ieee spectrum
THE MAGAZINE OF TECHNOLOGY INSIDERS
06.08

WILL SUPERSMART MACHINES LET US LIVE FOREVER OR RENDER US OBSOLETE?

THE RAPTURE OF THE GEEKS

SEPARATING SCIENCE FROM FICTION IN THE TECHNOLOGICAL SINGULARITY

A SPECIAL REPORT

IEEE

CLICK HERE TO ACCESS THE JUNE ISSUE OF the institute


F-35
AA-1

volume 45 number 6 international

ieee spectrum 06.08

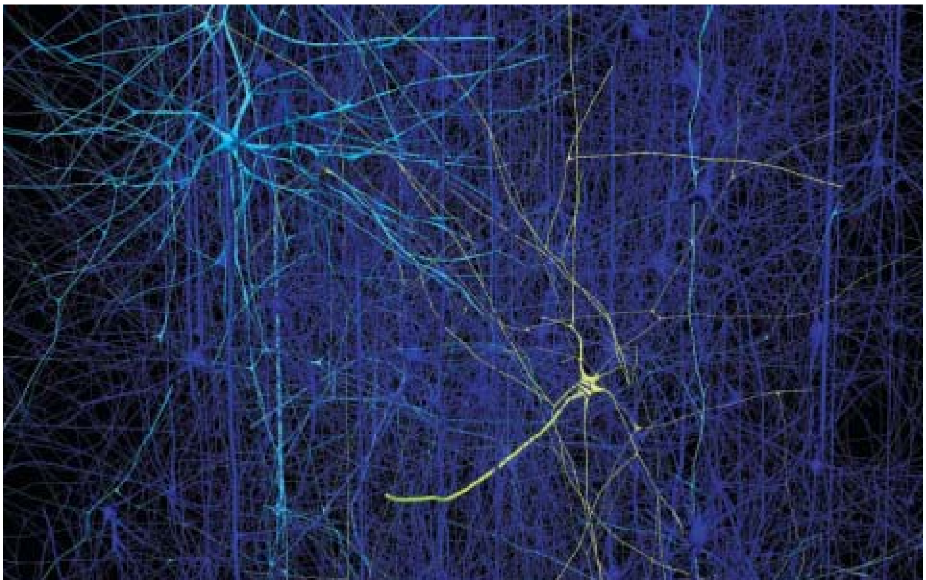

BRAIN EMULATION for machine intelligence is now under way. These blue and yellow traces, [above] formed in a computer simulation, show excited neurons in a rat neocortex.

COVER: BRYAN CHRISTIE DESIGN

THIS PAGE: PABLO DE HERAS CIECHOMSKI/ VISUALBIOTECH

SPECIAL REPORT: THE SINGULARITY

UPDATE

OPINION

DEPARTMENTS

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

volume 45 number 6 international

ieee spectrum 06.08

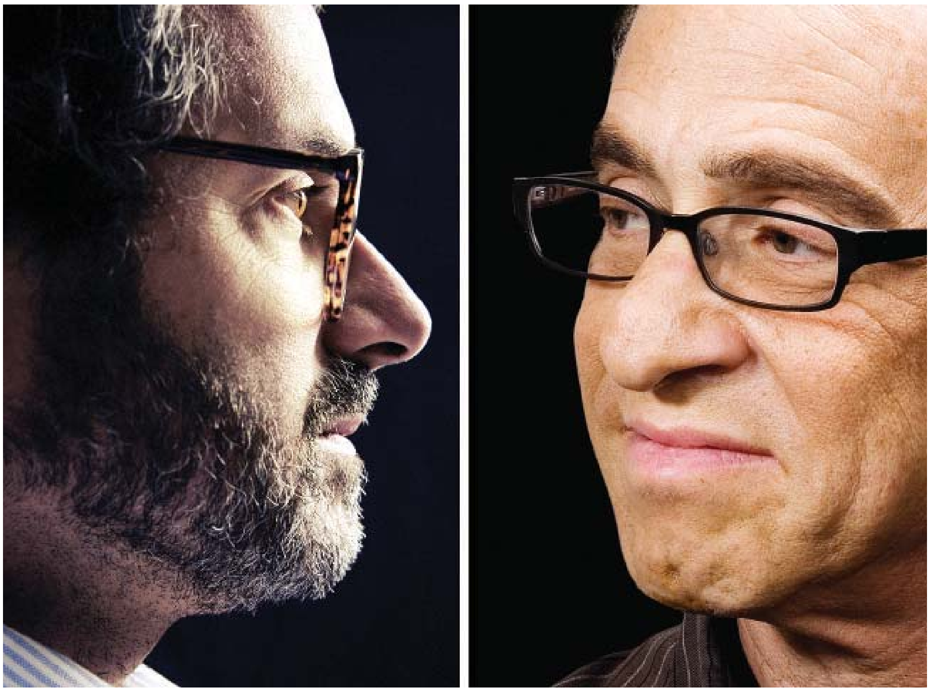

PHOTOS, FROM LEFT: STEVE PYKE/CONTOUR BY GETTY IMAGES; GREGG SEGAL; BILL CRAMER

WWW.SPECTRUM.IEEE.ORG

AVAILABLE 1 JUNE ON SPECTRUM ONLINE

TWO PATHS TO THE SINGULARITY

MIT professor Neil Gershenfeld [left] and technology futurist Ray Kurzweil have long worked at the leading edge of physical science and computer science. Today both believe that we are on the event horizon of a technological singularity. David Dalrymple, a child prodigy who claims both these luminaries as mentors, discovers that they came to this conclusion from two very different directions.

ONLINE FEATURES:

CAUGHT ON TAPE: Author Vernor Vinge, roboticist Rodney Brooks, and neuroscientist Christof Koch expand on their ideas about the singularity in exclusive video interviews with *IEEE Spectrum* editors Harry Goldstein and Erico Guizzo.

MAPPING THE BODY ELECTRIC: Human senses and body parts are increasingly augmented by a stunning array of high-tech devices. View our three-dimensional model to gauge progress toward the age of cyborgs.

ALSO ONLINE:

- Webcasts
- Radio
- News
- Audio Downloads
- Podcasts
- Jobs
- Career Accelerator
- IEEE Xplore® Digital Library
- Interviews
- Opinions
- More!

WWW.IEEE.ORG/THEINSTITUTE

AVAILABLE 6 JUNE ON THE INSTITUTE ONLINE

PART-TIME PASSIONS

IEEE members don't just work all day. Some have intriguing pastimes, like restoring old Mustangs and singing and playing drums in samba groups.

Q&A WITH KAM AND RAY

The IEEE annual election is right around the corner, so *The Institute* decided to get personal with the two candidates running for president-elect, Moshe Kam and Pedro Ray. Read about their most challenging projects, favorite books, hidden talents, and more.

ONLINE DIRECTORY OPENS DOORS TO NETWORKING

Ever wish you could connect with other IEEE members who share your interests? Want to find members who graduated from your alma mater or belong to the same technical society you do? IEEE memberNet, a new online directory, makes it all possible.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. 3 Park Avenue, New York, NY 10016-5997, U.S.A. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS-IEEE Members: $21.40 included in dues. Libraries/institutions: $205. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at W224-N3322 Duplainville Rd., Pewaukee, WI 53072-4195, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of American Business Media, the Magazine Publishers of America, and the Society of National Association Publications.

back story

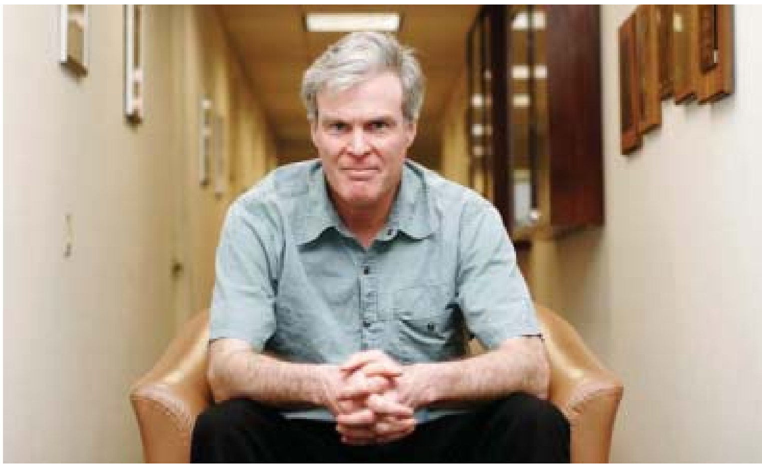

Of Two Minds

We knew early on that the lead article in this issue, which describes theories about how the brain creates the mind, was going to be an unusual challenge. It would have to explain one of the most elusive subjects in all of science, and it would also have to take a critical look at claims that technologists are on the verge of creating a mind *in silico*.

Executive Editor Glenn Zorpette knew exactly who should write it. "John Horgan is probably the only writer who is smart enough, cranky enough, and skilled enough to pull it off," he remembers thinking.

