usatoday.com
THE NATION'S NEWSPAPER
50 CENTS

PORTS

By Clive Brunskill, Allsport

ıbledon's on her mind

seeks third major in a row. Analysis ■ 10C

USA TODAY
NO. 1 IN THE USA

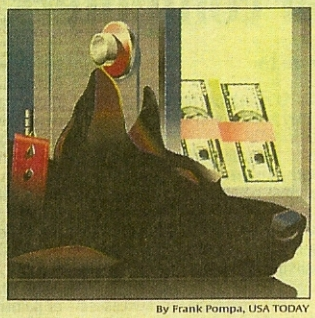

By Frank Pompa, USA TODAY

Is anyone guarding your mutual funds?

Well-paid boards of directors are supposed to, but critics say they don't. ■ Cover story, 1B

Artificial intelligence isn't just a movie

Machines, software that 'think' no longer folly of science fiction

By Kevin Maney
USA TODAY

Steven Spielberg's forthcoming *A.I.: Artificial Intelligence* is only a movie. Or is it?

The movie, set in the near future, is about a humanlike robot boy who runs on artificial-intelligence software — a computer program that doesn't just follow instructions, as today's software does, but can think and learn on its own. In some ways, the character is a fantasy. It's no closer to reality than the alien in Spielberg's earlier *E.T. the Extra-Terrestrial.*

Yet artificial intelligence is very real right now. It's far from recreating a human brain, with all its power, emotions and flexibility, though that might be possible in as little as 30 years. Today's AI can recreate slices of what humans do, in software that can indeed make decisions.

In recent years, this so-called narrow AI has made its way into everyday life. A jet lands in fog because of relatively simple AI programmed into its computers. The expertise written into the program looks at dozens of readings from the jet's instruments and decides, much as a pilot would, how to adjust the throttle, flaps and other controls.

Lately, AI has increasingly turned up in technology announcements. For example:

▶ Charles Schwab, the discount brokerage, recently said it has added AI to its Web site to help customers find information more easily.

Please see COVER STORY next page ▶

Illustration by Web Bryant, USA TODAY

History of AI

1950: Alan Turing publishes, *Computing Machinery and Intelligence.*

1956: John McCarthy coins the term, "Artificial Intelligence" at a Dartmouth computer conference.

1956: Demonstration of the first running AI program at Carnegie Mellon University.

1958: McCarthy invents the Lisp language, an AI programming language, at Massachusetts Institute of Technology (MIT).

1964: Danny Bobrow shows that computers can understand natural language enough to solve algebra word programs (MIT).

1965: Joseph Weizenbaum builds ELIZA, an interactive program that carries on a dialogue in English on any topic (MIT).

1969: Shakey, a robot, combines locomotion, perception and problem solving (Stanford Research Institute).

1979: The first computer-controlled autonomous vehicle, the Stanford Cart, is built.

1983: Danny Hillis co-founds Thinking Machines, the first company to produce massively parallel computers.

1985: The drawing program, Aaron, created by Harold Cohen, is demonstrated at AI conference.

1990s: Major advances in all areas of AI. Significant demonstrations in machine learning, intelligent tutoring, case-based reasoning, multi-agent planning, scheduling, uncertain reasoning, data mining, natural landscape understanding and translation, vision, virtual reality and games.

1997: IBM computer Deep Blue beats world champion Garry Kasparov in chess match.

Late 1990s: Web crawlers and other AI-based information-extraction programs become Web essentials.

2000: Interactive robot pets become commercially available. MIT displays Kismet, a robot with a face that expresses emotions. Carnegie Mellon robot Nomad explores remote regions of Antarctica and locates meteorites.

Source: American Association of Artificial Intelligence and Massachusetts of Technology Artificial Intelligence Lab.

Faster computers, focused programs solve early troubles with AI

Continued from 1A

► AT&T Labs is working on AI that can make robots play soccer and computer networks more efficiently.

► A computer program called Aaron, unveiled last month, has learned to make museum-quality original paintings. "It's a harbinger of what's to come," says technology pioneer Ray Kurzweil, who has licensed Aaron and will sell it to PC users. "It's another step in the blurring of human and machine intelligence."

Cover story

The commercial successes help fuel laboratory research that's pushing the fringes of AI ever closer to the equivalent of human intelligence. Software is getting better at cleverly breaking down the complex decision-making processes that go into even the simplest acts, such as recognizing a face. Hardware is marching toward brainlike capacity.

The fastest supercomputer, the IBM-built ASCI White at Lawrence Livermore National Laboratory in California, has about 1/1000th the computational power of our brains. IBM is building a new one, called Blue Jean, that will match the raw calculations-per-second computing power of a brain, says Paul Horn, IBM's director of research. Blue Jean will be ready in four years.

