

Artificial intelligence, once the stuff of sci-fi fantasies, is closing in on Wall Street. A corps of scientists is trying to teach computers to think like traders—and outsmart human stock pickers.

By Jason Kelly

THE ULTIMATE MONEY MACHINE

◀ Way up in a New York skyscraper, inside the headquarters of Lehman Brothers Holdings Inc., Michael Kearns is trying to teach a computer to do something other machines can't: think like a Wall Street trader.

In his cubicle overlooking the trading floor, Kearns, 44, consults with Lehman Brothers traders as Ph.D.s tap away at secret software. The programs they're writing are designed to sift through billions of trades and spot subtle patterns in world markets.

Kearns, a computer scientist who has a doctorate from Harvard University, says the code is part of a dream he's been chasing for more



MSE $\times 100$ MSE $\times 1000$

| | | |
|-----------------------------------|-----|-----|
| 1 $y = a \cdot b^x$ | 4.4 | 7.7 |
| 2 $y = (a-b)^x - a$ | 4.3 | 8.1 |
| 3 $y = a \cdot b^x$ | 3.4 | 5.9 |
| 2 $y = (a \cdot c) \cdot b^x + c$ | 3.9 | 7.0 |
| 3 | 2.5 | 4.0 |

$2 \frac{1}{2} \frac{1}{3} \frac{1}{3}$ $8 \frac{1}{3} \frac{1}{3} \frac{1}{3}$ 0 (bin)

fix value $p = P(S)$
order all leads by $\frac{v(1)}{c(p)} - \frac{v(2)}{c(p)}$

Analysis of children

- collector (vs. p.c.t.)
- semi-reflective (degree vs. clefts, influences; chain numbers, etc.)
- individual

| | |
|-----|----|
| 155 | ad |
| 152 | |
| 1 | |
| 175 | |

next experiments
Q: organ type?

| | | |
|------|----------|-----|
| K | hi | C |
| iron | in-bilan | O |
| P | low | low |

$v(1) > v(2)$

$\frac{c(p)}{v(1)}$

Michael Kearns
University of Pennsylvania

than two decades: to imbue computers with artificial intelligence, or AI. His vision of Wall Street conjures up science fiction fantasies of HAL 9000, the sentient computer in *2001: A Space Odyssey*. Instead of mindlessly crunching numbers, AI-powered circuitry one day will mimic our brains and understand our emotions—and outsmart human stock pickers, he says.

“This is going to change the world, and it’s going to change Wall Street,” says Kearns, who spent the 1990s researching AI at Murray Hill, New Jersey-based Bell Laboratories, birthplace of the laser and the transistor.

As finance Ph.D.s, mathematicians and other computer-loving disciples of quantitative analysis challenge traditional traders and money managers, Kearns and a small band of AI scientists have set out to build the ultimate money machine.

‘This is going to change the world, and it’s going to change Wall Street,’ Kearns says.

For decades, investment banks and hedge fund firms have employed quants and their computers to uncover relationships in the markets and exploit them with rapid-fire trades. Quants seek to strip human emotions such as fear and greed out of investing. Today, their brand of computer-guided trading has reached levels undreamed of a decade ago. A third of all U.S. stock trades in 2006 were driven by automatic programs, or algorithms, according to Boston-based consulting firm Aite Group LLC. By 2010, that figure will reach 50 percent, according to Aite.

AI proponents say their time is at hand. Vasant Dhar, a former Morgan Stanley quant who teaches at New York University’s Stern School of Business in Manhattan’s Greenwich Village, is trying to program a computer to predict the ways in which unexpected events, such as the sudden death of an executive, might affect a company’s stock price. Uptown, at Columbia University, computer science professor Kathleen McKeown says she imagines building an electronic Warren Buffett that would be able to answer just about any kind of investing question. “We want to be able to ask a computer, ‘Tell me about the merger of corporation A and corporation B,’ or ‘Tell me about the impact on the markets of sending more troops to Iraq,’” McKeown, 52, says.

