

PETER LYNCH ANSWERS READER MAIL • JIM ROGERS BETS ON CONFUCIUS

Worth

FINANCIAL INTELLIGENCE

Downloading

WARREN BUFFETT'S

Brain

Can a computer beat the Master?

**GET READY
FOR THE
CAPITAL-GAINS
TAX CUT**

**THE PRICE
OF HAPPINESS**

**THE SUPER
INVESTORS OF
SILICON VALLEY**

DECEMBER

SIMSON GARFINKEL
PO BOX 4188
VINEYARD HAVN MA 02568-4188

SEP 96

PDST

#EEL10

30

#GRFB4188S99 0#451956



0928101798111111
F-DIGIT 02568

0 0928101798111111

US \$3.00 CANADA \$3.50

**12 EXPERT
STOCK PICKS**



DOWNLOADING WARREN

WARREN BUFFETT IS, AS EVERYONE KNOWS, AN AVUNCULAR, RUMPLED, unpretentious man who is also, in the world of investing, the nearest equivalent to God. Forty years ago he started an investment partnership with \$100,000, then bought a small textile firm called Berkshire Hathaway. Diligently ferreting out and buying the stocks of companies he judged undervalued, Buffett eventually turned Berkshire Hathaway into a \$21 billion company,

shares of which trade, as of October 16, 1995, at \$29,650. Along the way, he amassed a personal fortune of about \$12 billion.

Now meet Warren Buffett's brain. It, too, is modest and unprepossessing—a Digital 486 personal computer that sits on the desk of one David Braverman, an investment officer with Standard & Poor's in New York City. One year ago, Braverman read Robert G. Hagstrom Jr.'s best-seller, *The Warren Buffett Way: Investment Strategies of the World's Greatest Investor*. Hagstrom never met Buffett while writing the book. Instead, he read Berkshire Hathaway's annual reports—with Buffett's famously informal chairman's letters—along with Buffett's speeches and interview transcripts to uncover what had made him so spectacularly successful.

Of course, many have tried to emulate Buffett. "The traditional approach is to buy stocks that Warren has already bought," says Braverman. "My complaint with that strategy is that you are buying things that have already moved up due to his actions." Similarly, Hagstrom's book purported to divulge Buffett's secrets, but "it didn't give me any information on what stocks to buy next."

To get what he wanted, Braverman did what any intrepid programmer would do: He started writing code, turning the investing axioms in the book into a series of equations that instructed his computer to screen for stocks meeting Buffett's criteria.

"What this does," he says of his creation, "is find stocks that he has not bought but that meet criteria similar to those he buys or that he may even buy in the future.... There is obviously an advantage to this over buying what he has already bought and made money in."

Thirty companies made the grade. Then Braverman tried to out-Buffett

Buffett by screening for strong consensus earnings estimates as well. Sixteen firms remained, which Braverman felt were truly undervalued.

"It looks like we are really on to something," says Braverman. He updated the list in July (see "Be Like Buffett"), replacing six stocks with new ones chosen by the computerized screen. Overall, the stocks were up an average of 42 percent through September 30, 1995. The S&P 500 rose by a little more than 21 percent over the same period.

Braverman also tested his model against markets over the past ten years. On July 1 of each year, the computer would buy the stocks that made it through the screen and sell the ones that didn't. Once again, the computer outperformed the S&P 500, though it underperformed Buffett over that period. "Over the last six months, we've done better [than Buffett]," Braverman says.

All this, he says, is less a tribute to his software than to Buffett's value style of investing. What's "amazing," he says, is "you can get returns like this just by going back to basics." Or, as the man himself wrote in the 1994 Berkshire Hathaway annual report, "In investing it is not necessary to do extraordinary things to get extraordinary results."

Of course, as Braverman is the first to admit, Buffett does a lot more than apply a formula to determine whether a stock is undervalued. "Buffett has bought and made a substantial amount of money in Capital Cities/ABC," says

Braverman. "Capital Cities/ABC has lower margins [net income divided by sales] than most of the stocks that were picked in our screen, yet he bought Capital Cities/ABC. Why? Because he is extremely impressed with the management, and no computer, whether a neural network or a simple stock-picking screen, can necessarily quan-



David Braverman loaded Warren Buffett into his PC.

tify being impressed with [Capital Cities/ABC chairman and CEO] Thomas Murphy."