"Like myself, a lot of AI researchers are driven by the pursuit of someday understanding intelligence deeply enough to create intelligences," says Eric Horvitz, who was a leading scientist in AI while at Stanford University and is now at Microsoft Research in Redmond, Wash. "Many of us believe we really are on a mission."

Horvitz and others also believe this is breakthrough time for AI, when the mission spins into a wide variety of technologies.

As an area of research, AI has been around since it was first identified and given its name during a conference at Dartmouth University in 1956. It hit a peak of excitement and media attention in the mid-1980s, when AI was overhyped as a technology that was about to change the world. One fervent branch at the time was expert systems — building a computer and software that could recreate the knowledge of an expert. A brewing company, for instance, could capture a master brewer in software, possibly making human master brewers less necessary.

The exuberance was hindered by a couple of snags that led to disenchantment with AI. For one, computers of the time weren't powerful enough to even come close to mimicking a human's processing power. Two, AI was trying to do too much. Creating a complete intelligence was too hard — and still is.

Knowing one thing well

These days, that's less of a barrier. Computers have gotten exponentially more powerful every year. Now, a PC is capable of running some serious AI software. And AI researchers have learned to aim at pieces of human capacity, building software that knows it can't know everything but can know one thing really well. That's how IBM's Deep Blue beat champion Gary Kasparov in chess. Together, the developments have "led to a blossoming of real-world applications," Horvitz says.

Those applications are taking on all forms.

In Littleton, Colo., a company called Continental Divide Robotics (CDR) is a result of work done at two AI labs — one at the Massachusetts Institute of Technology and the other at the Colorado School of Mines. CDR is about to offer a system that can locate any person or object anywhere in the world and notify the user if that person or object breaks out of a certain set of rules.

One of the first uses is for tracking parolees. The parolee would wear a pager-size device that uses Global Positioning Satellite technology to know where it is. Over wireless networks, the pager constantly notifies CDR's system about its location. If the parolee leaves a certain area or gets near a certain house, the CDR software will make decisions about the severity of the violation and whom to contact. That makes it more sophisticated than the electronic anklets now used on some parolees.

CDR's technology sounds simple, but it can involve a number of fuzzy choices. If a child being tracked goes just outside his limits, the system might decide to wait to see whether he comes right back in. And it might decide whether to send you a light caution or a major warning — or to call the police. "We are literally creating software that is reactive and proactive," says Terry Sandrin, CDR's founder. "It has the ability to make decisions."

At AT&T Labs, scientist Peter Stone spends a lot of his time preparing for Robocup, an annual robotic soccer challenge coming up in August. This year, it will be in Seattle and will pit AI research labs against one another. Rolling robots the size of pint milk cartons are armed with sensors and AI software. Like real soccer players, each of the 11 robots on a team has to know its job but also must react to situations and learn about the other team. At this point, the robots can pass the ball a little but still mostly act on their own. Their capabilities are improving quickly.

It seems frivolous, but getting AI-programmed robots to work as a team to achieve something would have real-world implications. One would be making the Internet more efficient. As Stone explains it, the Net is made up of thousands of computerized routers all moving data around but acting independently. If they could act as a team, they might figure out better ways to transmit the data, which would avoid clogged areas.

Aaron takes AI to the arts, which can be a little harder to believe. But Aaron creates original work on a computer screen — quite sophisticated work. Artist Harold Cohen taught the software his style over 30 years, feeding in little by little the ways he decides color, spacing, angles and every other aspect of painting.

After all that time, the program is finally ready, and computers are powerful enough to make it work. While still in development, it won fans such as computing legend Gordon Bell. Now, Kurzweil has licensed it and plans to sell it for $19.95. Load it on a PC and let the artist loose.

"There have been various experiments with having machines be an artist, but nothing of this depth," Kurzweil says. "Cohen has created a system that has a particular style but quite a bit of diversity — a style you'd expect of a human artist."

Other uses of AI range from the amazing to the mundane.

Computer as companion

At Microsoft, Horvitz is trying to make your computer more of a companion than an inanimate tool. His software lets the computer learn about you. It learns who is important to you and who's not. It learns how to tell whether you're busy — maybe by how much you type, or by using a video camera to see whether you're staring at the computer screen or putting golf balls across the carpet.

It can combine that information to help manage your workload. If an e-mail comes in from someone very important, the computer will always put it through. If it's from someone not so important and you're busy, it can save the e-mail for later.

The software can do that with all kinds of information, including phone calls coming in and going out of your office. The thinking at Microsoft is that these capabilities might someday be a part of every computer's operating system.