Some executives and scientists would rather not talk

about AI. It recalls dashed hopes of artificially intelligent machines that would build cities in space and mind the kids at home. In *2001*, the novel written by Arthur C. Clarke and made into a movie directed by Stanley Kubrick in 1968, HAL, a computer that can think, talk and see, is invented in the distant future—which is 1997 in the story. Things didn’t turn out as ’60s cyberneticians predicted. Somewhere between science fiction and science fact, the dream fell apart. People began joking that AI stood for *Almost Implemented*.

“The promise has always been more than the delivery,” says Brian Hamilton, chief executive officer of Raleigh, North Carolina-based software maker Sagemworks Inc., which uses computer formulas to automatically read stock prices, company earnings and other data and spit out reports for investors.

Hamilton, 43, says today’s AI-style programs can solve specific problems within a given set of parameters. Take chess. Deep Blue, a chess-playing supercomputer developed by International Business Machines Corp., defeated world champion Garry Kasparov in 1997. The rules of chess never change, however. Players have one goal: to capture the opponent’s king. There are only so many moves a player can make, and Deep Blue could evaluate 200 million such positions a second. Financial markets, on the other hand, can be influenced by just about anything, from skirmishes in the Middle East to hurricanes in the Gulf of Mexico. In computerspeak, chess is a closed system and the market is an open one. “AI is very effective when there’s a specific solution,” Hamilton says. “The real challenge is where judgment is required, and that’s where AI has largely failed.”

AI researchers have made progress over the years. Peek inside your Web browser or your car’s cruise control, and you’ll probably find AI at work. Meanwhile, computer chips keep getting more powerful. In February, Santa Clara, California-based Intel Corp. said it had devised a chip the size of a thumbnail that could perform a trillion calculations a second. Ten years ago, such a computational feat would have required 10,000 processors.

To believers such as Dhar, Kearns and McKeown, all of this is only the beginning. One day, a subfield of AI known as machine learning, Kearns’s specialty, may give computers the ability to develop their own smarts and extract rules from massive data sets. Another branch, called natural language processing, or NLP, holds out the prospect of software that can understand human language, read up on companies, listen to executives and distill what it learns into trading programs.

Collective Intellect Inc., a Boulder, Colorado-based startup, already employs basic NLP programs to comb through 55 million Web logs and turn up information that might make money for hedge funds. “There’s some nuggets of wisdom in the sea,” says Collective Intellect Chief Technology Officer Tim Wolters.

Another AI area, neural networking, involves building silicon versions of the cerebral cortex, the part of our brain that



Vasant Dhar
New York University

governs reason. The hope is that these systems will ape living neurons, think like people and, like traders, understand that some things are neither black nor white but rather in varying shades of gray.

Stock analyst Ralph Acampora, who caused a stir in 1999 by correctly predicting that the Dow Jones Industrial Average would top 10,000, says investment banks are racing to profit from advanced computing such as AI. "It's here, and it's growing," says Acampora, 65, chief technical analyst at Knight Capital Group Inc. in Jersey City, New Jersey. "Everybody's trying to outdo everyone else."

The computers have done well. A November 2005 study by Darien, Connecticut-based Casey, Quirk & Associates, an investment management consulting firm, says that from 2001 to '05, big-cap U.S. stock funds run by quants beat those run by nonquants. The quants posted a median annualized return of 5.6 percent, while nonquants returned an annualized 4.5 percent. Both groups beat the Standard & Poor's 500 Index, which returned an annualized negative 0.5 percent during that period.

Rex Macey, director of equity management at Wilmington

Trust Corp. in Atlanta, says computers can mine data and see relationships that humans can't. Quantitative investing is on the rise, and that's bound to spur interest in AI, says Macey, who previously developed computer models at Marietta, Georgia-based American Financial Advisors LLC, to weigh investment risk and project clients' wealth. "It's all over the place and, greed being what it will, people will try anything to get an edge," Macey, 46, says. "Quant is everywhere, and it's seeping into everything."