Or can it? Imagine a computer that would start out each day at midnight, scanning the Tokyo stock market. The system would also scan newspapers and analysts' reports and market data from throughout the industrialized countries, looking for companies that seemed to be performing unusually well. The program—call it E-Buff-

fett—would also read interviews with top corporate management and, through high-quality speech synthesis and a speech-recognition device connected to a telephone, even arrange for its own interviews with Thomas Murphy about Capital Cities/ABC's merger with Disney. Murphy, incidentally, would never realize that he was schmoozing with a silicon chip, not flesh and blood.

A computer that can understand a language and converse in it is one of the great, elusive goals of the computer discipline called artificial intelligence. It could even be considered the goal of computer science itself. More than 150 years ago, when Charles

the Germans' top-secret Enigma code during World War II and later helped create some of the earliest electronic computers, also said that his ultimate goal was to construct a thinking machine. Indeed, Turing devised a test that scientists still use to determine if a machine is sentient. The Turing test is simple: If you are sitting at a terminal and you cannot tell if the respondent at the other end of the wire is a human or a computer, then the computer program is intelligent.

These are matters that long ago ceased to be solely of academic or philosophical import. Businesses and venture-capital firms have invested many millions of dollars in artificial intelligence. The 1980s, in fact, began a positive AI boom, fueled by advances in three areas of research: knowledge-based systems (formerly called expert systems), neural networks, and genetic algorithms. Each offers ways of systemizing human or human-like intelligence, putting those smarts into a computer, and applying it to real-world problems.

One company that has successfully employed AI technology is Neuron Data, a small software firm in Mountain View, California. Neuron Data makes a program that companies can use to create their own knowledge-based systems.

The company's \$5,000 NEXPERT Object package is widely used in the financial world, says Alan Lund-

berg, the program's product manager. It is also being employed by banks to detect fraud and in emergency medicine to set up effective triage protocols.

Such systems are composed of elaborate sets of "if, then" statements, with special computer programs that make sure the proper ones get executed at the proper time. For example, a knowledge-based system for driving a car might have rules like "if the car needs gas, then look for a gas station" and "if there is a child directly in front of the

car, then step on the brake." By repeatedly interviewing people who are familiar with a particular domain of knowledge and collecting a large set of these rules of thumb, a programmer can endow an expert system with almost human-like decision-making capability.

However, knowledge-based systems encounter a problem when applied to trading: In order to build such programs, it's necessary for a "knowledge engineer"—the person who actually writes the expert system—to sit down with a "domain expert"—someone who understands the task that is to be computerized—and come up with each of the hundreds or thousands of "if, then" statements that make up the rule-based system. This is tedious and time-consuming, but more to the point, in the fast-paced, high-pressure world of financial trading, few of the domain experts—that is, the traders—understand what they are doing well enough to come up with a formal set of rules.

Christine Downton, a partner with London-based Pareto Partners and a former central banker and money manager, went through the process in 1992 when her firm formed a partnership with Hughes Aircraft Co. Hughes AI researchers painstakingly worked their way through Downton's expertise in a number of different markets, creating a global-investment knowledge-based system that can track up to 18 different variables in the currency and fixed-income markets of 12 separate economies. Each of the thousands of rules that constitute the program required a precise and quantifiable definition, even for terms like "high growth," "medium term," and "normal."

Today, Downton's disembodied brain manages \$110 million, and Pareto hopes to raise that to \$250 million by the end of the year. The program has "outperformed the relevant benchmarks," she says.

So knowledge-based systems have some utility. Unfortunately, like many human experts, they don't learn. They can't respond to changing conditions. And they don't know when they are making a mistake. Finally, says Andrew Lo, a professor at MIT's Sloan School of Management, "Expert sys-