Horgan [above] and Zorpette worked together in the mid-1980s at *IEEE Spectrum*. They would interview technologists and officials during the day and drink beer, eat Dominican food, and argue about politics after work. Once, having become lost driving to Albuquerque from Los Alamos National Laboratory, in New Mexico, they stumbled into a huge, exuberant ceremony in the desert with hundreds of Native Americans. Amid the feathers, bells, and body paint, the two journalists were conspicuous in blue blazers, wing-tip shoes, and Ray-Bans.

It was at *Spectrum* that Horgan began forging the probing, impious style that characterizes his best works. He did smart, tough pieces on underground nuclear testing, arms control, and biomedical devices. Later, at *Scientific American* magazine, where he and Zorpette worked and were occasional hockey teammates in the mid-1990s, he profiled scientists, technologists, and philosophers in stories that were free of the deference typical of that kind of article.

"At some point I felt that I could serve science better if I were skeptical rather than reverential," says Horgan, who now directs the Center for Science Writings at the Stevens Institute of Technology, in Hoboken, N.J. He is the author of three books and is a contributor to Bloggingheads.tv.

Zorpette says, "When I was in my early 20s, John made me understand that no career was more fun, worthwhile, and interesting than print journalism. His curiosity, skill, and fierce intelligence are what I have been trying to match for my whole career." ❑

CITING ARTICLES IN IEEE SPECTRUM

IEEE Spectrum publishes two editions. In the international edition, the abbreviation INT appears at the foot of each page. The North American edition is identified with the letters NA. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, the first Update page is in *IEEE Spectrum*, Vol. 45, no. 6 (INT), June 2008, p. 11, or in *IEEE Spectrum*, Vol. 45, no. 6 (NA), June 2008, p. 13.

ieee spectrum

EDITORIAL

EDITOR IN CHIEF Susan Hassler, s.hassler@ieee.org

EXECUTIVE EDITOR Glenn Zorpette, g.zorpette@ieee.org

MANAGING EDITOR Elizabeth A. Bretz, e.bretz@ieee.org

SENIOR EDITORS Harry Goldstein (Online), h.goldstein@ieee.org; Jean Kumagai, j.kumagai@ieee.org; Samuel K. Moore (News), s.k.moore@ieee.org; Tekla S. Perry, t.perry@ieee.org; Philip E. Ross (Resources), p.ross@ieee.org; William Sweet, w.sweet@ieee.org

SENIOR ASSOCIATE EDITOR Steven Cherry, s.cherry@ieee.org

ASSOCIATE EDITORS Erico Guizzo, e.guizzo@ieee.org; Joshua J. Romero (Online), j.j.romero@ieee.org; Sandra Upson, s.upson@ieee.org

ASSISTANT EDITOR Willie D. Jones, w.jones@ieee.org

SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org

COPY EDITOR Michele Kogon, m.kogon@ieee.org

EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org

ADMINISTRATIVE ASSISTANTS Ramona Gordon, r.gordon@ieee.org; Nancy T. Hantman, n.hantman@ieee.org

IEEE SPECTRUM JOURNALISM INTERN Sally Adee, s.adee@ieee.org

INTERN Francesco Ferorelli, f.ferorelli@ieee.org

CONTRIBUTING EDITORS John Blau, Robert N. Charette, Peter Fairley, Alexander Hellemans, David Kushner, Robert W. Lucky, Paul McFedries, Kieron B. Murphy, Carl Selinger, Seema Singh, John Voelcker

ART & PRODUCTION

SENIOR ART DIRECTOR Mark Montgomery

ASSISTANT ART DIRECTORS Laura H. Azran, Brandon Palacio

PHOTO EDITOR Randi Silberman

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy

EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia

SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani

WEB PRODUCTION COORDINATOR Jacqueline L. Parker

WEB PRODUCTION SPECIALIST Michael Spector

EDITORIAL CORRESPONDENCE

IEEE Spectrum, 3 Park Ave., 17th Floor, New York, NY 10016-5997
Attn: Editorial Dept. Tel: +1 212 419 7555 Fax: +1 212 419 7570
Bureau: Palo Alto, Calif.; Tekla S. Perry +1 650 328 7570
Responsibility for the substance of articles rests upon the authors, not the IEEE or its members. Articles published do not represent official positions of the IEEE. Letters to the editor may be excerpted for publication.

ADVERTISING CORRESPONDENCE

IEEE Spectrum, 3 Park Ave., 17th Floor, New York, NY 10016-5997
Attn: Advertising Dept. +1 212 419 7760
The publisher reserves the right to reject any advertising.

RANDI SILBERMAN

contributors

RODNEY BROOKS is a machine. Or so he says in his article on p. 62. A professor at the Massachusetts Institute of Technology, he researches the engineering of intelligent robots capable of operating in real-world environments and how to understand human intelligence by building humanoid robots. Brooks is also the chief technical officer of iRobot Corp.

ROBIN HANSON, author of "The Economics of the Singularity" [p. 36], is an associate professor of economics at George Mason University, in Fairfax, Va. He trained as a physicist and worked in the aerospace industry before getting a Ph.D. in social science at Caltech in 1997. He is a pioneer of idea futures markets, among them the Foresight Exchange Prediction Market.

RICHARD A.L. JONES, author of "Rupturing the Nanotech Rapture" [p. 56], is a professor of physics at the University of Sheffield, in England, and senior nanotechnology advisor for the UK government's physical sciences and engineering funding agency. His book *Soft Machines: Nanotechnology and Life* (2004) argues that nanotechnology needs to learn as much from biology as from engineering.

CHRISTOF KOCH is a professor of cognitive and behavioral biology at Caltech.

GIULIO TONONI is a professor of psychiatry at the University of Wisconsin, Madison. In "Can Machines Be Conscious?" [p. 46], the two neuroscientists discuss how to assess synthetic consciousness. Koch became interested in the physical basis of consciousness while suffering from a toothache. Why should the movement of certain ions across neuronal membranes in the brain give rise to pain? he wondered. Or, for that matter, to pleasure or the feeling of seeing the color blue? Contemplating such questions determined his research program for the next 20 years.

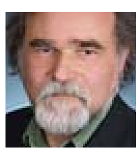

ALFRED NORDMANN, author of "Singular Simplicity" [p. 52], is a professor of philosophy and the history of science at Darmstadt Technical University, in Germany. His interests include the philosopher Ludwig Wittgenstein, the physicist and philosopher of science Heinrich Hertz, and the birth of new scientific disciplines, such as nanotechnology.

VERNOR VINGE, who wraps up this issue on p. 68, first used the term *singularity* to refer to the advent of superhuman intelligence while on a panel at the annual conference of the Association for the Advancement of Artificial Intelligence in 1982. Three of his books—*A Fire Upon the Deep* (1992), *A Deepness in the Sky* (1999), and *Rainbows End* (2006)—won the Hugo Award for best science-fiction novel of the year. From 1972 to 2000, Vinge taught math and computer science at San Diego State University.

CLOCKWISE FROM TOP LEFT: IROBOT CORP.; GIULIO TONONI; ALFRED NORDMANN; VERNOR VINGE; CALTECH; D. JONES; GEORGE MASON UNIVERSITY

For more special features related to our singularity report,
go to http://www.spectrum.ieee.org/jun08/singularityspecialreport

spectral lines

Un-assuming The Singularity

GIVEN THE current state of computer science and robotics, it's hard to understand how "the singularity" meme has become lodged in the serious discourse of the technosphere. This is the idea that, as a consequence of exponentially accelerating technological innovation and continuously self-improving artificial intelligence, computer power will outstrip human brainpower, leading to the end of human culture as we know it. Not a century from now, mind you, but somewhere between 2030 and 2045, depending on whom you talk to.

BOSTON DYNAMICS' DARPA-funded robotic dog, BigDog.
PHOTO: BOSTON DYNAMICS

The concept was framed in its most tech-savvy form by computer scientist and science-fiction writer Vernor Vinge in 1983 in *Omni* magazine. It has since morphed into a complicated "theory" that for some, notably prolific inventor Ray Kurzweil, includes a posthuman afterlife in which we abandon our biological selves and are uploaded into digital and possibly robotic vessels, there to spend eternity as cybernetic Methuselahs. It is also thought by its followers to be inevitable, not merely one of many possible future scenarios.

The singularity represents an untestable set of assumptions about our near future. So why are so many willing to take it seriously? That's what we set out to discover in our special report, "The Rapture of the Geeks," in this issue. Given that it's the 25th anniversary of Vinge's seminal work, it seemed like a good time to call upon the science and technology experts—including Vinge—to get a sense of the merits and the demerits of the singularity case. We were particularly interested to learn what, if any, technology supports the extraordinary claims made by the singularity's proponents.

