nonetheless, put a bunch of them together in a complex world requiring cognition, learning and memory, and the multiprocessing idiots are king. Or so Minsky's agents would have us believe.

Such agents were hardly the original intent of Oliver Selfridge, whose work is credited with inspiring the first autonomous computer interface agent, the "Oliver." Selfridge says of the agent: "When I used it first, I wanted it to mean something more than a subroutine. I wanted it to take responsibility for handling something, so that I wouldn't have to or a programmer wouldn't have to. In the sense that it cares that what it's responsible for gets done, agents today don't quite meet that definition." Selfridge, who at one time named a whole class of agents "demons" after the squawking devils in Milton's *Pandemonium*, understood the irony of giving beneficent software the same name as one of the banes of show business. "That kind of agent is a leech, living off the good works of others, that's true. But to me an agent is a representative. It essentially means adopting somebody else's purposes."

Agents took a showbiz turn in the 1980s when at its height the Atari division of Warner Communications (now part of Time Warner Inc.) funded research into "intelligent" encyclopedias and video games. A veteran of that effort, now manager of user system ergonomics research at IBM's Almaden (Calif.) Research Center, Ted Selker teaches a course on "Proactive and Reactive Agents in User Interfaces" at Stanford University. "The popular usage in the last year," he says, "has been that an agent is any time users perceive there to be a persona inside the computer doing something for them. One kind of agent uses anthropomorphism, having the computer do some animation that makes you think there's somebody there helping you, and the other is simply macros." Selker distinguishes different functions for today's agents: as assistants that "learn your needs but can become so personalized no one else can benefit from them," or as more generic advisers, essentially smart help facilities. He expects Hollywood stars and cartoon characters to appear soon as on-screen computer agents, "but there's a question of whether this is distracting or helpful to the user."—*S.D.*

one reason we quit was because they stuck to the idea that they were going to build million-dollar machines for big companies to publish newspapers. So we missed out on desktop publishing."

As cofounder of what is now Logo Computer Systems Inc., a Montreal company trying to market a computer language for children, Minsky was in for another disappointment. "Again, we had a very nice thing, the Logo language, done by Seymour Papert. I developed some of the accessories and mechanical things. We had a company [then] called General Turtle and we made 20 of these special Logo computers, which worked wonderfully. But that company failed because the machine was too expensive for schools to buy. A friend of ours, a Canadian, said, 'I can do it more effectively.' So we refounded the company in Montreal and it became Logo Computer Systems, which only makes software and some parts." Minsky has few regrets about this enterprise. "Last I heard, 30 million children have used Logo all over the world. It's very popular."

Minsky's third commercial venture followed a sobering attempt at building hardware in the MIT Artificial Intelligence Lab. As he tells it, "a couple of students invented this kind of processor called the LISP machine. They built one and it worked amazingly well. We built four more. Then everybody wanted one and we built 14 more. The lab turned into a factory. There was demand from the outside, so the students started two companies. They should have started one, but there was a disagreement about leadership. One company was called Symbolics and it sold a great many machines, but they were not mass-produced and cost about $100,000. Then machines like Sun [Microsystems Inc.] workstations started to speed up, and a Sun could do almost as much for a tenth the cost." When Symbolics Inc., Concord, Mass., filed for Chapter 11 protection on January 23, 1993, it marked the biggest bankruptcy in AI history and the end of an era, when the LISP language, developed in part by Minsky and long a lingua franca of AI, had lost out to programs written in C language running on Unix workstations.

The inspiration for his current business involvement came from Danny Hillis, a former graduate student and one of a select cadre of inhabitants of

Minsky's by-now-legendary "attic." Over the years, in addition to distinguished visitors such as science-fiction author Clarke and Jaron Lanier, guru of virtual reality (see UPSIDE, June 1993), the Minsky attic in Brookline, Mass., has housed students such as Michael Hawley and Ken Haase, who went on to become MIT professors, as well as Hillis. Hillis' PhD thesis included the design for a new type of computer using as many as 65,536 processors in parallel to solve AI pattern-recognition problems that can't be solved with traditional single-processor computer systems.