BE LIKE BUFFETT



David Braverman's "Warren Buffett" stock picks

| NAME | RECENT PRICE |
|-------------------------------|--------------|
| OAK INDUSTRIES (OAK) | \$22 |
| COMPUTER ASSOCIATES (CA) | 49 |
| STURM, RUGER & CO. (RGR) | 27 |
| UST INC. (UST) | 30 |
| FRANKLIN RESOURCES (BEN) | 54 |
| INTERNATIONAL GAME TECH (IGT) | 13 |
| PFIZER (PFE) | 61 |
| AMGEN (AMGN) | 46 |
| GREEN TREE FINANCIAL (GNT) | 30 |
| MICRON TECHNOLOGY (MU) | 70 |
| MEDIA GENERAL (MEG/A) | 31 |
| LAWTER INTERNATIONAL (LAW) | 11 |
| AMERICAN HOME PRODUCTS (AHP) | 88 |
| INTEL (INTC) | 67 |
| SCHERING-PLOUGH (SGP) | 54 |
| SOFTKEY INT'L (SKEY) | 39 |

Price as of October 20, 1995. The stocks are listed in order of relative strength, as assessed by the program.

Babbage, an English mathematician, embarked on his project to build a steam-powered computing machine, his dream was to build a machine that could think. In a letter dated June 21, 1833, Lady Augusta Ada Byron, daughter of the poet and, as Babbage's collaborator, the world's first computer programmer, described Babbage's invention as a "thinking machine (for such it seems)."

Alan Turing, the British mathematician who aided the effort to crack



t MIT's Andrew Lo (left) dreams of thinking software. Michael de la Maza has moved from theory to reality with the Redfire fund.

tems are only as expert as the [person creating the rules]." And that can be very difficult to assess.

Michael Prietula has spent years studying stock traders and investment bankers. Both a psychologist and a computer scientist, Prietula is a professor at the Center for Accounting Research and Professional Education at the University of Florida at Gainesville.

According to Prietula, even big trading firms can't predict who will be a good trader. "I went into a trading house and had about 50 people on the floor. Then I said, 'Where are your superstars?' There were maybe 3 out of all those 50. And I said, 'Well, how do you trade?' And they said, 'You got to be real careful.'

"There is a type of cognition called situated cognition," Prietula says. "A lot of things in the environment contribute to your performance. We just aren't exactly sure about that environment, except that it changes a lot. Somehow the good traders are good at perceiving and adapting to the change."

Of course, learning unwritten rules is fundamental to many human endeavors. It's what learning and thinking are all about. It follows, then, that creating computers that can learn is another element of AI's overarching ambition of replicating human intelligence.

Over the past ten years, a relatively simple AI technique for finding rules of thumb and unseen correlations has gained many adherents. Called neural networking, this technique bears little resemblance to the way a human brain works. Indeed, the only net-

works in most neural systems are complex networks of nonlinear mathematical equations.

Linear relationships are the kind of mathematical rules that most people learn in grade school. Nonlinear relationships, on the other hand, describe things like S curves, normal distributions, and exponential growth. "Warren starts with \$100,000, works for ten years, and earns \$100 million. Thirty years from now, how much is Warren worth?" If the answer is \$400 million, then Warren's work-to-earnings relationship is linear. But if Warren is worth \$12 billion, the relationship is definitely nonlinear.

Neural networks do an excellent job of representing complex combinations of nonlinear equations. Even better, thanks to a mathematical technique discovered in the 1980s called back propagation, it's possible to teach a neural network to "learn" a complex relationship between a set of variables—even find relationships humans

can't see on their own.

To make a basic neural network, a programmer constructs a series of equations with a few dozen input variables, some nonlinear functions, a few dozen more tuning parameters, and a very small number of outputs. This sounds more abstruse than it is. For instance, a bank might use a neural network to determine the credit-worthiness of a loan applicant. "Inputs" would be values such as the applicant's household income, age, gender, and financial makeup. The "output" might be a single number—the predicted likelihood of the applicant defaulting on the loan. To train the network, the bank might subject the system to all of its loan records, hoping that the neural network could learn its own tuning parameters—its way of figuring the odds—more accurately than a domain expert.

That, at least, is the ideal. Casey Klimauskas, product manager at NeuralWare, a Pittsburgh firm that makes a neural-network construction kit for personal computers, says it's still "more an art than a science" to make these programs work effectively. You can't just plug them in.

NeuralWare's product, called Neural Works Predict, is being used to forecast the rate of return on individual stocks and to calculate the amount of risk in a portfolio of stocks and bonds. Klimauskas estimates that there are perhaps 50 people around the world who are using trainable neural networks for portfolio management.