What we found is that there's a lot of hyperbole distracting us from the real work under way in nanotechnology, brain implants, and machine learning. Researchers are, with some success, making machines more intelligent and responsive to solving real-world problems. The explosion of disciplines involved in these pursuits gives you some sense of their complexity. Robotics departments have now added *developmental*, *epigenetic*, or *evolutionary* to their names; control and systems are becoming more and more intelligent; AI is coursing through the blood of embodied cognitive science.

But we're still a very long way from understanding how consciousness arises in the human brain, let alone figuring out how to re-create it in a machine. We're even a long way from the much simpler goal of creating autonomous, self-organizing, and perhaps even self-replicating machines.

Simple locomotion—like walking—has only recently been conquered by roboticists. And there's still a lot of work to be done to integrate walking with other functions, like seeing and hearing. An example that's been causing a sensation on YouTube is Boston Dynamics' DARPA-funded robotic dog duo, BigDog and LittleDog. They are the braindogs of IEEE member and company founder Marc Raibert, who figured out in the 1980s that robotic running could be controlled by a few decoupled control laws.

BigDog—which actually looks more like a large headless spider that's been chopped in half—is designed to help soldiers in the field carry heavy equipment and supplies across rocky terrain. It can run, walk, and climb and is able to right itself after stepping into a hole or tripping over a branch, but it can't figure out how to avoid these obstacles—at least not yet.

LittleDog, on the other hand, is learning to "see" its environment before taking a walk. Software and sensors help the robot evaluate the surface it's about to step on so it can decide where to go next. The interesting thing, as our journalism intern Sally Adee noted in her blog post about LittleDog, is just how much processing time it takes the robopup to make "walk here, not there" decisions. It's certainly not ready for Tokyo's busy crosswalks.

This is hard scientific and technical work, and it involves the deep understanding and use of dozens of biological and electromechanical systems. The ability to remain upright while avoiding potholes is important, but it's just an infinitesimal step along the road to anything resembling human consciousness.

Wireless communication, ubiquitous computing, nanotechnology, distributed sensing, and embedded systems are going to converge and deliver wonders. Electronic prosthetic devices and biopharmaceuticals will help us correct or expand our physical and mental capabilities. Ultimately, we may even learn enough about consciousness to re-create it in a machine and create artificial vessels for our own minds. But with all we have to do over the next 30 to 40 years, we don't expect to be hitting the "Upload to digital heaven" button anytime soon.

Do you? Let us know. —SUSAN HASSLER

forum

LETTERS do not represent opinions of the IEEE. Short, concise letters are preferred. They may be edited for space and clarity. Additional letters are available online in "And More Forum" at http://www.spectrum.ieee.org. Write to Forum, *IEEE Spectrum*, 3 Park Ave., 17th Floor, New York, NY 10016-5997, U.S.A.; fax, +1 212 419 7570; e-mail, n.hantman@ieee.org.

ALL THE NEWS THAT FITS...YOU

GREG LINDEN'S article, "People Who Read This Article Also Read..." [March] presents personalized news article recommendations as the logical extension of the book and movie recommendation systems used by sites like Amazon and Netflix. News, however, is more than entertainment—it largely defines people's perceptions of the world. News recommendation algorithms based entirely on user preferences are therefore inappropriate. For example, say an algorithm determines that a user enjoys articles featuring mostly positive information about a certain political or religious figure. Would the algorithm then avoid showing negative articles about that figure to that user?

If a handful of personalized news-aggregation sites do indeed become the public's primary portal to information on current events, the risk of bias and censorship—inadvertent or intentional—will be tremendous. The role of a free press sometimes includes telling people exactly what they don't want to hear.

MICHAEL RUTBERG
IEEE Member
New York City

WHERE HAVE ALL THE ENGINEERS GONE?

THE U.S. culture of abundance and short-term thinking is at fault for the dearth of engineers decried in Robert W. Lucky's column ["U. S. Engineers and the Flat Earth," Reflections, March]. U.S. society reached a tipping point about 20 years ago, when we became so wealthy as a nation that we stopped deferring our gratification to a future generation and began to believe that we could have it all in our lifetime. Today we lack scientists and engineers who have the persistence to endure the challenges and frustration of difficult but engaging work, because the societal goals have changed. In the 1960s there was a national purpose—to respond to Soviet challenges, win the Cold War, and put a man on the moon before the end of the decade. Now we have no overweening objective to drive us except "consume mass quantities" and the ethos of "the one who dies with the most toys wins."

We need a new Marshall Plan for education that will inspire our youth to strive for goals beyond portable gadgets and flashy video games. Technology drives progress, but technology requires a numerate, literate populace, and we are eating our own seed corn if we do not plant the love of knowledge and science early on in the hearts and minds of our young people.

ROSS BETTINGER
IEEE Member
Seattle

FIGHTING OBSOLESCENCE

I JUST READ "Trapped on Technology's Trailing Edge" [April]. Until I retired, I worked with military missiles and space applications. One of the techniques engineers use to get around the problem of obsolete parts is to insist that the interfaces meet well-established standards. That way, subsystems can be replaced with new components without replacing the entire system—or so the idea goes. Unfortunately, the standards now seem to have a lifetime not much longer than the components'. It is also notable that the standards often change solely for competitive reasons, not because of technical inadequacy.

NORMAN WORTH
IEEE Life Member
Los Alamos, N.M.

LOSING OUR GRIP

IN "TOP 10 Tech Cars" [April], the text concerning the 2009 Chevrolet Corvette ZR1 states: "The factory did confirm that the cornering grip is more than 1 *g*—enough to make you feel twice your weight in a perfectly banked curve." Not correct. The forces add vectorially. If the cornering (centrifugal) force is 1 *g*, then you must add it to the downward force due to gravity. The resulting force is 1.41 *g*, directed at 45 degrees to the vertical. This is also the case if the curve isn't perfectly banked, but if it is, the force is vertical on the rotated car frame, and any car could probably achieve that. It may, however, be a noteworthy achievement to manage this on a perfectly horizontal (unbanked) curve.

CLIVE WOODS
IEEE Senior Member
Baton Rouge, La.

BYE, BYE, *BURAN*

READERS OF "Copying NASA's Mistakes," the book review by James Oberg about the Soviet space shuttle [March], may be interested in a sequel. On 11 and 12 April 2008, the German news was full of reports about the delivery of a *Buran* test vehicle to the Technik Museum in Speyer, Germany, where it will be exhibited—a fate more fitting than what befell the vehicle Oberg mentions at the end of his review: being turned into a restaurant.

EDITH BORIE
IEEE Member
Karlsruhe, Germany

ILLUSTRATION: DAN PAGE

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

update

more online at http://www.spectrum.ieee.org

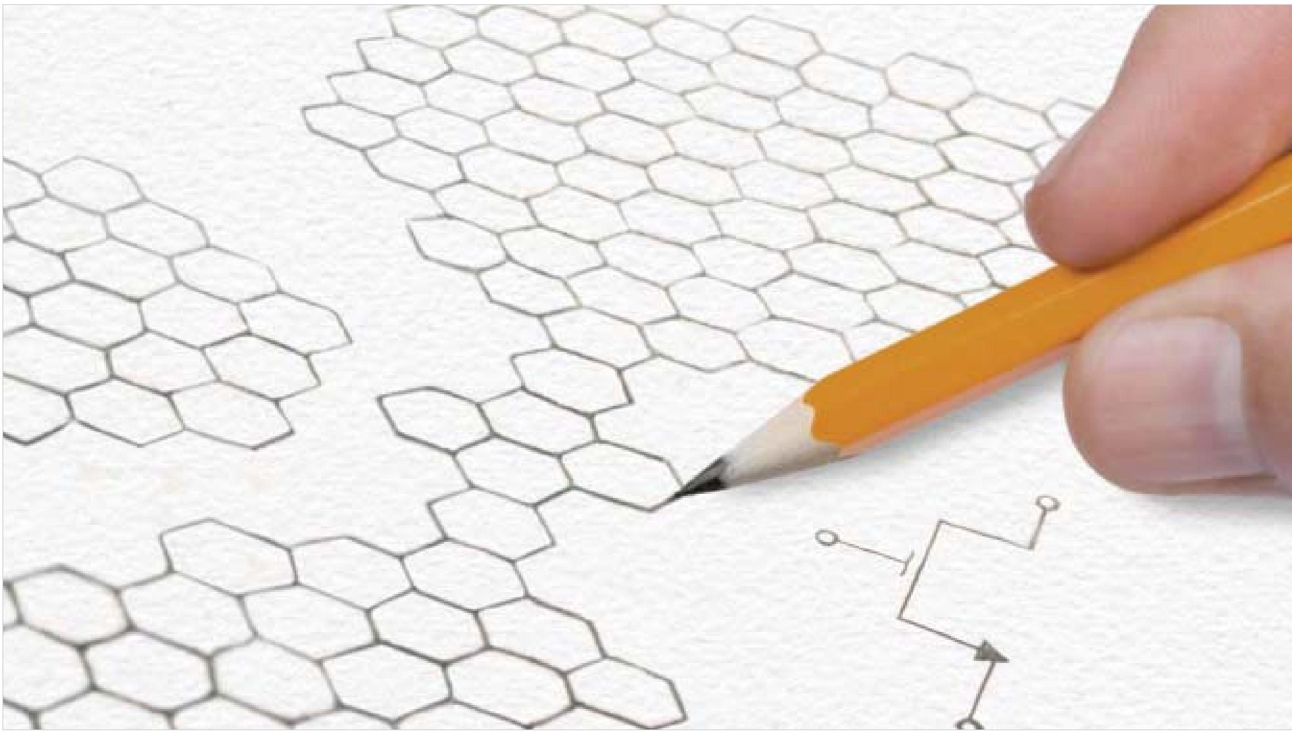

Researchers Pencil In Graphene Transistors

Graphene's weird electrical properties allow for smallest transistor yet

THE LITTLE smudges you leave behind whenever you use a pencil could be the key ingredient of the next revolution in computer circuitry, according to experts around the globe. Part of what shears off from the graphite in a pencil is a substance known as graphene, a one-atom-thick crystal with remarkable electrical properties that may overcome the physical limits silicon faces as transistors shrink to ever-smaller sizes.