AI proponents are positioning themselves to become Wall Street's hyperquants. Kearns, who previously ran the quant team within the equity strategies group at Lehman Brothers, splits his time between the University of Pennsylvania in Philadelphia, where he teaches computer science, and the New York investment bank, where he tries to put theory into practice. Neither he nor Lehman executives would discuss how the firm uses computers to trade, saying the programs are proprietary and that divulging information about them would cost the firm its edge in the markets.

On an overcast Monday in late January, Kearns is at work in his cubicle on the eighth floor at Lehman Brothers when a

AI: Fact and fantasy



Turing

1951 Alan Turing proposes his Turing Test for recognizing machine intelligence.

1956 The term "artificial intelligence" is coined; the first running AI program is demonstrated.

1958 Lisp, an AI programming language, is invented.


1964 Computers understand natural language enough to solve algebra word problems.

1965 ELIZA, an interactive program, proves capable of carrying on a dialogue in English on any topic.



2001: A Space Odyssey

1968 Stanley Kubrick's *2001: A Space Odyssey* features intelligent computer HAL 9000.



Star Wars

1977 *Star Wars* droids C-3PO and R2-D2 become pop icons.

1979 The Stanford Cart, an autonomous vehicle, navigates across a room full of obstacles.

1987 Automated selling helps trigger **stock market crash**.



'87 crash

1990s AI advances in machine learning, case-based reasoning, data mining, virtual reality, games.




Deep Blue

1997 IBM computer **Deep Blue** beats world chess champion **Garry Kasparov**.

Late 1990s Web crawlers, other AI-based information programs become Internet mainstays.

1999 Sony **AIBO**, a robotic pet dog, understands 100 voice commands, learns and matures.



2001 Steven Spielberg's *Artificial Intelligence: AI* is nominated for two Oscars.

2006 Carnegie Mellon opens its Machine Learning Department.

few members of his team drop by for advice. At Lehman, Kearns is the big thinker on AI. He leaves most of the actual programming to a handful of Ph.D.s, most of whom he's recruited at universities or computer conferences. Kearns himself was plucked from Penn. Ian Lowitt, who studied with Kearns at the University of Oxford and is now co-chief administrative officer of Lehman Brothers, persuaded him to come to the firm as a consultant in 2002.

Kearns hardly looks the part of a professor. He has closely cropped black hair and sports a charcoal gray suit and a crisp blue shirt and tie. At Penn, his students compete to design trading strategies for the Penn-Lehman Automated Trading Project, which uses a computerized trading simulator.

Tucking into a lunch of tempura and sashimi at a Japanese restaurant near Lehman Brothers, Kearns says AI's failure to live up to its sci-fi hype has created many doubters on Wall Street. He says people should be skeptical: Trading requires institutional knowledge that is difficult, if not impossible, to program into a computer.

AI holds perils as well as promise for Wall Street, Kearns says. Right now, even sophisticated AI programs lack common sense, he says. "When something is going awry in the markets, people can quickly sense it and stop trading," he says. "If you have completely automated something, it might not be able to do that, and that makes you subject to catastrophic risk."

The dream of duplicating human intelligence may be as old as humanity itself. The intellectual roots of AI go back to ancient myths and tales such as Ovid's story of Pygmalion, the sculptor who fell so in love with his creation that the gods brought his work to life. In the 19th century, English mathematician and proto-computer scientist Charles Babbage originated the idea of a programmable computer.