Minsky wanted a massively parallel machine and thought about building one, "but we didn't want the lab to turn into a factory again," he remembers. "There were negotiations with IBM [Corp.] that maybe they would build the machine for us. We talked to Digital [Equipment Corp.] and others, but we couldn't get any existing computer company to try it." Hillis and he "decided that if we could form a company, it would get done. So forming a company is in fact an easy way out, or at least it was for me." Minsky's "easy way out" in 1983 turned into Thinking Machines Corp., Cambridge, Mass., a leader in the $270 million massively parallel-processing supercomputer market.

"That idea of getting a machine to learn by bringing it up as a child is a nice idea, instead of all this boring programming. We once had a young cat that I decided to bring up like a child, and we started with chess. I hung a king . . . to one of those window shades, and I taught the cat to attack the king. He would do this from time to time, and he looked as if he was getting better at it. Of course, as soon as he had mastered this goal, I was going to teach him to defend the king with a bishop."

— MARVIN MINSKY

As the company's chief scientist, Danny Hillis takes pride in its current installed base of 100 Connection Machines. He admits to not having anticipated the impact of his computer design on such areas as statistical analysis and graphics for scientific visualizations. "I didn't realize how general-purpose it was," he says. In fact, it's such non-AI applications as these that have kept Thinking Machines from suffering the fate of other AI hardware firms. Among the company's happy customers: a wholly owned subsidiary of American Express called Epsilon, which uses a Connection Machine to mine databases and improve the yield of direct-mail solicitations by both American Express and outside clients. Hillis spends much of his time with his latest Connection Machine model, the CM-5, researching "genetic algorithms," software that improves itself through evolutionary

processes. Actively involved in designing the next generation of Connection Machines with other senior architects, he can't wait to get beyond the CM-5's current limit of 15,384 processors to the full potential of 65,536. "As soon as you get the latest box, you're always wishing for the next box," Hillis remarks.

Thinking Machines' technical growth has been surprisingly smooth, but the business side has been bumpy in recent months. In October 1993, Sheryl Handler, a cofounder and president for nine years, was eased out of day-to-day control of Thinking Machines and replaced by Richard Fishman, a former Millbank, Tweed legal partner specializing in taking companies public. Handler remains chairman. With Fishman apparently brought on board to preside over a much-anticipated IPO, financial observers were surprised by a subsequent wait-and-see policy, with another round of private financing in November postponing any immediate need to go public.

Minsky professes to be unperturbed by the changes at Thinking Machines. All he says is that "the company is reorganizing. I don't have anything to do with the management of the company at all." He does express his irritation at IBM and others who wouldn't build parallel computers 10 years ago but are now targeting Thinking Machines' market. Still, Minsky seems unconcerned about IPOs and stock prices. He confesses that, despite all his business connections over the years, he's only managed to survive financially thanks to his professor's salary at MIT and the Media Lab—it took a 1990 Japan Prize check, awarded by the Emperor of Japan, to clear all his debts.

FIGHTING THE NEURAL NET

Obviously no stranger to controversy, Minsky can't shake one long-standing fracas: he first pioneered, then abandoned technology that has recently reemerged as one of the most promising areas for AI applications—neural networks. His PhD thesis in 1951 included building the first randomly wired neural-network learning machine. Based on the structure of neurons in the brain, this promising "bottom-up" approach acquired knowledge by discerning patterns in input. It was all but stifled by *Perceptrons*, a 1969 book written by Minsky and AI Lab codirector Seymour Papert. The result was the dominance for a generation by "top-down" AI, based on symbolic programming by humans. The "connectionists," as today's neural-net enthusiasts are known, are still taking Minsky to task for *Perceptrons*.