Silicon's remarkable run as ruler of the chip world may be nearing an end as engineers eventually lose the ability to make faster silicon transistors by making them smaller. In the hunt for what comes next, carbon nanotubes have gotten a big chunk of the attention, but if the current explosion of research activity is any indication, it may be graphene

that wins in the end. This spring saw a flurry of breakthroughs surrounding graphene, culminating in the creation of what may be the smallest transistor ever made—one atom thick by 10 to 50 atoms wide.

Like carbon nanotubes, graphene is a crystal structure of carbon atoms but arranged in a flat plane instead of a cylinder. The electrons in graphene behave as if they have no mass. Like photons—but unlike electrons in other materials—the electrons move at a constant speed, regardless of how much energy each one has.

A transistor built out of graphene, therefore, should operate much faster than a comparable one made from silicon. Michael S. Fuhrer, a physicist at the University of Maryland's Center for Nanophysics and Advanced Materials, recently showed that at room temperature

PUSHING PENCILS: Graphene, found in pencil marks, is a candidate material for making future transistors. It's extracted from graphite crystals [below] using sticky tape.

TOP: ILLUSTRATION: ANNA DEMIAN/RANDI SILBERMAN; LEFT: GRAPHENE INDUSTRIES

update

electrons in graphene move at 200 000 centimeters per second for every volt per centimeter of electric field, 100 times faster than in silicon. "All other things being equal, that would translate into a 100 times faster transistor," he says.

Graphene has been known for decades as a single plane of graphite, but it was only in 2004 that Andre Geim and Kostya Novoselov of the University of Manchester, England, were able to isolate it by the simple act of pressing a piece of tape to a graphite crystal and placing it on a silicon substrate. In April, the two researchers described their transistor, 10 to 50 atoms wide and built by etching a pattern into graphene.

The substance is not a natural choice for a transistor material, in that it has no electrical bandgap, so applying an external electric field doesn't block the transport of electrons. In other words, a graphene transistor would be hard to turn off. Geim and Novoselov overcame that obstacle by etching away some of the graphene to create narrow constrictions, which, in the odd physics of two-dimensional materials, produced an artificial bandgap.

Funding for graphene research is flowing, notably from the EU's Graphene-based Nanoelectronic Devices (GRAND) program, which among other things is looking at whether graphene would still work its wonders when integrated with the silicon CMOS process. In the United States, the Defense Advanced Projects Research Agency is funding graphene research in the quest for better RF circuits. Industrial heavyweights such as IBM are also exploring the material. The company constructed a graphene transistor last fall.

Graphene may find a home outside ICs as well. Its high electrical conductance could lead to more sensitive chemical sensors, and it could prove to be a cheaper, more flexible substitute for the indium tin oxide used as transparent electrodes in LCDs and touch screens. It may also allow batteries to pack in more energy.

Despite widespread excitement over graphene's potential, Maryland's Michael Fuhrer warns that the research is still in its early days. "Right now we're at the stage where we've got this interesting material to play with," he says, "but I think it's all blue sky."

Robert Westervelt, a Harvard University expert on how electrons behave in 2-D materials, doubts that graphene will entirely replace silicon computing devices, but he says the material could lead to specialty circuits that compute in new ways, perhaps using some characteristic of electrons other than charge. "There are probably more questions than answers now, but it's got a lot of people excited," he says. "All of the rules are really changed about the ways the electrons behave."

—NEIL SAVAGE

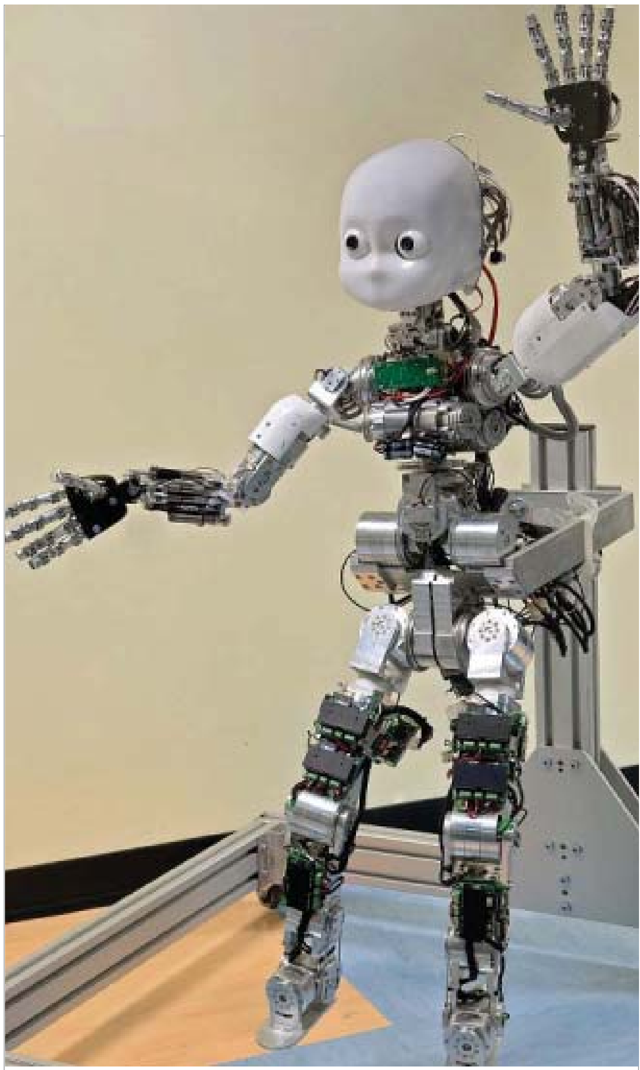

OPEN-SOURCE BABY

RESEARCH GROUPS ACROSS EUROPE are becoming parents to bouncing baby robots. By teaching them to walk, open doors, shake hands, and even talk, they hope to figure out how human children learn to do the things we adults take for granted. iCub—the size of a 3-year-old child—was developed by 11 separate research groups. So unlike competing robotic children, it had to be designed using modular hardware and open-source software. "It's more effective to have a critical mass of researchers tackling an area as complex as this," says Giorgio Metta, an assistant professor at Italy's University of Genoa. The iCub adoption program is sending robots to six European research labs this year and the next.

More online at http://www.spectrum.ieee.org/jun08/icub.

PHOTO: ROBOTCUB

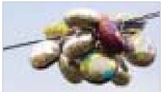

63 The number of blackouts caused by Mylar balloons in Burbank, Calif., since 1993. The city is considering a ban.

Putting Wireless Power To Work

Smart sensors harvest radio-frequency energy

When the U.S. Navy decided it needed to monitor the condition of its aging aircrafts' wings, it tried embedding wireless sensors inside them. Each sensor, attached to energy-harvesting circuitry, periodically checks the wings for damaging stress and strain, says Zoya Popovic, a professor of electrical engineering at the University of Colorado, Boulder, who worked on the project. To recharge or activate the sensors, a technician holds a transmitter a meter away from a wing to create a low-energy electric field within range of the sensors' energy-harvesting circuits.

In solving this problem, engineers dusted off a decades-old idea: radio-frequency energy harvesting, be it from strategically placed transmitters or from the ambient energy emitted by cellphone towers and television stations. The concept was once dismissed as unfeasible because of the rapid dissipation of electromagnetic waves as they travel from their source. But even microwatts, if trickled into a battery or supercapacitor, can be enough to power some sensors for more than a decade. The combination of extremely low-power microprocessors, increasingly affordable supercapacitors for energy storage, and budding markets for sensors that make buildings more energy efficient and monitor inventory has enabled a new generation of energy-harvesting devices.

Typically, wireless sensors are designed to observe environments in a more flexible way than wired ones can—tracking cattle in the middle of a field, generating early warnings of impending earthquakes, and assessing the structural health of bridges, for example. But a sensor's power supply is the most confounding problem. Each option has its limitations: a battery alone has a short lifetime, and solar cells, the most common energy-harvesting technique, can't soak up photons from inside an airplane's wing.