It wasn't until 1951, however, that British mathematician Alan Turing proposed a test for a machine's capability for thought. In a paper titled "Computing Machinery and Intelligence," Turing, a computer pioneer who'd worked at Bletchley Park, Britain's World War II code-breaking center, suggested the following:

A human judge engages in a text-only conversation with two parties, one human and the other a machine. If the judge can't reliably tell which is which, the machine passes and can be said to possess intelligence. No computer has ever done that. Turing committed suicide in 1954. Two years later, computer scientist John McCarthy coined the phrase *artificial intelligence* to refer to the science of engineering thinking machines.

The Turing Test, as it's now known, has fueled almost six decades of controversy. Some computer scientists and philosophers say human-like interaction is essential to human-like intelligence. Others say it's not. The debate still shapes AI research and raises questions about whether traders' knowledge, creativity, intuition and appetite for risk can ever be programmed into a computer.



Thomas Mitchell
Carnegie Mellon University

During the 1960s and '70s, AI research yielded few commercial applications. As Wall Street firms deployed computer-driven program trading in the '80s to automatically execute orders and allow arbitrage between stocks, options and futures, the AI world began to splinter. Researchers broke away into an array of camps, each focusing on specific applications rather than on building HAL-like machines. Some scientists went off to develop computers that could mimic the human retina in its ability to see and recognize complex images such as faces. Some began applying AI to robotics. Still others set to work on programs that could read and understand human languages.

Thomas Mitchell, chairman of the Machine Learning Department at Carnegie Mellon University in Pittsburgh, says many AI researchers have decided to reach for less and accomplish more. "It's really matured from saying there's one big AI label to being a little more refined and realizing there are some specific areas where we really have made progress," Mitchell, 55, says.

Financial service companies have already begun to deploy basic machine-learning programs, Kearns says. Such programs typically work in reverse to solve problems and learn from mistakes. Like every move a player makes in a game of chess, every trade changes the potential outcome, he says. Machine-learning algorithms are designed to examine possible scenarios at every point along the way, from beginning to middle to end, and figure out the best choice at each moment. Kearns likens the process to learning to play chess. "You would never think about teaching a kid to play chess by playing in total silence and then saying at the end, 'You won' or 'You lost,'" he says.

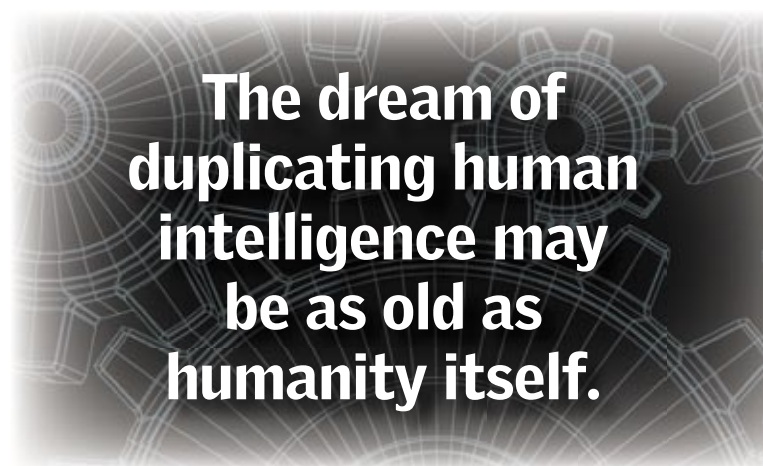
As an exercise, Kearns and his colleagues at Lehman Brothers used such programs to examine orders and improve how the firm executes trades, he says. The programs scanned bids, offers, specific prices and buy and sell orders to find patterns in volatility and prices, he says. Using this information, they taught a computer how to determine the most cost-effective trades. The program worked backward, assessing possible trades and enabling trader-programmers to evaluate the impact of their actions. By working this way, the computer learns how to execute trades going forward.