The neural-net controversy comes up during the visit to Minsky's home. Within this Georgian-style house, virtually every surface of the techno-Victorian living room is cluttered with gadgets. On the mantel is the last surviving remnant of the electromechanical SNARC, Minsky's pioneering neural net. "I realized that a single neural net could only do certain things," he recalls. "If you wanted to do the kinds of things people do, you might have to start top-down to understand the structure of those things and then figure out what kind of brain machinery would do that."

Again bringing up *Perceptrons*, Minsky says, "When Papert and I started the book—it took four years to write—we just tried to explain why these machines weren't doing well." But that's not how a 1990 book, *Neurocomputing*, by Robert Hecht-Nielsen, recalls it. That book describes "a campaign led by Minsky and Seymour Papert to discredit neural-network research and divert neural-network research funding to the field of 'artificial intelligence.' . . . The campaign was waged by means of personal persuasion by Minsky and Papert and their allies, as well as by limited circulation of an unpublished technical manuscript (which was later de-venomized and, after further refinement and expansion, published in 1969 by Minsky and Papert as the book *Perceptrons*)."

In his living room, Minsky bristles at the mention of HNC Inc., San Diego, the neural-net company started by Hecht-Nielsen. "He wrote a book in which he repeats these inaccurate, almost libelous statements about Papert and I proving these theorems to take money away from people. Until he apologizes publicly, I'm certainly not going to say anything about him or his company." Reached at his headquarters, HNC chairman Hecht-Nielsen is surprised but hardly concerned by Minsky's reaction. After applying neural nets to military problems at TRW Inc., Hecht-Nielsen cofounded HNC in 1986 and was buoyed by two rounds of venture capital. According to him, the privately held firm is "a profitable, ongoing business." With 80 employees on the payroll, HNC is finding a niche for neural nets in vertical markets—a strategy shared by a number of successful AI firms.

HNC's hot application is the detection of credit-card fraud, a nearly $3 billion-a-year problem for U.S. companies. The HNC Falcon system, software that runs on mainframes and workstations, examines close to 25 million credit-card accounts for nine major customers. In 1994 HNC will roll out a merchant credit fraud detection system developed with Visa USA Inc., San Mateo, Calif. At the heart of Falcon are two neural-net models that compare each account transaction against credit-card fraud data that's constantly updated, as well as against a cardholder's past account data, producing a numerical score that rates the possibility of fraud. The result is the ability to detect 60 percent of credit-card fraud due to cards lost or stolen before delivery, cards stolen from or lost by cardholders and counterfeit cards, all within minutes of a fraudulent transaction.

HNC also has a trademark on one of the latest buzzwords in AI applications: the DataBase Mining Workstation.

"What's consciousness? People say, 'Could a machine be conscious? Could a machine be aware of itself?' Well, first of all, you're not aware of yourself. You're only sort of baby-aware of yourself. Say a word. . . . What made you think of that word? You haven't the foggiest idea. You're not self-aware. . . . You don't know how you recognize a word when you hear it."

— MARVIN MINSKY

Database mining—the uncovering of statistical relationships within large masses of data using neural nets—can spot likely prospects for direct-mail solicitations, relate individual consumer preferences to demographics and weigh seeming intangibles in assessing real estate values. HNC's DataBase Mining Workstation consists of a Sun workstation or a fast PC with an HNC proprietary coprocessor. After HNC starts receiving its newest chip, called SNAP, in quantity, neural net applications are expected to be speeded up. "We consider ourselves a software firm, though hardware is an important adjunct to our work," says Hecht-Nielsen, whose firm is one of several in position to parlay software expertise into neural-net hardware deals.

"It's a good time to be in neural nets," observes Ed Rosenfeld, editor and publisher of the New York-based newsletter *Intelligence*. After starting his newsletter in 1984, he became convinced that neural nets were "the future of computing" (his newsletter's subtitle), even though the first neural systems weren't offered for sale until 1986. By 1988, he counted some 200 different neural-net startup firms. Because of their late arrival on the scene, many of these companies have been spared the bloodletting that occurred among highfliers during the bleak AI winter. Lean and mean, the neural technology companies are opening new markets for machine intelligence, especially in cooperation with advanced-technology programs at financial firms. "Financial-services companies are trying to squeeze greater yields out of their products," Rosenfeld explains. "They're open to statistical solutions to their problems, and that's where neural nets excel."