The technology for harvesting wireless power is essentially based on radio-frequency identification, or RFID. A transmitter sends a burst of radio-frequency energy that both carries information to a chip and can be converted to dc electricity to power it. A tag consisting of an antenna and a microchip responds by sending back data about the object it is attached to.

Turning those simple tags into fancier monitoring devices requires more power, so the RF energy would need to be captured and stored or transmitted continuously. Some food companies, for example, have begun tracking their delivery trucks more closely, according to one RFID technology company. A truck outfitted with a transmitter can both recharge and query RFID-based sensors that periodically check temperatures inside the truck or perform antitheft surveillance. When the sensors detect a change in their environments, the tags relay that data back to the transmitter.

A key development has been a steady growth in the distances over which the tags can communicate. Power companies hoping to install sensors along electric power lines, for example, would like to gather data from those devices without having to check each one individually. Norman D. McCollough Jr., a director of technology at Hendrix Wire & Cable Inc., a power-distribution systems company in Milford, N.H., anticipates that his remotely located energy-harvesting sensors, which now can transmit data across 100 meters, will soon be able to return information to a transmitter more than a kilometer away. McCollough has been using a chip introduced in January that can both pick up lower-power signals and operate at a higher power to send data back over longer distances, thanks to a redesigned power-amplifying circuit.

POWER TOWERS: Wireless sensors could scavenge power from radio signals.

LESTER LEFKOWITZ/GETTY IMAGES

With significantly more flexible sensor installations come new applications. For example, using energy-harvesting sensors to control thermostats in office buildings could improve the energy efficiency. To assess whether sensors in an office building could be powered ambiently, Marlin Mickle, an electrical engineering professor at the University of Pittsburgh, mapped the radiated power from the radio stations around his city. "When we talk about ambient energy, the source is not under our control. But it's still out there, and we're showing that we can make it work," Mickle says.

But even so, one essential problem remains: a reliable source of RF energy isn't always available. In addition to RF harvesting and attaching solar cells to microprocessors, engineers have explored converting temperature changes or mechanical movements—perhaps from vibrations or the flipping of a light switch—into electrical energy. Greg Durgin, an assistant professor of electromagnetics at the Georgia Institute of Technology, in Atlanta, thinks that ultimately the solution will be more complex.

"A really emerging area is getting hybrid power supplies to work on this," Durgin says. Linking two energy-harvesting techniques on one sensor might finally make the devices truly independent. —Sandra Upson

JANET BLAKE

"Now all the EE textbooks need to be changed" —LEON CHUA, who predicted the existence of the recently discovered memristor, the final fundamental circuit element. More at http://spectrum.ieee.org/may08/6207

update

Tata Hopes Its Supercomputer Is A Money Machine

Will Eka, the most powerful privately owned supercomputer in the world, turn a profit?

WHEN EKA, the 117.9-teraflop supercomputer built by the Computational Research Laboratories (CRL) in Pune, India, was named the fourth most powerful machine in the world last November, the global computing community—and even the computer's developers—were surprised. Though a new top 10 list, due out this month, may see Eka demoted, CRL has already proved that top rankings in this field, typically the domain of national laboratories in the richest countries, cannot be taken for granted. Next month, the company should see its larger goals achieved as well: turning the US $30 million Eka into a revenue-generating supercomputer for hire and its team into a supercomputer services consultancy. Eka is the only privately funded supercomputer in the top 10, and it is the only one built specifically to make money.

"We believed that high-performance computing can earn money, be a profitable business, and make a difference to the economy of the nation," says N. Seetha Rama Krishna, head of high-performance computing (HPC) operations at CRL, a wholly owned subsidiary of Tata Sons Limited, a division of the $28.8 billion-per-year conglomerate Tata Group, based in Mumbai.

True to India's software and services tech culture, rather than try to outdo Cray, IBM, Hewlett-Packard, or Silicon Graphics at designing and selling supercomputers, CRL will provide end-to-end supercomputing services—renting computer time, adapting and fine-tuning applications, and offering analytical services.

Today Eka is testing more than 15 applications for customers, and the company is in talks with several clients from the automobile, aerospace, financial, oil and gas exploration, and life sciences sectors, including aerospace giants Boeing, Embraer-Empresa Brasileira de Aeronautica, and Airbus.

"That commercial outfits are willing to bet huge money on HPC is surprising," says S. Ramakrishnan, director general of the Center for Development of Advanced Computing, another supercomputing center in Pune. "CRL could be a game changer in this arena."

To get the job done, CRL started with a conventional cluster model of supercomputer, a group of 1800 Hewlett-Packard servers made up of more than 14 000 processors. The innovation lies in how the servers are connected to one another. The interconnect scheme relies on the concepts of projective geometry—a non-Euclidean method in which parallel lines intersect at infinity. The result is a simpler interconnect layout that both increases the computer's bandwidth and lowers its power consumption and cost of construction.

As a result of the interconnect scheme, the supercomputer is basically circular. At the center are data storage and networking hardware. Encircling those is the computing hardware, and surrounding it is a system of powerful air conditioners. This shape contrasts sharply with that of other cluster-type supercomputers, which are arranged as rows upon rows of computers.

The Eka team is still optimizing the machine, designing its own network switches and building specialized mathematical software. CRL has also begun work on the architecture and the key pieces of Eka's successor. —SEEMA SINGH

Eka Supercomputer Layout

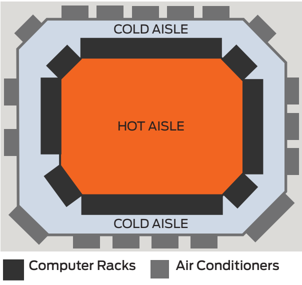

Typical Supercomputer Layout

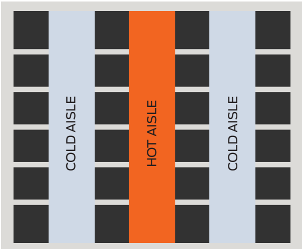

THE RING: Eka's interconnects give it an unusual layout.

FROM TOP: R. STANLEY WILLIAMS; TATA; ANNA DEMIAN

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

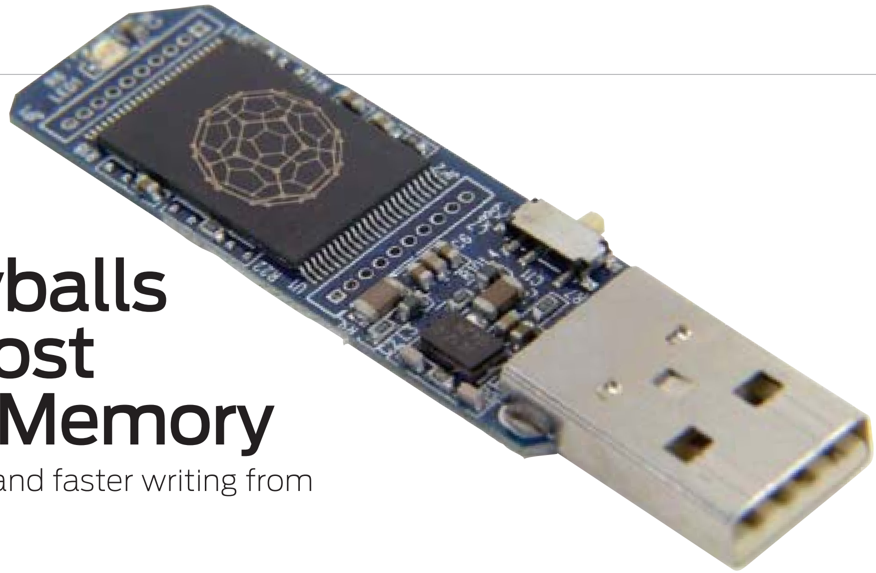

Buckyballs To Boost Flash Memory

Lower power and faster writing from a dash of C_{60}

WANT FASTER flash? Sprinkle in a little carbon. Researchers at Cornell University, in Ithaca, N.Y., suggest that a thin layer of buckminsterfullerene (C_{60}), a.k.a. buckyballs, embedded in an ordinary flash-memory cell, can increase how long the memory holds a bit, boost the speed at which a bit is written or erased, and decrease the memory's drain on a battery.

Like ordinary transistors, flash-memory cells are made up of a source, a drain, and a gate, whose voltage controls the flow of current between the two. The gate is separated from the rest of the transistor by a thin layer of insulation, the gate dielectric. The difference in flash is that the cell contains an additional, "floating" gate embedded within the dielectric. Put enough voltage on the main gate and electrons will jump the dielectric barrier and get stuck inside the floating gate. Reverse the polarity of the voltage and the charges will jump back out. The stuck charge, or its absence, is the stored bit.