Schwab's AI implementation seems less grand but no less helpful. It's using AI technology from iPhrase that can comprehend a typed sentence. More than just looking for key words, it can figure out what you really mean, even if you make spelling mistakes. So you could type, "Which of these has the most revenue?" and get the answer you were looking for. Based on the page you have up, it would know what you mean by "these." On Schwab's Web site, www.schwab.com, this is supposed to help users find information.

Beyond all the near-term uses of AI, there's the nearly unfathomable stuff.

The trends that brought AI from the failures of the mid-1980s to breakthrough success 15 years later will continue. Computers will get more powerful. Software will get more clever. AI will creep closer toward human capabilities.

If you want a glimpse of where this is heading, look inside MIT's AI lab. Among the dozens of projects there is Cog. The project is trying to give a robot humanlike behaviors, one piece at a time. One part of Cog research is focused on eye movement and face detection. Another is to get Cog to reach out and grab something it sees. Another involves hearing a rhythm and learning to repeat it on drums.

A brain like a cat's

In Belgium, Starlab is attempting to build an artificial brain that can run a life-size cat. It will have about 75 million artificial neurons, Web site Artificialbrains.com reports. It will be able to walk and play with a ball. It's supposed to be finished in 2002.

Labs all over the globe are working on advanced, brainlike AI. That includes labs at Carnegie Mellon University, IBM and Honda in Japan. "We're getting a better understanding of human intelligence," Kurzweil says. "We're reverse-engineering the brain. We're a lot further along than people think."

But can AI actually get close to human capability? Most scientists believe it's only a matter of time. Kurzweil says it could come as early as 2020. IBM's Horn says it's more like 2040 or 2050. AT&T's Stone says his goal is to build a robotic soccer team that can challenge a professional human soccer team by 2050. He's serious.

In many ways, an artificial brain would be better than a human brain. A human brain learns slowly. Becoming fluent in French can take years of study. But once one artificial brain learns to speak French, the French-speaking software code could be copied and instantly downloaded into any other artificial brain. A robot could learn French in seconds.

A tougher question is whether artificial intelligence could have emotions. No one knows.

And a frightening question is whether AI robots could get smarter than humans and turn the tables on us. Kurzweil, technologist Bill Joy and others have been saying that's possible. Horn isn't so sure. Though raw computing power might surpass the brain, he says, "that doesn't mean it will have any of the characteristics of a human being, because the software isn't there to do that."

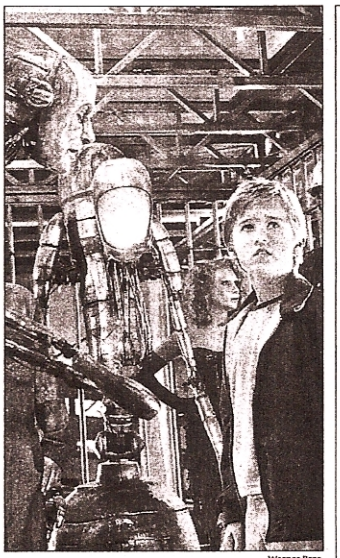

Warner Bros.

I see smart robots: Haley Joel Osment stars in *A.I.: Artificial Intelligence,* which opens June 29.

Horvitz has a brighter outlook, which at least makes the AI discussion more palatable. He says humans are always getting better at guiding and managing computers, so we'll stay in control. "Most of us (in AI) believe this will make the world a better place," he says. "A lot of goodness will come of it."

How AI could work

By Kevin Maney
USA TODAY

At Microsoft Research, scientist Eric Horvitz has been working on artificial-intelligence software that would let your PC help manage your workload.

The experimental software can learn about what you're doing at any given moment and make decisions about how to give you incoming information or messages. How it does that:

► The AI program scans the sender and text of all incoming e-mail and gives each one a score, from high priority to low. An e-mail from someone new concerning lunch next week would get a low score. An e-mail from your boss containing words such as "due today" and "fired" would get a high score.

► It would track your keyboard and mouse use, learning that how much you're typing could mean you're busy on deadline, so you don't want to get any incoming messages.

► It watches your calendar and contacts. If you're in a meeting across town with a client, your PC would forward high priority messages to your cellphone.

► A video camera on the PC would track your movements. If you're staring at the screen, it might mean you're thinking and shouldn't be disturbed.

► If you haven't moved in a long time, it might deduce that you're sleeping.

► Audio sensors would know whether you're talking on the phone or whether several people are in the room talking.

► The program could build a database about what e-mail you read and respond to and what you delete, learning what's important to you.

Using all that information, the AI program would screen incoming messages and make decisions about which ones to send you at what times.