As an example, he points out that more than $3 billion in assets in five of Boston-based Fidelity Investments' mutual funds are controlled by neural networks. After Fidelity fund manager Bradford Lewis attended a seminar on neural-net applications by NeuralWare Inc., Pittsburgh, he mastered NeuralWare's NeuralWorks Professional and wrote his own code for weighing market indicators. Lewis' picks have been beating market indicators by 2.3 to 5.6 percentage points in the last two years. Neural-net evangelist Casey Klimasauskas, president and cofounder of NeuralWare with his wife, Jane Klimasauskas, who is now vice president of sales and marketing, started the company in 1986 to bring neural technology to potential users like Lewis. The closely held NeuralWare, which is "profitable more years than not," according to Casey, employs 40 people, including two who teach several seminars a month on applying neural computing to everything from government to process control.

Casey Klimasauskas is eager to spread neural-nets wider with NeuralWare's development tools, which range in price from the turnkey NeuralWorks systems running on PCs, Macintoshes or workstations ($3,995 to $7,995) to modestly priced neural-net learning tools, the NeuralWorks Explorer ($149 to $179). For would-be stock-market analysts, he notes encouragingly that "early on, users would try to relate 200 different indicators, but later we found out that just five give the same accuracy."

Among those close to commercialization of neural nets is Nobel laureate Leon Cooper, who won the 1972 physics prize for his work on superconductivity. In 1975 he cofounded pattern-recognition specialist Nestor Inc. in Providence, R.I., with former Brown University physics department chairman Charles Elbaum; they took the company public in 1983 in a $1.5 million IPO. Later, a private placement of $6 million was necessary to get the company through lean times, when development of new neural-net products took longer than expected. Today, Nestor has yet to return to profitability, though its NestorWriter software for character recognition is highly rated.

Also about to boost the bottom line for Nestor is a 15-year deal with Intel Corp., which has licensed Nestor algorithims for use in Intel neural-net chips. The Ni1000, Intel's second-generation neural-network chip, which includes the equivalent of 1,024 artificial neurons, has already been delivered in beta versions to the federal government's Advanced Research Projects Agency (ARPA), the source of the chip project's original funding. Considered a must in defense applications such as target recognition and unmanned vehicle control, the chips should also become available in commercial quantities this year for use in character recognition, voice recognition and data entry, thus heralding a major revenue stream for Nestor by 1995. The Ni1000 has 3.7 million transistors—more than a Pentium microprocessor—making it one of the most complicated chips ever manufactured. And just in case that project takes longer than expected, Nestor has also joined the lucrative ranks of credit-card fraud detection with its FDS neural-net software running on IBM mainframe computers. The software was initially developed for Pittsburgh-based Mellon Bank, which has used it to reduce the number of fraudulent transactions per card before detection by over 20 percent.

"Something's gone wrong both in AI, I think, and in all of engineering: we're getting worse and worse as we try to make more complicated things. We don't have the slightest idea why the *Mars Explorer* disappeared. It may have blown up; it may have been eaten. Heaven only knows. But in the case of *Galileo* [telescope], isn't it peculiar that for all the talk about reliability and zero defects, we get things that need only one defect [to go wrong]."

—MARVIN MINSKY

THE AI COMEBACK

Other companies are also finding fertile markets for machine intelligence. Waltham, Mass.-based Kurzweil Applied Intelligence Inc., also known as Kurzweil AI, specializes in clinical reporting systems, including speech and handwriting

recognition, targeted at the health care market. After starting 11 years earlier as a general speech-recognition firm, Kurzweil went public in August 1993. It quickly attracted buyers for its 2.1 million shares of stock at an asking price of $10 each. Four months later, the stock is selling at $15, thanks to the company's latest profitable quarterly report and news of an agreement with IBM to form a long-term strategic alliance.