Tuo-Hung Hou, of the Cornell group, likens the floating gate to a tiny island inside the chip. To change the bit's value, the water level surrounding the island—the barrier that the gate dielectric presents to an electron trying to jump from the transistor channel into the floating gate—must be reduced so that the floating gate can be easily accessed.

Having followed Moore's Law for years, flash memory is bumping up against the physical limits of speed, size, and power consumption. To make flash chips that are faster and consume less power during the writing process, engineers must find a way to get electrons into and out of the floating gate using less voltage. One way is to make the insulation thinner so that it takes less voltage to induce an electron to jump across the insulation and less time for electrons to fill the floating gate.

Today the insulator is just 7 to 8 nanometers thick. But, says Sanjay Banerjee, director of the University of Texas at Austin's Microelectronics Research Center, if it gets any thinner, using the same insulation technology, "that insulator becomes so thin that it inevitably has some defects, which causes the charge [holding the bit]...to leak away."

Hou and his colleagues' solution was to build a layer of buckyballs into the dielectric between the floating gate and the transistor channel. "The buckyballs are like putting stepping stones in the middle of the river," Hou says. "Then electrons can jump onto the stepping stones and then jump to the other side."

Without any applied voltage, the dielectric barrier is still too big for electrons to cross. But introducing just a few volts reduces the barrier enough to make the buckyball stepping stones accessible, allowing electrons to move into and out of the floating gate more quickly and easily. Researchers had earlier made the stepping stones using silicon nanocrystals, but their effectiveness depends on the nanocrystals' size, which is difficult to keep consistent. Buckyballs, which act as a semiconductor in this case, are always the same size.

The Cornell group says that the presence of C_{60} can cut the voltage needed to operate flash-memory chips from today's 10- to 13-volt range down to 5 V. Less voltage required for every bit being written or erased means substantially less drain on the battery, says Banerjee, who researches flash memory and consults for leading flash manufacturers. The improved writing speed could open up new kinds of applications that previously used faster memories, such as DRAM.

Banerjee thinks that adding buckyballs would add to the cost of manufacturing flash chips. "But if it works the way we hope it does, then it should improve the performance so much that I think it'd be worth the cost." —MARK ANDERSON

ORIGINAL PHOTO: IMAGEBROKER/ALAMY; PHOTO-ILLUSTRATION: ANNA DEMAN

news brief

NANO-PINE TREES

University of Wisconsin materials scientists grew these lead-sulfide nanowires [above] in the shape of pine trees, with a main trunk and branches spiraling around it like a helical staircase. Structures like this were difficult to make—until the researchers found that a single screw-shaped dislocation embedded in the trunk wire causes a spiral of branches to sprout.

PHOTO: AAAS/SCIENCE

update

Radiation Sensor Fine-Tunes Cancer Treatments

An implantable detector could monitor changes in tumors

KEEPING TABS: Babak Ziaie's implant monitors radiation doses and wirelessly delivers the data to doctors.

A NEW IMPLANTABLE device promises to offer doctors more precision in treating patients with cancerous tumors. Babak Ziaie, a professor of biomedical engineering at Purdue University, in West Lafayette, Indiana, has created a wireless radiation detector that measures from within a tumor itself how much therapeutic radiation the tumor is getting.

Ziaie designed the tiny dosimeter, which is about 2 centimeters long, to fit inside a hypodermic needle, making it easy to inject the device into the body. The detector is simple—a modified capacitor attached to an inductor, both of which are encased in a glass capsule. His breakthrough was to use micromachining techniques to shrink the device.

About half of all cancer patients in the developed world receive radiation therapy, according to the International Agency for Research on Cancer, in Lyon, France. In one common type of therapy, a cluster of individually sculpted beams of high-energy photons are shot at a prone patient. Before a series of treatments begins, a computer creates models of the patient's body to tailor the paths of the thin beams to match the shape of the cancerous target.

The challenge is knowing exactly where the tumor is, down to a fraction of a millimeter, on any given day. "When a patient comes in for radiation day in and day out, his prostate may be off by 5 or 10 millimeters from where it was the day before," says Allan Pantuck, a urologic oncologist at the University of California, Los Angeles. "So the question is, are you really hitting the target you think you're hitting?"

If an implanted dosimeter indicates that a tumor received less radiation than the treatment plan had specified, the hospital staff could revise the patient's therapy, Ziaie says. As a result, the tumor would receive a corrected dosage, and the surrounding tissue—which would presumably absorb the missing radiation—could be spared more damage.

The radiation-sensing capacitor consists of two plates separated by an air gap. The top plate is mounted on a flexible membrane that allows it to move. The bottom plate is fixed in place and consists of a layer of Teflon on top of glass. The Teflon plate forms an electret. Often referred to as the electric equivalents of magnets, electrets are materials that are able to maintain an electric field, sometimes for longer than a decade. The electret's electric field pulls the top capacitor plate toward it.

During a treatment, the radiation ionizes the air between the electrode plates, and the resulting charged particles weaken the electret's electric field. That in turn causes the top plate to move away from the bottom one.

To read the dosimeter, an external antenna delivers an oscillating electric field to the implanted device. The inductor in the device picks up this field and resonates at a frequency proportional to the distance between the plates in the capacitor. Detecting this resonance frequency gives a measure of how much radiation the device has absorbed.

One company, Sicel Technologies, in Morrisville, N.C., has also developed an implantable dosimeter, but it contains a microchip and other circuitry. Ziaie anticipates that his device will be significantly cheaper and easier to manufacture, on the order of US $5 each.

However, not all doctors are convinced that patients need an implanted dosimeter. Generally, doctors treat their patients according to the initial computer simulations rather than measuring the dose each patient is actually receiving. So they don't necessarily know that there may be a problem to fix. Niko Papanikolaou, the director of medical physics at the University of Texas Health Science Center, in San Antonio, tested Sicel's dosimeters in several of his patients. "All we discovered was that we were doing things right," he says.

But not everyone does things right. In April 2008, Ottawa Hospital Cancer Centre, in Canada, revealed that over the course of three years 326 cancer patients had received almost 20 percent less radiation than they were prescribed, due to an incorrectly programmed machine.

Whether or not doctors warm up to real-time radiation monitoring with implanted dosimeters remains to be seen. To improve its odds, Ziaie hopes to shrink his device even further, in time to start testing it in pigs by the end of 2008. —SANDRA UPSON

PURDUE UNIVERSITY

technically speaking BY PAUL MCFEDRIES

Homo Nerdus

"Nerds are the ones who don't go to the party so they can stay home and do homework; geeks bring their homework to the party."
—David Anderegg, *Nerds: Who They Are and Why We Need More of Them* (2007)

Is THE word *nerd* an insult or not? Until recently, there was no doubt; in fact, most dictionaries call *nerd* an offensive term, used to insult a person's appearance, hygiene, or social skills. That sense of the term has been around since at least the early 1950s. The 28 October 1951 issue of *Newsweek* tells us that "in Detroit, someone who once would be called a drip or a square is now, regrettably, a nerd." The word *nerd* also appears in the 1950 Dr. Seuss story *If I Ran the Zoo*, but he was referring to a fictional animal, not a socially inept person.

Now, however, most reference guides also include a second definition for *nerd* that's practically a compliment. For example, Encarta defines a nerd as a "single-minded enthusiast: somebody who is considered to be excessively interested in a subject or activity that is regarded as too technical or scientific." The phrases "excessively interested" and "too technical or scientific" still give the definition an odor of insult, but that bit about being a "single-minded enthusiast" doesn't sound bad at all. Wikipedia's definition is similarly ambiguous: "a person who passionately pursues intellectual activities, esoteric knowledge, or other obscure interests that are age inappropriate rather than engaging in more social or popular activities."

Some folks are taking the positive aspects of the word's definitions and running with them. That is, people are enthusiastically embracing their inner (and outer) nerd. For example, the online merchandiser Cafe Press has a Geek and Nerd Gifts section where you can buy T-shirts and other items with slogans like "Talk Nerdy to Me," "Nerd Girl," and "I [Heart] My Nerd." There's even a Nerd Pride Day (also called Geek Pride Day), which is celebrated on 25 May, the day the first *Star Wars* movie was released, in 1977.

All this **pro-nerd** feeling is spilling over into the language, too, with nerd-related coinages popping up like pocket protectors at a comic-book convention. For example, the population of nerds taken as a whole is called **nerdom**, and a person's **nerdy** traits and characteristics represent their **nerdity**. The latter term is used often by the psychologist David Anderegg in his engaging book *Nerds: Who They Are and Why We Need More of Them* [Tarcher, 2007]. The whole nerd-is-cool meme is often summarized in the formerly oxymoronic phrase **nerd chic**.

Any long and nerd-oriented activity is known as a **nerdathon**, and if that activity happens to be a computer game or a LAN party (a gathering where people bring their own computers, connect them together into a local area network, and then play games against one another), it's called a **nerdstorm**.

