Maybe We Should Leave That Up to the Computer

Humans Have Flaws as Decision Makers

By DOUGLAS HEINGARTNER

AMSTERDAM — Do you think your high-paid managers really know best? A Dutch sociology professor has doubts.

The professor, Chris Snijders of the Eindhoven University of Technology, has been studying the routine decisions that managers make, and is convinced that computer models, by and large, can do a better job of it. He even issued a challenge late last year to any company willing to pit its humans against his algorithms.

"As long as you have some history and some quantifiable data from past experiences," Mr. Snijders claims, a simple formula will soon outperform a professional's decision-making skills. "It's not just pie in the sky," he said. "I have the data to support this."

Some of Mr. Snijders's experiments from the last two years have looked at the results that purchasing managers at more than 300 organizations got when they placed orders for computer equipment and software. Computer models given the same tasks achieved better results in categories like timeliness of delivery, adherence to the budget and accuracy of specifications.

No company has directly taken Mr. Snijders up on his challenge. But a Dutch insurer, Interpolis, whose legal aid department has been expanding rapidly in recent years, called in Mr. Snijders to evaluate a computer model it had designed to automate the routing of new cases — a job previously handled manually by the department's in-house legal staff.

The manager in charge of the project, Ludo Smulders, said the model was much faster and more accurate than the old system. "We're very satisfied about the results it's given our organization," he said. "That doesn't mean there are no daily problems, but the problems are much smaller than when the humans did it by hand. And it lets them concentrate more on giving legal advice, which is what their job is."

Mr. Snijders's work builds on something researchers have known for decades: that mathematical models generally make more accurate predictions than humans do. Studies have shown that models can better predict, for example, the success or failure of a business start-up, the likelihood of recidivism and parole violation, and future performance in graduate school.

They also trump humans at making various medical diagnoses, picking the winning dogs at the racetrack and competing in online auctions. Computer-based decision-making has also grown increasingly popular in credit scoring, the insurance industry and some corners of Wall Street.

The main reason for computers' edge is their consistency — or rather humans' inconsistency — in applying their knowledge.

"People have a misplaced faith in the power of judgment and expertise," said Greg Forsythe, a senior vice president at Schwab Equity Ratings, which uses computer models to evaluate stocks.

The algorithms behind so-called quant funds, he said, act with "much greater depth of data than the human mind can. They can encapsulate experience that managers may not have." And critically, models don't get emotional. "Unemotional is very important in the financial world," he said. "When money is involved, people get emotional." Many putative managerial qualities, like experience and intuition, may in fact be largely illusory. In Mr. Snijders's experiments, for example, not only do the machines generally do better than the managers, but some managers perform worse over time, as they develop bad habits that go uncorrected from lack of feedback.

Other cherished decision aids, like meeting in person and poring over dossiers, are of equally dubious value when it comes to making more accurate choices, some studies have found, with face-to-face interviews actually degrading the quality of an eventual decision.