Another AI technique for picking stocks is being pioneered by Michael de

**SOMETIMES
BRAD LEWIS WILL
OVERRULE
THE COMPUTER.
"USUALLY,"
HE SAYS, "IT'S
A MISTAKE."**

la Maza and Deniz Yuret, two graduate students at the MIT Artificial Intelligence Laboratory. A couple of years ago, de la Maza and Yuret were sitting

their own money at the end of 1993. In November 1994 they got a two-paragraph mention in *Upside* magazine that said Redfire was generating 8 percent monthly returns. This was followed by a brief story in *Worth* (February 1995), and the calls started coming in. This March, with 14 investors, the two MIT students turned Redfire into a hedge

technology is Brad Lewis, a fund manager at Fidelity Investments with \$4.3 billion under management. "I'm running five funds controlled by something that is a form of AI," says Lewis. Until February, the funds were controlled by a neural-network system, but he recently changed over to a new technology that he won't reveal. "It is a pretty heavy-duty form of AI," Lewis says. "The fact that it is still very nonlinear, that's the claim to fame of all of these neural nets: nonlinearity."

NEURAL-NETWORK TECHNOLOGY

"IS TOTALLY WORTHLESS

SAYS LOUIS NAVELLIER.

"WE'VE BEEN TESTING IT."

around bemoaning the state of artificial-intelligence research. "So we made a list of problems that we could work on to apply AI techniques to the real world," says de la Maza. The two settled on the idea of an autonomous computer program that would learn to make money investing in the market.

They code-named their project Redfire.

For the first six months of 1993, de la Maza and Yuret developed a program using the most "gee-whiz" AI technology they could find: genetic algorithms. Genetic algorithms create hundreds or thousands of randomly generated computer programs, then try out each one. Some programs work; most don't. The computer then takes pieces from the successful programs to make new programs, sort of interbreeding them. At the end of a few thousand generations, the resulting programs can work quite well. Frequently it's hard to figure out why.

The beauty of genetic algorithms is that de la Maza and Yuret didn't have to know why. All they had to do was feed their program with the prices at the close of the market. When the program wanted to make a trade, it printed a message on the screen.

Once the system was operational, de la Maza and Yuret did a six-month dry run, then put up \$10,000 of

fund in order to avoid SEC reporting regulations. The fund is brokered by Furman Selz of New York.

The two founders are keeping their returns confidential. De la Maza will only say they are "positive."

Another trader using this sort of AI

THERE ARE PLENTY OF STORIES ON Wall Street about kids without a college education making a million dollars their first year as a trader and theoreticians with Ph.D.s in economics losing their shirts. These stirring tales may amuse traders having a beer after work, but they're positively depressing for college professors. After all, if there's no benefit to a college education (at least on Wall Street), then why pay for one?

That question is driving John O'Brien, an associate professor of accounting and finance at Carnegie Mellon University's Graduate School of Industrial Administration. O'Brien is director of technology at CMU's Center for Financial Analysis and Securities Trading. FAST, in effect, is a **cyber-classroom** whose purpose "is designed around integrating theory with practice," says O'Brien. "If we teach a concept, then it is immediately applied."

The program gives would-be traders an intimate knowledge of networked computer systems, program trading, and derivatives, and makes them feel at home with the mathematical foundations of market dynamics. O'Brien calls his fledgling market masters "the Cyber Warriors."

The program teaches students to be traders the same way Wall Street does: putting them in front of a terminal and telling them to

TRADES FOR GRADES

make money. Of course, since O'Brien is at CMU, the money isn't **cash—it's grades**. "You earn grade cash," says O'Brien. "Grade cash accumulated over the duration of a course translates into a percentage of their grade. That makes it very real and competitive."

The trading is grouped into a series of projects, some running for a few hours, others running for weeks or an entire semester. Students trade against each other, against traders built with artificial-intelligence technology, and against students at other schools: Copies of the trading room are now running in Japan, England, Hong Kong, Russia, Ukraine, Mexico, Australia, Singapore, and South Korea.

Each project is designed to reveal an important facet of the market. In one project, says O'Brien, students start out with "a very undiversified position, and through trading they can hedge their risks." **As** the project continues, students **get to see**

how their actions determine the prices of securities in the virtual market. "You can diversify at stupid prices and make yourself worse off, or you can diversify at reasonable **prices and make yourself better off**," O'Brien says.