As yet another example of the digital DIY movement I talked about in my column last June, nerds are embracing crafts of various kinds. For example, some nerds are baking cakes in the shape of Sonic the Hedgehog or an Xbox 360 console. These are known as **gamecakes**, and the people who bake them are **gamecakers**. The desserts are examples of a larger genre called **nerdcraft**, and the people who engage in such activities are called **nerdcrafters**.

On the music front, there are artists who specialize in a form of rap music with lyrics relating to computers, technology, and engineering (I am not making this up), a genre known as **nerdcore** (from its original association with the hard-core music genre), though many people prefer the term **geeksta** (a play on *gangsta*).

Nerds are even starting to congregate in the same areas (outside of Silicon Valley, that is), a trend first recognized by the urban analyst Joel Kotkin. He uses the term **nerdistan** to refer to any upscale and largely self-contained suburb or town with a sizable population of high-tech workers employed in nearby office parks that are dominated by high-tech industries. Those employees with vested stock options in successful tech start-ups are known as **millionerds** or, if they started the company, **entreprenerds**.

All these nerdologisms can't hide the fact that, for nongeeks, the word *nerd* is still something of an insult (more so than the now almost neutral term *geek* but less so than the truly insulting terms *dork* and *dweeb*). The difference is that now the nerds simply shrug their shoulders, push up their glasses, and go back to whatever they were obsessing about. They're proud of their **nerdhood**, and they know that living **nerdily** is the best revenge. ❑

DAVID BRABYN/CORBIS

careers

DENEE BUSBY: TECHNICAL ACTRESS

She can sing, she can dance, she can act, she can redesign your oil refinery

TO EXPERIENCE the full fabulousness of Denee Busby, first you have to see her in action, then you have to meet her, and finally you have to listen to her story.

I recently saw her onstage, playing a scantily clad comedic floozy opposite Vivica A. Fox. When we met after the show, Busby was resplendent in a neon orange jumpsuit with rhinestone hoop earrings and wild, curly shoulder-length hair. And that was when she told me how she got to be both a working engineer and an actress.

Since 1999, when she graduated with double bachelor's degrees in electrical and mechanical engineering from the University of California, Los Angeles, Busby has assisted with the design of refineries, plants, and pipelines along the West Coast for Shell and Chevron. She has also appeared on the TV series "Eve" and "Charmed," had a role in the 2004 movie *The Cookout* with Ja Rule, and was in last year's sold-out touring musical *Whatever She Wants* with Fox. She says her cast mates would tap her for computer tech support on the road.

"People tell me, 'You don't act like an engineer, but you operate like one,'" Busby says with her husky voice. "I was attracted to electrical engineering because it was so structured. It was circuit design, and there was only one way it worked. But because my brain thinks in absolutes, the subjective nature of entertainment is hard for me to deal with.

BILL BALLERINI

"I always knew that the analytical side of my personality would be how I'd make my living. But I still danced and sang through college, as a release. It didn't happen until after I graduated college that the acting bug really got unbearable."

Busby grew up in northern California and began dancing at age 2. In high school she acted in school productions and spent a summer performing with the Alvin Ailey American Dance Theater. But an affinity for math coupled with two academically oriented parents pushed her to the safer path of engineering. She began at MIT, but says she found it too cold, too white, and too staid, so she transferred to UCLA. From there she went on to Shell and Chevron.

She says she did not blend in at the office.

"Imagine a boardroom of white men over 60 who have been in this business forever," she says, laughing. "I walk in with a miniskirt, 4-inch stilettos, and a very bright smile. They're, like, 'So, uh, when is the head engineer gonna get here so we can start the meeting?' I'm, like, 'No, no, Boo. That's me. Now turn to page one.' But after I started talking shop, they realized I knew what I was talking about."

Then came the epiphany she had in 2001 while watching the "Soul Train" dance show. She began crying, realizing how much she missed performing. The experience persuaded her to start dancing and to take acting classes, where she was spotted by a casting director for *The Cookout*. She soon landed an agent.

To get the flexibility she needed to schedule clients around her auditions, she started working as a contractor, primarily for the Encino, Calif., branch of Go Engineer, a consulting agency headquartered in Salt Lake City. Her boss doesn't mind that she has a second life as an entertainer, but it took her parents six months to reconcile themselves to her career moves. She has advised such big-name clients as Boeing, Raytheon, and Northrop Grumman, mostly on product design and the streamlining of communication between manufacturing and production divisions.

She says her background in engineering helps her as an actress, and vice versa. "After memorizing formulas and theories, I can memorize a script in a matter of minutes," she says. "And being an actress helps me adapt to each client's personality."

—SUSAN KARLIN

mini-profile
By Susan Karlin

GORDON CLAPP: THE VOICE IN YOUR COMPUTER

After starring turns on the television series "NYPD Blue" and Broadway's *Glengarry Glen Ross*, actor Gordon Clapp has found a way across the digital divide. He voices General Randall—a rough, buff, self-serving zealot—in *Prototype*, Sierra Entertainment's highly anticipated video game, slated for release in August. (See the trailer at http://www.prototypegame.com.) Randall tries to rein in Prototype, a young man granted superhuman powers who is attempting to recover his memory. "Randall's a bad guy, but he's got a full head of hair, so it's a tradeoff," Clapp says, laughing. "A whole new generation of fans will think this is what I look like. Even my son suddenly thinks I'm cool. Now I just have to figure out how to work one of these games."

PHOTO: DAVID M. ROSENTHAL

HOW TO TALK LIKE A SALESMAN

Short answer: by learning to think like one

OF ALL the courses I took while earning my three engineering degrees, the most valuable one was in marketing management. Every engineer should be acquainted with the subject.

To be sure, there is often chronic tension between salespeople and us engineers. They always seem to promise too much and to want it delivered too fast, and they never seem to care for our input. I imagine the salespeople have equally harsh things to say about engineers. The lack of understanding is the heart of the problem.

With that in mind, here are tips for engineers to integrate sales and marketing into their careers:

Build a relationship with the sales and marketing staff. Take the initiative to explain your product or service in layman's terms. Don't promise features that can't be delivered at a certain price. Give a realistic due date. Avoid the conflicts between engineers and marketing that so often delay software releases or send products off that have not been properly tested for bugs.

Get the terminology straight. For instance, "sales" is often used synonymously with "marketing," but there are big differences, best understood by placing things in the right order. The first step is market research, which determines what potential user groups need or want. Next is market planning, which estimates how many users a given product will attract at a given price. Then there's market development, which addresses strategic problems such as when to enter a market.

Last comes market promotion, the active selling, or "marketing." This involves the so-called 4 Ps—the attributes of the product, its price, the place where it can be bought, and the way it is promoted.

Segment the market. Define the potential users who should have the most influence on the design and use of your product, then "position" your offering for them. A classic success story was General Electric's estimate that airlines would eventually want quieter, more fuel-efficient engines. The company spent a decade developing such engines, then entered the airline business, shouldering aside such traditional stalwarts as Rolls-Royce.

Use focus groups to study target users. This is where a facilitator discusses product concepts or actual prototypes with a group of people. I've gained insights listening to potential customers' reactions at such gatherings.

Observe how customers behave with new products. I introduced prepaid phone cards into U.S. airports when they were unknown. One day an 11-year-old girl shadowed me at work. I handed her a phone card and told her to call her mother but gave her no instructions on how to use it; she read the instructions and got through on the second try using my phone. I later told our senior executives that phone cards were so easy "an 11-year-old could use them."

Read competitors' ads. Advertisements often give valuable clues to how a competitor wants to appeal to customers and which customers it is targeting.

Use a venture-capitalist model to develop products. Explain the product in plain English so that anyone can understand it. Identify real users, and show where the product has actually been implemented.

If we engineers can mesh our know-how with sales and marketing, we can better deliver products to markets that will actually use them. —CARL SELINGER

books

Just What Do You Think You're Doing, Dave?

How Apollo's astronauts learned to work with—and around—their computers

IN 1961, the average rocket-borne computer ran on average for 15 hours before an electronics failure crashed it. That dismal performance record didn't matter much to the military, whose suborbital missiles required only minutes of computer on-time. But a manned moon shot required that computers run 1500 hours between failures.

As David Mindell points out in *Digital Apollo: Human and Machine in Spaceflight*, NASA's project managers not only met that 1500-hour goal but greatly overshot it. When Neil Armstrong and his compatriots strode on the lunar surface between 1969 and 1972, the total mean time between failures of the onboard computers turned out to be closer to 50 000 hours.

The story is more than a historical footnote. Computer systems reliability was perhaps the biggest spin-off the space program ever had. The entire world has benefited from the achievement.

Yet getting there was half

ROBIN BARTHOLICK/GETTY IMAGES

the fun. The most compelling part of the book is its careful examination of the ways computers occasionally misbehaved. All six moon landings, for instance, took place with the mission commander manually overriding an overwhelmed computer—which carried just 4 kilobytes of RAM and clocked in at 1 megahertz.