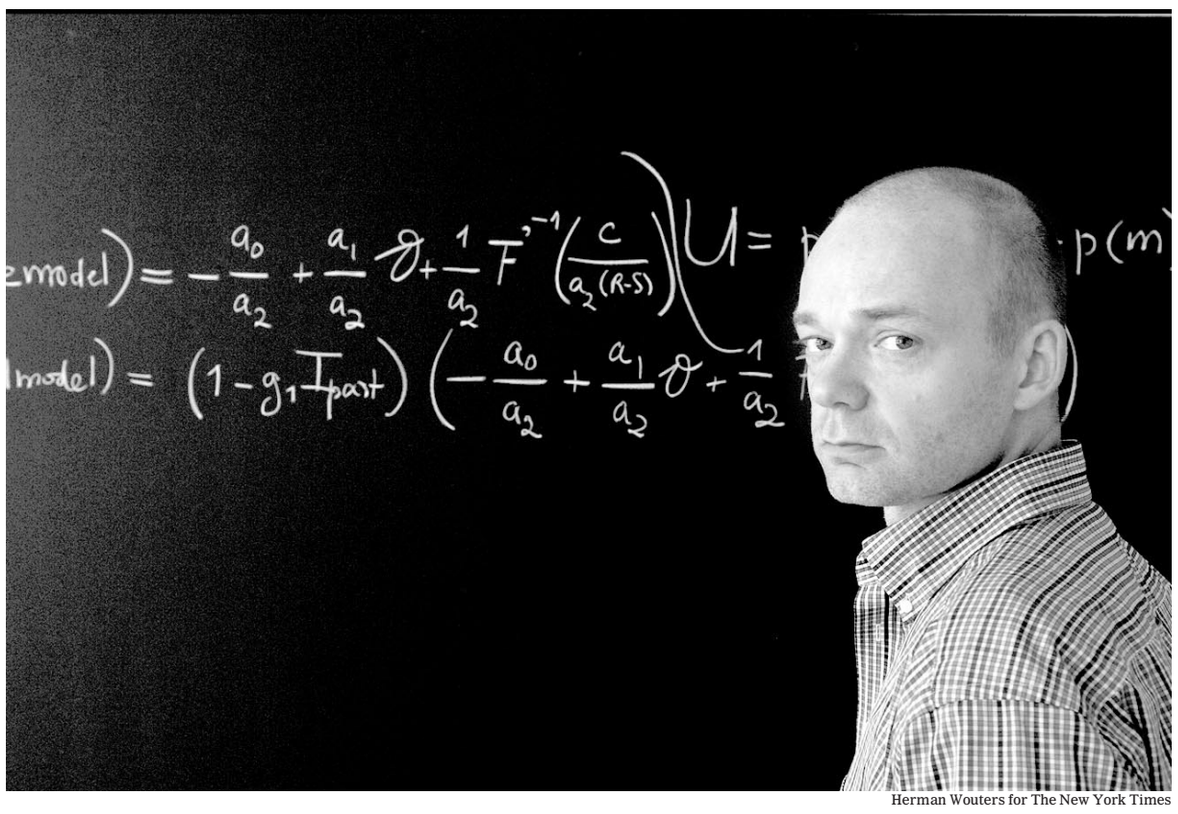

Herman Wouters for The New York Times

Chris Snijders, a Dutch sociology professor, challenges any company to pit its humans against his algorithms.

"People's overconfidence in their ability to read someone in a half-an-hour interview is quite astounding," said Michael A. Bishop, an associate professor of philosophy at Northern Illinois University who studies the social implications of these models.

And the effects can be serious. "Models will do much better in predicting violence than will parole officers, and in that case, not using them leads to a more dangerous society," he said. "But people really don't believe that the models are as accurate as they are."

Models have other advantages beyond their accuracy and consistency. They allow an organization to codify and centralize its hard-won knowledge in a concrete and easily transferable form, so it stays put when the experts move on. Models also can teach newcomers, in part by explaining the individual steps that lead to a given choice. They are also faster than people, are immune to fatigue and give the human experts more time to work on other tasks beyond the current scope of machines.

So if they're so good, why aren't they already used everywhere?

Not everyone is convinced that managers are incorrigibly myopic. "I've never seen any evidence that there is a pattern of decline at all, and it just doesn't fit with the way management literature is going, which is all around the emotional intelligence angle," said Laura Empson, the director of the Clifford Chance Center of the Said Business School at Oxford University.

"I think there are a lot of people who have a strong technological orientation who would agree life would be a lot simpler if it weren't for the humans," she said. "But the reality is, organizations do have a lot of very intense and complicated human issues within them."

Max H. Bazerman, a professor at Harvard Business School, wonders how many managerial decisions can actually be modeled. "The vast majority of decisions that we make in professional life don't have this quality," he said.

He agrees that models can make better decisions about credit card applications and college admissions, he said, "but there are many decisions that are much more unique, where that database doesn't exist. I'm as skeptical about human intuition as these folks, but it's not only a computer model that we replace it with. Sometimes it's thinking more clearly."

Many in the field of computer-assisted decision-making still refer to the debacle of Long Term Capital Management, a highflying hedge fund that counted several Nobel laureates among its founders. Its algorithms initially mastered the obscure worlds of arbitrage and derivatives with remarkable skill, until the devaluation of the Russian ruble in 1998 sent the fund into a tailspin.

"As long as the underlying conditions were in order, the computer model was almost like a money machine," said Roger A. Pielke Jr., a professor of environmental studies at the University of Colorado whose work focuses on the relation between science and decision-making. "But when the assumptions that went into the creation of those models were violated, it led to a huge loss of money, and the potential collapse of the global financial system."

In such situations, "you can never hope to capture all of the contingencies or variables inside of a computer model," he said. "Humans can make big mistakes also, but humans, unlike computer models, have the ability to recognize when something isn't quite right."

Another problem with the models is the issue of accountability. Mr. Forsythe of Schwab pointed out that "there's no such thing as a 100 percent quantitative fund," in part because someone has to be in charge if the unexpected happens. "If I'm making decisions," he said, "I don't want to give up control and say, 'Sorry, the model told me.' The client wants to know that somebody is behind the wheel."

Still, some consider the continuing ascendance of models as inevitable, and recommend that people start figuring out the best way to adapt to the role reversal. Mark E. Nissen, a professor at the Naval Postgraduate School in Monterey, Calif., who has been studying computer-vs.-human procurement, sees a fundamental shift under way, with humans becoming increasingly peripheral in making routine decisions, concentrating instead on designing ever-better models.