"Why did psychology start in 1895? It's very puzzling, and I'm very angry, because if it had started in minus 1895 [1895 B.C.], we would all be immortal. We'd be made of little chips with quintillions of synapses, and whenever one broke, it would be automatically replaced, and we wouldn't have death or taxes or any of those things."

—MARVIN MINSKY

This is chairman and co-CEO Raymond Kurzweil's third go-round as an entrepreneur. The MIT grad and friend of Minsky started his first company, Kurzweil Computer Products Inc., to develop the original print-to-speech reading machine for the blind (the Kurzweil Reading Machine) in 1975. He sold the company in 1980 to Xerox Corp., which now operates it as Xerox Imaging Systems Inc. In 1984 his second firm, Kurzweil Music Systems Inc., produced the first musical synthesizer (the Kurzweil 250) capable of accurately reproducing the sounds of the grand piano and other orchestral instruments. It was sold to the Japanese firm Young Chang Akki Ltd.

At Kurzweil AI, the two major product lines are KurzweilVoice, a PC-based voice-recognition word-processing system with a 40,000-word vocabulary, and VoiceMed, a turnkey voice-enabled medical reporting system for automated transcription of doctors' reports. With such promising technology, why would a company's publicists shun the words *artificial intelligence* in favor of *applied intelligence*?

Over the years, AI researchers have tended to be overly optimistic in their predictions, resulting in a healthy skepticism among the press and analysts. In addition, it's difficult to get accurate business totals for the artificial intelligence industry as a whole because its boundaries are in such flux. According to Minsky, "artificial intelligence is just advanced computer programming. As soon as something works, it's not called AI anymore. Optical character recognition used to be AI. Now I just slide a document in my scanner and something called Accutext or Omnipage reads it, and I can't find out how it works because it's proprietary. Speech recognition used to be AI. Whenever something gets an application that stands on its own, they call it something else."

Nonetheless, it can seem at times that artificial intelligence is the computer science that dares not speak its name. For instance, executives of privately held Inference Corp., El Segundo, Calif., a leader in knowledge-based systems since its founding in 1983, refused to cooperate with this story so long as it had the words *artificial intelligence* in the title—this despite the success of Inference's trendsetting help-desk technology, which has drawn favorable attention by automating help queries for Compaq Computer Corp.'s customer support phone line and has recently been licensed by Microsoft Corp. to further automate user assistance in the next version of Windows.

Any aversion to AI is understandable to Harvey Newquist, publisher and editor of the newsletter *Critical Technology Trends*, Scottsdale, Ariz. Until the end of 1993 he called it *Artificial Intelligence Trends*. Speaking of the recent hard times for mainstream AI firms, Newquist notes "it wasn't just an 'AI winter,' as some people call it. It was an 'AI Donner Pass,' where companies were cannibalized to survive." In his own case, he changed from the artificial intelligence rubric to include a broader range of technologies, an eclecticism reflected in the breaking down of once-rigid categories within AI.

A number of hard-hit AI pioneers have repositioned themselves for a broader market, as systems software suppliers catering to client-server computing with object-oriented programming (OOP) environments. OOP, originally developed for AI programming, is now viewed as the must-have program-development tool of the '90s. This is Inference's strategy, as well as that of rival Trinzic Corp., Palo Alto, Calif. According to Larry Harris, chief technology officer of the publicly traded Trinzic, "AI programming languages have always been among the most advanced computer tools. Our strength is in working with large amounts of data across platforms."