Those with access to the Internet can learn more about the FAST program on the World Wide Web. Just point your web browser to <http://fastweb.gsia.cmu.edu>.—S.G.



Cyber Warrior John O'Brien

times he 'will even overrule the computer. "Usually when I do, it is a mistake," Lewis says. "It seems that way. I am overriding the box a lot less than I used to seven years ago."

One problem with this sort of investing, says Lo at MIT, is that it ignores the factors that drive financial prices. All too often, says Lo, people attempting to do computer-based investing get caught up with their numbers and forget the realities for which those numbers stand.

"Ultimately," says Lo, "the problem of making sound financial investment decisions is not a statistical problem but an economic problem, and although statistical inference and AI technology can add tremendous value, they alone will never provide a complete solution for financial-market participants—at least not until software becomes self-aware."

Of course, there's no mention of nonlinearities or neural nets in Berkshire Hathaway's annual reports. Nor will you find them in the work published by Louis Navellier, editor of *MPT Review* and president of Navellier & Associates. Navellier manages \$1.2 billion.

Neural-network technology is "totally worthless," says Navellier. "We've been testing it." The problem, he says, is the time frame. "Neural nets might be good for short-term trend followers," but they can't pick up long-term trends.

Navellier is also suspicious about groups like Redfire that are claiming outrageously high earnings. He wants to know the liquidity limit of such technology. While it might be possible to get a 100 percent or 200 percent yearly return on a million dollars, try to run the same strategies with a billion dollars and you'll start to affect the market, rather than making money from it.

There is another big difference be-

tween the AI-based trading systems in use today and our hypothetical E-Buffer program. A truly intelligent artificial trader would probably watch TV, listen to the radio, and read the newspa-


ly surmising that demand for oil would rise with a nuclear power plant suddenly down and that so many potato fields would be contaminated that potato futures would climb.

Indeed, it is possible that in order to mimic Warren Buffett's investment skills, it would be necessary for a computer program to share the man's interest in the Omaha Royals minor-league baseball team. But for most people working on AI investing, things like knowledge, news events, and even institutional memory are just noise—something to be shut out.

Katia Sycara, a professor of computer science at Carnegie Mellon University, is trying to keep the noise. "Most of the computerized investment strategies are sophisticated analytical techniques for prediction," says Sycara. "What we would like to do is have some more content-based ways, rather than just statistical manipulation."

To build that system, Sycara and her students are working with a new kind of AI technology: intelligent, autonomous agents. An agent is a computer program that specializes in a small set of tasks and is linked to a bunch of other agents. The agents can ask questions of one another, send answers, and work with one another toward a common goal. The initial system developed by Sycara will have five main agents: a portfolio-management agent; fundamental-analysis agent; a technical-analysis agent; a financial-news agent; and an analysts'-estimate tracking agent. Eventually there may be dozens of agents, each one running on its own separate computer.

But where do the knowledge and heuristics given to the agents come from? From successful investors—both human and computer. The advantage of using an agent-based approach is that different agents can be built with different kinds of AI technologies—one agent can be built with a neural network, another with genet-



IT IS POSSIBLE THAT TO MIMIC BUFFETT, A COMPUTER MUST SHARE HIS INTEREST IN THE OMAHA ROYALS.



Katia Sycara (top) of Carnegie Mellon thinks computers need noise. Her CMU student traders compete in a global virtual market.

per, like Alexander, the trader in Michael Lewis's book *Liar's Poker*. When the Chernobyl nuclear reactor exploded, Alexander immediately bought oil and potato futures, correct-

ic algorithms, another using expert systems, and perhaps others using case-based reasoning. The system can then look at new situations as they arise and use the agents that have done well in the past with similar situations.

Sycara's research, incidentally, is funded by the Department of Defense's Advanced Research Projects Agency, the folks who created the ARPANET, precursor to the Internet. The reason: The techniques that she is developing for rapidly evaluating and acting upon incomplete and

sometimes contradictory information that arrives in real time are equally applicable to the stock market and the battlefield, both of which are forever shrouded in what Clausewitz called "the fog of war."