"The newest space, and the one that's most exciting, is where machines are actually in charge, but they have enough awareness to seek out people to help them when they get stuck," he said — for example, when making "particularly complex, novel, or risky" decisions.

The ideal future, then, may lie in letting computers and people each do what they do best. One way to facilitate this development is to train people to identify the typical cognitive foibles that lead to bad choices. "I've now worked with these models for so long," Mr. Snijders said, "that my instincts have changed along the way."

As Mr. Bishop of Northern Illinois University puts it, by making smart use of computer models' advantages, "you'll become like the crafty A student who doesn't work that hard but gets mostly right answers, rather than the really hard-working student who gets lots of wrong answers and as a result gets C's."

From Fantasy and Fact

Artificial intelligence has come far since its inception, both in its practical applications and in the popular imagination.

1950 The mathematician Alan Turing proposes a test for machine intelligence.

1956 John McCarthy coins the term "artificial intelligence" as the topic of the Dartmouth Conference, the first conference devoted to the subject.

1959 Arthur Samuel's checkers program wins games against the best human players.

1962 First industrial robot company, Unimation, founded.

1967 "HAL" (above) stars in "2001: A Space Odyssey."

1969 Stanford Research Institute: Shakey the Robot demonstrated combining movement, perception and problem solving.

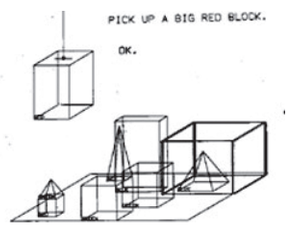

1971 Terry Winograd's Ph.D. thesis (M.I.T.) demonstrated the ability of computers to understand English sentences in a restricted world of children's blocks, in a coupling of his language-understanding program with a robot arm that carried out instructions typed in English.

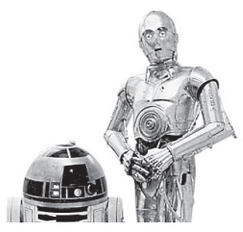

1977 C3PO and R2D2 star in "Star Wars."

1981 The Fifth Generation Computer Systems project by the Ministry of International Trade and Industry of Japan heralds a generation of intelligent machines.

1985 A Kawasaki robot kills a Japanese mechanic during a malfunction.

1997 I.B.M.'s Deep Blue chess program beats the world chess champion Garry Kasparov in a widely followed match.

2005 Stanley, a robot Volkswagen Touareg, designed by a team of Stanford University engineers, wins the DARPA Grand Challenge award by traveling autonomously for 132 miles through a desert.

The New York Times/Photographs by Associated Press (HAL), Professor Terry Winograd (blocks), 20th Century Fox (Star Wars), Agence France-Presse (Deep Blue), PRNewsFoto (Volkswagen)

As Artificial Intelligence Gains, Brainy Robots Start Stepping Into Daily Life

Continued From Page A1

ticated. Today some scientists are beginning to use the term cognitive computing, to distinguish their research from an earlier generation of artificial intelligence work. What sets the new researchers apart is a wealth of new biological data on how the human brain functions.

"There's definitely been a palpable upswing in methods, competence and boldness," said Eric Horvitz, a Microsoft researcher who is president-elect of the American Association for Artificial Intelligence. "At conferences you are hearing the phrase 'human-level A.I.,' and people are saying that without blushing."

Cognitive computing is still more of a research discipline than an industry that can be measured in revenue or profits. It is pursued in various pockets of academia and the business world. And despite some of the more startling achievements, improvements in the field are measured largely in increments: voice recognition systems with decreasing failure rates, or computerized cameras that can recognize more faces and objects than before.

Still, there have been rapid innovations in many areas: voice control systems are now standard features in midpriced automobiles, and advanced artificial reason techniques are now routinely used in inexpensive video games to make the characters' actions more lifelike.

A French company, Poseidon Technologies, sells underwater vision systems for swimming pools that function as lifeguard assistants, issuing alerts when people are drowning, and the system has saved lives in Europe.

Last October, a robot car designed by a team of Stanford engineers covered 132 miles of desert road without human intervention to capture a $2 million prize offered by the Defense Advanced Research Projects Agency, part of the Pentagon. The feat was particularly striking because 18 months earlier, during the first such competition, the best vehicle got no farther than seven miles, becoming stuck after driving off a mountain road.

Now the Pentagon agency has upped the ante: Next year the robots will be back on the road, this time in a simulated traffic setting. It is being called the "urban challenge."