More competitors are targeting OOP, like IntelliCorp, Mountain View, Calif.; Neuron Data Inc., Palo Alto, Calif.; Gensym Corp., Cambridge, Mass.; and Sapiens USA Inc., Cary, N.C. "We feel there's room for two or three object-oriented development firms and we feel we have the best tools," says IntelliCorp President Ken Haas. "Our object and rules technologies give us a distinct advantage in client-server environments," claims Alain Rappaport, president and chief scientist of the privately owned Neuron Data. Despite their confidence, the OOP strategy could prove disastrous for many competitors. And that's without the probability that Microsoft

"What is pleasure? I think pleasure is actually an extremely negative phenomenon. . . . Pleasure is not an enlarging and ennobling thing; it's a narrowing thing. My motto is you don't have pleasure, pleasure has you. Namely, it's some neurological mechanism that has cut out all the other things that you might possibly do. . . . I suspect it's part of a learning mechanism."

—MARVIN MINSKY

think there's a taboo against thinking about how the mind works. We all pretend to be working very hard [at it], but we're all doing something wrong."

—MARVIN MINSKY

could pull the market away from everyone else by establishing its own standard in object-oriented development environments—it may have already started doing so with its recently announced Common Object Model (COM).

With pure AI disappearing among specialized firms, artificial intelligence is becoming more of a factor among computer industry giants. In addition to its alliance with Inference, since 1987 Microsoft has invested $2 million in Natural Language Inc., a Berkeley, Calif.-based supplier of English-language interfaces to relational databases, and supports its own natural-language laboratory in Redmond, Wash. IBM, which has blown hot and cold over AI, is now selling a breakthrough, under-$1,000 PC-based Personal Dictation System that's the culmination of 21 years of vocal-recognition work at its Thomas J. Watson Research Center. IBM's recently announced "human-centered technology" strategy for upcoming PCs using the PowerPC chip includes speech and handwriting recognition as well as agents with artificial intelligence to make computing more accessible.

According to Minsky, truly accurate speech recognition will have to wait until a computer has a knowledge base of common sense, which is now imminent with the near completion of Cyc, an 11-year project by Doug Lenot, principal scientist at the Microelectronics and Computer Technology Corp., Austin, Tex. The result will be a 200-Mbyte database, well within the storage capacity of most office PCs, with all the human assumptions required to read a dictionary or screen data entries for obvious mistakes. Lenot recalls that Minsky was present at the birth of the Cyc project in a lounge at Stanford University in 1984: "We were doing calculations of what order of magnitude would be required for a commonsense database. Marvin sent somebody out to get some envelopes so he could do some actual back-of-the-envelope calculations for the occasion." Sponsored over the years by Bellcore, Digital, Eastman Kodak Co., NCR and Apple Computer Inc., with first rights to commercializing this technology, Cyc contradicts the stereotype that American firms can't think long-range. However, adds Minsky, "my concern is that there ought to be a dozen such projects. That's a tragedy. People say they're very expensive. That's a big project, but it's not any bigger than what a toy company puts into a new toy or what a cosmetic company puts into a new lipstick."

As for his own identification with Toshiba Corp. as the holder of the endowed chair at the Media Lab, Minsky says, almost wistfully: "Toshiba may be the only place in the world seriously trying to make systems using some of the principles in my book, *Society of Mind*. There's one group doing pattern recognition for speech and for drawings. They liked the idea that the best way to build something and debug it would be to use several different ways and combine them." And yet it would seem that the nonconfrontational Japanese may never understand the contentiousness inherent in Minsky's community of mind. Therein may lie the answer to why the Japanese have never been able to beat American researchers in developing artificial intelligence, an avowed national priority from the much-vaunted Fifth Generation computing initiative to the present.

Still, why confrontation? Minsky explains, as he walks a visitor to the door: "I get my best ideas when cornered. If I can get people mad at me, then maybe I'll get a new idea. There are a lot of people who, in their entire adult lives, have never argued with anyone. I like to break up a party and get people fighting over something important." ■

Steve Ditlea, a freelance writer living in New York, has written several articles for UPSIDE. *The most recent was "Notes Has Lotus Humming" in September 1993.*