Language represents one of the biggest gulfs between human and computer intelligence, Dhar says. Closing that divide would mean big money for Wall Street, he says. Unlike computers, human traders and money managers can glimpse a CEO on television or glance at news reports and sense whether news is good or bad for a stock. In conversation, a person's vocal tone or inflection can alter—or even reverse—the meaning of words. Let's say you ask a trader if he thinks U.S. stocks are cheap and he responds, "Yeah, right." Does he mean stocks are inexpensive or, sarcastically, just the opposite? What matters is not just what people say, but how they say it. Traders also have a feel for what other investors are thinking, so they can make educated guesses about how people will react.

For Dhar, the markets are the ultimate AI lab. "Reality is the acid test," says Dhar, a 1978 graduate of the Indian Institutes of Technology, or IIT, whose campuses are India's best schools for engineering and computer science. He collected his doctorate in artificial intelligence from the University of Pittsburgh.

A professor of information systems at Stern, Dhar left the school to become a principal at Morgan Stanley from 1994 to '97, where he founded the data-mining group and focused on automated trading and the profiling of asset management clients. He still builds computer models to help Wall Street firms predict markets and figure out clients' needs. Since 2002, his models have correctly predicted the stock prices from month to month 61 percent of the time, he says.

Dhar says AI programs typically start with a human hunch about the markets. Let's say you think that rising volatility in stock prices may signal a coming "breakout," Wall Street-speak for an abrupt rise or fall in prices. Dhar says he would select market indicators for volatility and stock prices, feed them into his AI algorithms and let them check whether that intuition is right. If it is, the program would



**The dream of
duplicating human
intelligence may
be as old as
humanity itself.**

look for market patterns that hold up over time and base trades on them.

Surrounded by stacks of papers and books in his Greenwich Village office, Dhar, wearing jeans and a black V-neck sweater, says many AI scientists are questing after NLP programs that can understand human language. "That's the next frontier," he says.

At Columbia, McKeown leads a team of researchers trying to make sense of all the words on the Internet. When she arrived at the university 25 years ago, NLP was still in its infancy. Now, the Internet has revolutionized the field, she says. Just about anyone with a computer can access news reports, blogs and chat rooms in languages from all over the world.

Rather than flowing sequentially, from point A to point B, information moves around the Web haphazardly. So, instead of creating sequential rules to instruct computers to read the information, AI specialists create an array of rules and try to enable computers to figure out what works.

McKeown, who earned her doctorate from Penn, has spent the past 10 years developing a program called NewsBlaster,



Kathleen McKeown
Columbia University

which collects and sorts news and information from the Web and draws conclusions from it.

Sitting in her seventh-floor office in a building that's tucked behind Columbia's Low Library, McKeown describes how NewsBlaster crawls the Web each night to produce summaries on topics from politics to finance. She decided to put the system on line after the terrorist attacks of Sept. 11, 2001, to monitor the unfolding story.

'We want to ask a computer, "Tell me about the merger of corporation A and corporation B,"' McKeown says.

NewsBlaster, which isn't available for commercial use, can "read" two news stories on the same topic, highlight the differences and describe what's changed since it last scanned a report on the subject, McKeown says. The program can be applied to market-moving topics such as corporate takeovers and interest rates, she says.

McKeown is trying to upgrade her program so it can answer broad "what-if" questions, such as, "What if there's an earthquake in Indonesia?" Her hope is that one day, perhaps within a few years, the program will be able to write a few paragraphs or pages of answers to such open-ended questions.

Dhar says computer scientists eventually will stitch together advances in machine learning and NLP and set the combined programs loose on the markets. A crucial step will be figuring out the types of data AI programs should employ. The old programmer principle of GIGO—garbage in, garbage out—still applies. If you tell a computer to look for relationships between, say, solar flares and the Dow industrials and base trades on the patterns, the computer will do it. You might not make much money, however. "If I give an NLP algorithm ore, it might give me gold," Dhar says. "If I give it garbage, it'll give me back garbage."

Collective Intellect, financed by Denver-based venture capital firm Appian Ventures Inc., is trying to sell hedge funds and investment banks on NLP technology. Wolters says traders and money managers simply can't stay on top of all the information flooding the markets these days.

