



Artificial intelligence gets real

On a recent visit to the doctor, Edward Feigenbaum had the eerie experience of seeing one of his inventions used in a way he never expected: His 25-year-old concept was being used to diagnose a problem with his own breathing.

"It's using artificial intelligence," the doctor patiently explained about the spirometer, which measures airflow.

"Oh, I see," said Feigenbaum.

A professor of computer science and co-scientific director of the Knowledge Systems Laboratory at Stanford University, Feigenbaum is a pioneer of artificial intelligence (AI) -- the science of making machines think like humans. Dozens of applications have their roots in the Stanford lab he started in 1965 and in related software programs that solve complex problems the same way human experts do.

Feigenbaum was the first person to realize that human intelligence springs not from rules of logic but from knowledge about particular problems (whether it's chemistry or auto mechanics) and about the world in general. For a computer to think the way a human does, Feigenbaum theorized, it would have to know all of the commonsense things that people take for granted -- that a ball thrown in the air will return to earth, for example; or that a cup holding liquid should be carried right-side up.

There are millions of such "rules." Douglas Lenat, a Feigenbaum protégé at Stanford, has spent 15 years coding commonsense rules into software. He figures it will take another 25 years for his company, Austin, Texas-based Cycorp, to finish the job.

Lenat is constructing a massive "expert system." If he succeeds, sometime around 2025 he will produce what he calls "a generally intelligent artifact" that can think like a human. "The big enchilada," is what Feigenbaum calls it.

That's for the future. For now the practical applications of AI are concerned not with replicating human intelligence but with recreating small pieces of it. Feigenbaum-inspired expert systems are used to design buildings and sell products over the Web. They configure airplanes and personal computers and scan the ocean for enemy submarines. They troubleshoot problems in copy machines and make Web browsers easier to use. And, of course, diagnose Prof. Feigenbaum's breathing problems.

Says Feigenbaum, 62: "I'm not interested in theoretical concepts. I like to see my work used in the real world." So much so that he often invests in companies started by his protégés and has done so well at it that he lives in San Francisco on a scale more fitting to a business magnate than a university professor.

That practical bent has sometimes caused colleagues to underestimate his work. He invented the first expert system in 1967, an AI program that could determine the molecular structure of chemical compounds.

Feigenbaum practically got hooted at when he gave a demonstration at a national conference. "Ed, that sounds like great chemistry," a scientist in the audience scoffed, "but what does it have to do with AI?"

After all, AI wasn't about analyzing chemical compounds. It was about creating a working model of the human mind. AI was conceived in 1950 by the brilliant British mathematician Alan Turing, who, before his tragically early death at 41, published a paper best known as "Can a Machine Think?" If computers could be programmed to reason like people, imagine what we could learn about human intelligence.

But early groundbreaking work didn't advance the ball very far. AI researchers were teaching computers to solve logic puzzles and play chess. "Get out into the real world and solve real-world problems,"

Feigenbaum told an audience at Carnegie Mellon University.

Thirty years later researchers are nowhere close to creating a truly intelligent machine. By contrast, Feigenbaum succeeded by thinking small. Unlike his rivals, he didn't set out to recreate all of human intelligence in a computer. His idea was to take a particular expert -- a chemist, an engineer, a pulmonary specialist -- and figure out how that person solved a single narrow problem. Then he encoded that person's problem-solving method into a set of rules that could be stored in a computer.

In case after case it turned out that computers could do as good a job at these highly specialized tasks as humans could, if not better.

Furthermore, expert systems didn't retire, quit or go stale. Better yet, expert systems could combine the knowledge of several specialists and take on jobs that no single human could tackle alone.

As a boy in North Bergen, New Jersey, Feigenbaum was fascinated by a hand-cranked mechanical calculator his father used in his accounting practice. In the mid-1950s Feigenbaum studied at Carnegie Mellon with Herbert Simon (FORBES GLOBAL, Nov. 16), who was conducting the first research into artificial intelligence. Feigenbaum recalls Simon announcing to the class one day that he and his colleague Allen Newell had invented a thinking machine.

To help explain, Simon handed out a manual for an early IBM computer.