Whether or not the technical wizards ever reach the Holy Grail of a thinking computer, there is irony in the attempt to clone the minds of old-fashioned investors like Warren Buffett. The awesome calculating power of computers—and the ingenuity of the software creators—is directed at

ceaselessly shuffling a deck of investments, acquiring and discarding, looking to maximize returns. Even Standard and Poor's David Braverman each year sells those stocks that no longer meet his program's Buffett criteria.

But that is precisely what Buffett himself "wams against in his annual chairman's letter. Speaking for himself and Berkshire's vice chairman, Charles T. Munger, Buffett writes, "Gin rummy managerial behavior (discard your least promising business at each turn) is not our style." •

"NOW WE ENTER THE BOOLEAN TIME-SERIES CONDITIONS here," says Dave Hirschfeld, tapping at his keyboard. A three-dimensional shape appears on the screen, blue foothills and valleys mapped against a black-and-white grid, with a sizable crater in the middle. "That's the '87 crash," says Hirschfeld, who is head of research for Tudor Investments and president of Tudor Software for the legendary futures trader Paul Tudor Jones. "We can zoom in on that tick by tick if you want," Hirschfeld says as another shape appears. "Here's London"—click—"Tokyo"—click—"Frankfurt."

This is Wall Street's hottest piece of technology. Rivals refer to it as "Tudor in a box"—the brain of Paul Tudor Jones shooting off sparks in a vat. In fact, the system, in development since February 1990, is a highly sophisticated analytical tool that combines approximately 50 data streams, embracing markets from currency to cotton, every 15 minutes, every trading day, back to 1982 and before. "In 20 seconds you can use this to dismiss statements that have a lot of logic to them but no reality," says Rich Jaycocks, former head of technology for Tudor Investments, as he presses a button and instantly disproves the conventional wisdom that the currency market is unusually volatile this year.

Even more dazzling are the analytics that can say in seconds how to pursue a given strategy with the least risk and the greatest chance of profit. "We all have our views of the world," Jaycocks says. "This tells you what is the cost of your view."

One Tudor employee—who declined to be named—describes the system. "The idea was to model an expert system based on Paul's knowledge and his trades in the market," he says. "The genesis came from the chess program developed at Carnegie Mellon, DeepThought, which beat chess champion Gary Kasparov. We recruited several Carnegie people. We went to Paul: We reviewed his past trades. And then we'd say, 'Okay, here's a market, what would you do?' We'd then take what he did and see if it applied to other markets, and so on."

All futures trading is based on the belief that past performance is predictive—that if 20 times over the past 20 years a Fed rate increase coupled with a rising dollar pushed up the price of

WALL STREET'S HOTTEST HOT ROD

coal in Germany, which in turn pushed down the price of German utility stocks, then the same variables, behaving in the same way, will produce the same result. One might, then, short German utilities.

This illustrates the potential advantage of computerized trading systems over their human counterparts: the ability to instantaneously calculate the effect, on a historical basis, of the innumerable variables—political, environmental, economic, psychological, etc.—that affect the odds of a given trade or investment paying off.

Not everyone agrees with this. As the late Fischer Black, a partner at Goldman, Sachs and one of the fathers of modern options trading, told *The Economist*, "There are things that machines are good at, but trading does not appear to be one of them.... The list of factors that matter are changing all the time."

To a limited extent, Peter Borish, who developed the system at Tudor before leaving to found Computer Trading Corp., would agree. "Paul Tudor Jones, Louis Bacon, and the other hedge-fund heads take a big speculative position, but they're actually not risking very much," he says, because they have an extraordinary ability instantly to intuit the likely impact of a given piece of news, or a market move, on the minds of several thousands of their trading peers. This is something that may never be matched by a machine.

Risk and leverage, on the other hand, can be usefully modeled, and applied by the push of a button. "Someone like Paul doesn't make one trade always based on the same things," Borish says. "So we tried to break it into subsets, to model different

ideas. And if that was successful, you could put different ideas into a portfolio and devise an allocating scheme to weight which idea is a little bit better at which time."

After that, you have to trust your program's calculations. "If someone told me—and believe me, I've heard everything—that the phases of the moon affected the soybean market, I can run a historical data series of new-moon and full-moon soybean prices," says Borish. "If there's any statistical significance to that, and if the risk-reward is there, I'd probably trade off it."

—David Samuels