Collective Intellect seeds its NLP programs with the names of authors, Web sites and blogs that its programmers think might yield moneymaking information. Then, the company lets the programs search the Web, make connections and come up with lists of sources they can monitor and update. Collective Intellect is pitching the idea to hedge funds, Wolters says.

Technology has upended the financial services industry before. Just think of automated teller machines. Michael Thiemann, CEO of San Diego-based hedge fund firm Investment Science Corp., likens traditional Wall Street traders to personal loan officers at U.S. banks back in the '80s. Many of these loan officers lost their jobs when banks began assigning scores to customers based on a statistical analysis of their credit histories. In the U.S., those are known as FICO scores, after Minneapolis-based Fair Isaac Corp., which developed them.

Computers often did a better job of assessing risk than human loan officers, Thiemann, 50, says. "And that is where

Wall Street is going,” he says. Human traders will still provide insights into the markets, he says; more and more, however, those insights will be based on data rather than intuition.

Thiemann, who has a master’s degree in engineering from Stanford University and an MBA from Harvard Business School, knows algorithms. During the ’90s, he helped HNC Software Inc., now part of Fair Isaac, develop a tracking program called Falcon to spot credit card fraud. Falcon,

you let it go through, but maybe follow up with a call a day or so later.”

Thiemann says he’s taking a similar approach with a trading system he’s building. He calls his program Deep Green. The name recalls IBM’s Deep Blue—and money. His program evaluates market data, learns from it and scores trading strategies for stocks, options and other investments, he says. Thiemann declines to discuss his computerized

hedge fund, beyond saying that he’s currently investing money for friends and family and that he plans to seek other investors this year. “This is hard, like a moon launch is hard,” Thiemann says of the task ahead of him.

As AI invades Wall Street, even the quants will have to change with the times. The kind of conventional trading programs that hunt out arbitrage opportunities between stocks, options and futures, for example, amount to brute-force computing. Such programs, much like Deep Blue, merely crunch a lot of numbers quickly. “They just have to be fast and comprehensive,” Thiemann says. AI systems, by contrast, are designed to adapt and learn as they go.

Dhar says he doubts thinking computers will displace human

traders anytime soon. Instead, the machines and their creators will learn to work together. “This doesn’t get rid of the rule of human creativity; it actually makes it more important,” he says. “You have to be in tune with the market and be able to say, ‘I’m smelling something here that’s worth learning about.’”

At Collective Intellect, Vice President Darren Kelly, a former BMO Nesbitt Burns Inc. stock analyst, says tomorrow’s quants will rely on AI to spot patterns that no one has imagined in the free-flowing type of information that can be found in e-mails, on Web pages and in voice recordings. After all, such unstructured information accounts for about 80 percent of all the info out there. “The next generation of quant may be around unstructured analytics,” Kelly says.

After more than 50 years, the quest for human-level artificial intelligence has yet to yield its HAL 9000. Kearns says he’d settle for making AI pay off on Wall Street.

“We’re building systems that can wade out in the human world and understand it,” Kearns says. Traders may never shoot the breeze with a computer at the bar after work. But the machines just might help them pay the bill. ▶

JASON KELLY covers technology at Bloomberg News in Atlanta.
jkelly14@bloomberg.net

For the top technology-related news stories, type TOP TECH <Go>.



Michael Thiemann
Investment Science Corp.

which today watches over more than 450 million credit and debit cards, uses computer models to evaluate the likelihood that transactions are bogus. It weighs that risk against customers’ value to the credit card issuer and suggests whether to let the charges go through or terminate them. “If it’s a customer with a questionable transaction and you don’t mind losing them as a customer, you just deny it,” Thiemann says. “If it’s a great customer and a small transaction,

**Artificial intelligence
researchers are
jockeying to become
the hyperquants
of Wall Street.**