Feigenbaum, then 19, stayed up all night reading the manual. "By the next morning, when the sun came up, I was born again," he says. "Here was this beautiful modern electronic version of what I was so intrigued with as a kid -- that little mechanical calculator of my dad's. And furthermore, Herb Simon had said, 'I'm going to show you how to make them think.'"

With Simon as his adviser, Feigenbaum completed a Ph.D. at Carnegie Mellon. Even then AI was breaking up into specializations -- machine vision, speech recognition, robotics. Feigenbaum went to Stanford, where he focused on expert systems. Working with colleagues in other departments, he created applications in organic chemistry, molecular biology and medicine. In the late 1970s Feigenbaum and his graduate students launched a handful of companies that applied expert system technology to business problems. Two of them, Teknowledge and Intellicorp, are still in existence.

A new generation of Feigenbaum's students are leading a second wave of artificial intelligence companies. They differ from their predecessors in one very significant way -- they don't call themselves AI.

"The new generation are not nerds, they're not entranced by the technology -- they're emphasizing business solutions," says Feigenbaum.

"AI got a bad rap," says Monte Zweben, founder and chief executive of Blue Martini Software, a San Mateo, California-based startup that uses AI technology for data mining, sifting through the ages, zip codes, buying habits and so forth of millions of consumers for retailers over the Web. "There was all this hype. Then the AI community didn't deliver."

Zweben, 34, left Stanford in 1990 with a master's degree. He continued his work at NASA, designing an expert system for planning and scheduling the maintenance of the space shuttle. The application saved millions of dollars for NASA -- and made millions for Zweben, who commercialized the system at a startup called Red Pepper Software, which he sold to PeopleSoft Inc. in 1996 for \$250 million. PeopleSoft now uses the technology to streamline the production supply chain at companies like Toyota and Bausch & Lomb.

"There has always been this battle in the AI community between those who want to develop a unified theory of cognition and those who are interested in solving business problems," Zweben says. "It's intellectually such a grand goal to understand the human mind. But there's also a lot of value in an application that solves a business problem."

That limited promise appeals to a lot of companies that are seeking to maintain quality and reduce costs in a competitive business environment. For Hewlett-Packard, Trilogy Software developed an application to reduce the number of incorrectly configured workstations HP was shipping.

Trilogy's president, Joseph Liemandt, is a former Feigenbaum student who dropped out of Stanford to become an entrepreneur. Trilogy, in Austin, Texas, sells expert systems to companies such as Boeing, IBM, Lucent Technologies and Xerox, which use Trilogy software to configure airplanes, mainframe computers, phone switches and copiers.

"People are realizing that in the real world a little AI goes a long way," says Peter E. Hart, a Stanford Ph.D. and former AI researcher who now runs Ricoh Silicon Valley, Inc., a Menlo Park, California, subsidiary of the giant Japanese copier company. "Most of the valuable applications of AI have been cases where AI was an important component -- but only one component -- of a larger system."

Take the help desk program Hart developed for Ricoh's tech support folks. He and a partner sat down with Ricoh's smartest tech support staff and found out how they solved problems. They coded that know-how into software that Ricoh's support team now uses to get faster answers to customer problems.

If you use a computer you are probably using AI applications.

Consider the "What's Related" button on the latest version of Netscape Navigator, which lets the user jump from one Web page to other pages with related content -- say, skipping from milk to dairy farming but passing over a page about the Milky Way and other galaxies. Seems simple, but it isn't, as anyone who has used a search engine to call up dozens of wrong answers can tell you.

Instead of looking for key words, Netscape's expert system understands the content of pages. The button was designed by 33-year-old Ramanathan V. Guha, who completed a Ph.D. at Stanford in 1991.

The runaway popularity of the World Wide Web has provided lots of business for AI -- and AI has helped the Web become more user-friendly.

Palo Alto, California-based Teknowledge, for example, was founded 17 years ago by Feigenbaum and others from Stanford to commercialize expert systems technology. It now provides an artificial sales agent that sells stereos over the Web. "The program asks questions, figures out configurations, makes recommendations -- it's as active as a real salesperson," says Frederick Hayes-Roth, chairman and chief executive of Teknowledge.

Alan Turing asked rhetorically: Can machines think? In a limited way, they already do.