At Microsoft, researchers are working on the idea of "predestination." They envision a software program that guesses where you are traveling based on previous trips, and then offers information that might be useful based on where the software thinks you are going.

Tellme Networks, a company in Mountain View, Calif., that provides voice recognition services for both customer service and telephone directory applications, is a good indicator of the progress that is being made in relatively constrained situations, like looking up a phone number or transferring a call.

Tellme supplies the system that automates directory information for toll-free business listings. When the service was first introduced in 2001, it could correctly answer fewer than 37 percent of phone calls without a human operator's help. As the system has been constantly refined, the figure has now risen to 74 percent.

More striking advances are likely to come from new biological models of the brain. Researchers at the École Polytechnique Fédérale de Lausanne in Lausanne, Switzerland, are building large-scale computer models to study how the brain works; they have used an I.B.M. parallel supercomputer to create the most detailed three-dimensional model to date of a column of 10,000 neurons in the neocortex.

"The goal of my lab in the past 10 to 12 years has been to go inside these little columns and try to figure out how they are built with exquisite detail," said Henry Markram, a research scientist who is head of the Blue Brain project. "You can really now zoom in on single cells and watch the electrical activity emerging."

Blue Brain researchers say they believe the simulation will provide fundamental insights that can be applied by scientists who are trying to simulate brain functions.

Another well-known researcher is Robert Hecht-Nielsen, who is seeking to build an electronic butler called Chancellor that would be able to listen, speak and provide in-home concierge services. He contends that with adequate resources, he could create such a machine within five years.

Although some people are skeptical that Mr. Hecht-Nielsen can achieve what he describes, he does have one successful artificial intelligence business under his belt. In 1986, he founded HNC Software, which sold systems to detect credit card fraud using neural network technology designed to mimic biological circuits in the brain. HNC was sold in 2002 to the Fair Isaac Corporation, where Mr. Hecht-Nielsen is a vice president and leads a small research group.

Last year he began speaking publicly about his theory of "confabulation," a hypothesis about the way the brain makes decisions. At a recent I.B.M. symposium, Mr. Hecht-Nielsen showed off a model of confabulation, demonstrating how his software program could read two sentences from The Detroit Free Press and create a third sentence that both made sense and was a natural extension of the previous text.

For example, the program read: "He started his goodbyes with a morning audience with Queen Elizabeth II at Buckingham Palace, sharing coffee, tea, cookies and his desire for a golf rematch with her son, Prince Andrew. The visit came after Clinton made the rounds through Ireland and Northern Ireland to offer support for the flagging peace process there."

The program then generated a sentence that read: "The two leaders also discussed bilateral cooperation in various fields."

Artificial intelligence had its origins in 1950, when the mathematician Alan Turing proposed a test to determine whether or not a machine could think or be conscious. The test involved having a person face two teleprinter machines, only one of which had a human behind it. If the human judge could not tell which terminal was controlled by the human, the machine could be said to be intelligent.

In the late 1950's a field of study emerged that tried to build systems that replicated human abilities like speech, hearing, manual tasks and reasoning.

During the 1960's and 1970's, the original artificial intelligence researchers began designing computer software programs they called "expert systems," which were essentially databases accompanied by a set of logical rules. They were handicapped both by underpowered computers and by the absence of the wealth of data that today's researchers have amassed about the actual structure and function of the biological brain.

Those shortcomings led to the failure of a first generation of artificial intelligence companies in the 1980's, which became known as the A.I. Winter. Recently, however, researchers have begun to speak of an A.I. Spring emerging as scientists develop theories on the workings of the human mind. They are being aided by the exponential increase in processing power, which has created computers with millions of times the power of those available to researchers in the 1960's — at consumer prices.

"There is a new synthesis of four fields, including mathematics, neuroscience, computer science and psychology," said Dharmendra S. Modha, an I.B.M. computer scientist. "The implication of this is amazing. What you are seeing is that cognitive computing is at a cusp where it's knocking on the door of potentially mainstream applications."

At Stanford, researchers are hoping to make fundamental progress in mobile robotics, building machines that can carry out tasks around the home, like the current generation of robotic floor vacuums, only more advanced. The field has recently been dominated by Japan and South Korea, but the Stanford researchers have sketched out a three-year plan to bring the United States to parity.

At the moment, the Stanford team is working on the first steps necessary to make the robot they are building function well in an American household. The team is focusing on systems that will consistently recognize standard doorknobs and is building robot hands to open doors.

"It's time to build an A.I. robot," said Andrew Ng, a Stanford computer scientist and a leader of the project, called Stanford Artificial Intelligence Robot, or Stair. "The dream is to put a robot in every home."

Peter DaSilva for The New York Times

Andrew Ng of Stanford University, left, leads a project to create a robot with artificial intelligence.