Digital Apollo is certain to interest readers of a technical bent, especially those curious to delve into the gritty details of some of the first portable computers. Mindell joyfully plumbs the deep history of Apollo's decade-long clash between the MIT eggheads who built the computers and the thrill-jockey military test pilots who used them. Even I, who last programmed a computer in the early 1990s, found myself swept up in the account of the Apollo 11 mission. The entire world watched as the *Eagle*'s computer blurted out strings of alarms to its pilots, who simply ignored the warnings and landed the thing anyway. Neil Armstrong's cool control, leaving just 40 seconds of fuel in the tank, deserves shared credit with all the hardware and software Mindell cites as having saved *Apollo 11* from disaster.

Unfortunately, the book devotes just one paragraph to the most extraordinary test the project's computers ever faced: the Apollo 13 mission, which after a midflight oxygen tank rupture became a storied quest to rescue three men from spacebound oblivion. By the time Mindell gets that far in his narrative, he seems to want to concentrate only on moon landings—even though he leads off his book with a poetic Antoine de Saint-Exupéry quote about man versus machine. Mindell ought to have known that inflexible schedules and programs shouldn't always be allowed to carry the day.

—Mark Anderson

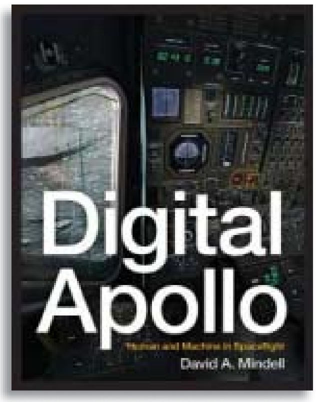

DIGITAL APOLLO
By David A. Mindell; MIT Press, May 2008; 456 pp.; US $29.95; ISBN: 13:978-0-262-13497-2

If It's Tuesday, This Must Be Los Alamos

A husband-and-wife team of tourists sees the sights of Cold War weapons projects

Sharon Weinberger and her husband, Nathan Hodge, spent two years touring the world's nuclear weapons sites. What they found calls to mind William Faulkner's famous remark "The past is never dead. It's not even past."

Though the public debate over nuclear weapons is alive, its terms are quaint, having hardly changed since the end of the Cold War and the subsequent illusion of being on a "holiday from history." Neither those who favor the continued existence of a nuclear deterrent force nor those who oppose it seem to have bothered to update their arguments.

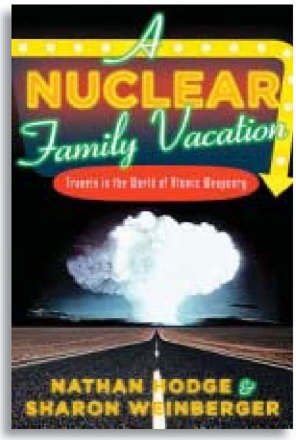

A NUCLEAR FAMILY VACATION: TRAVELS IN THE WORLD OF ATOMIC WEAPONRY
By Sharon Weinberger & Nathan Hodge; Bloomsbury USA, 2008; 336 pp.; US $24.99; ISBN 978-1-596-91378-3

Those who argue for the weapons tend to base their reasoning on the Cold War policy of "mutually assured destruction," an incomplete argument in a time of suicide bombers. Those who argue against the weapons generally demand unilateral U.S. disarmament, which ignores the complexities posed by North Korea and Iran. But the real contrast is between those who care at all and the vast disengaged majority, which treats nuclear weapons and their regulation as a problem for the history books.

Meanwhile, the Pentagon is making policy decisions to address its aging nuclear stockpile and deteriorating facilities—decisions that barely get mentioned in the newspapers. "What happens when a war ends," the authors ask, "but the warriors don't go home?"

The couple's defense-writing credentials are as impeccable as they come. Weinberger is editor in chief of *Defense Technology International*, and Hodge writes for industry gold standard *Jane's Defence Weekly*. Their book contains as much history as a college text, but years of magazine writing have given Weinberger and Hodge the spoonful of sugar they need to make the medicine go down.

That sugar is delivered in the form of jaunty, you-are-there travel writing big on local color from such undeniably colorful places as Sandia National Laboratories, in Albuquerque, the Kwajalein atoll, in the Marshall Islands, and Semipalatinsk Test Site, in Kazakhstan. The tone is neither pro- nor antinuclear, but the authors can't resist making the occasional bone-dry aside in the footnotes. It comes rather as a shock to learn that the U.S. Department of Energy got around to canceling the annual deer hunt at Tennessee's Y-12 National Security Complex only after 9/11. Before then, it seems, nobody had even worried about letting armed, uncleared individuals roam a nuclear installation.

The book, which goes on sale this month, succeeds in its attempt to resurrect the nuclear dialogue largely because it refrains from drowning the reader in facts. The fancy wrapping around the complex package makes this a good beach book for engineers.

—Sally Adee

invention

THE PITFALLS OF PATENT SEARCHES

Knowing about patents can sometimes be worse than not knowing

STEPHEN SAVAGE

YOU CAN learn a lot by searching patents, but what you learn can sometimes be dangerous. If a court should ever find that you infringed on a patent *knowingly*, you might have to pay triple the damages, together with attorney fees.

You can protect yourself by obtaining a lawyer's written opinion stating either that there is no infringement or that the patent in question is invalid. However, such an opinion can cost tens of thousands of dollars. So if you search patents regularly, getting an opinion for each patent you know about is just not practical. And what does it really mean to *know* about a given patent?

At one extreme on the knowledge continuum is the large, evil corporation that has studied a smaller competitor's patent and then ignored it, deciding, say, to outspend the competitor in court if sued for patent infringement. At the other extreme is the company accused of knowing about a patent simply because one of its engineers has stored a copy of it in his files, unbeknownst to bosses and colleagues. In between these two extremes lie the hard cases.

It's perfectly legal to study existing patents in order to know how to design around one of them. But what if a jury decides that you haven't really side-stepped a patent? You can use that design-around attempt as evidence that your infringement was not willful. Unfortunately, an unsuccessful design-around attempt does not automatically produce a finding of no willful intent.

It used to be that you had a duty to obtain an opinion from a patent attorney once you knew about a given patent. Then in 2004, the United States Court of Appeals for the Federal Circuit in Washington, D.C., ruled that an opinion of counsel was no longer required. Better yet, the court recently handed down a decision that makes "objective recklessness" the new standard for willful infringement instead of mere negligence.

Here are a few tips:

- Focus your search. If you are interested in analog-to-digital (A/D) converters, don't print out and file away the thousands of patents with "A/D converter" in the title or in the abstract.
- Become patent savvy. Learn how to read patents, how to tell if one has expired, and how to zero in on the patent's claims. Understand, too, that even expired patents you know about must be cited to the patent office when you apply for your own patent.
- If you know about a patent, tell your patent attorney what you've learned and its source.
- Follow your company's policies. That way, if the company gets in trouble, at least it shouldn't be counted against you.
- Finally, weigh in with your representatives in Congress. Many bills have been introduced to limit a finding of willful patent infringement to cases in which a company receives notice of the patent from the patent owner. So far, none of the bills has made it into law.

—KIRK TESKA

FATHER'S DAY gifts are no-brainers if Dad's an engineer. What? You say he's already got an oscilloscope—*and* a USB-powered hand warmer? Well, then get him a silk tie that *really* fits him, from ThinkGeek's online store (US $39.99). Our favorite, on the bottom right, depicts formulas.

PHOTO: THINKGEEK

tools&toys

Securing Your Laptop

The plug-in Yoggie Gatekeeper Pico promises to supersede standard antivirus and antispam software, but it's not for everyone

I'm a paranoid computer user. The first thing I do with a PC is install a full suite of "anti-" software programs—antivirus, antispy, antispam, you name it. I even leave Microsoft Vista's "Annoy me constantly" mode turned on. So when I got a browser virus anyway, I lost my faith in my software security shields.

Just in time, along came the Yoggie Gatekeeper Pico. It's a USB stick that bills itself as a replacement for all the security software we ordinarily run under Windows, designed with laptops in mind. All network traffic, wired or wireless, goes through the Pico before any Windows software sees it. And because the Pico is itself a complete computer, running Linux on an Intel XScale processor, it promises to bump up performance by supplanting the security software that now sucks cycles from your laptop's central processing unit.

That claim got my attention, because I both design and play games, and what gamer doesn't crave better performance? The price seemed right, too—I found it for US $149 up front and $30 a year for automatic updates (including, for example, new virus profiles), with the first year's updates free.

Installation is supposed to be straightforward: just insert the Pico and install a driver from a CD. However, it didn't work that way for me. The CD's installer wouldn't run on its own, so I ran it from Windows Explorer. Then, when the program launched Internet Explorer to register my Pico

A SECURITY stick may ward off malware well enough for the casual business user, but the serious user will still need to run standard cycle-sapping antivirus software. *PHOTOS: YOGGIE*

on Yoggie's Web site, the browser reported an invalid site certificate—not a good sign for a security product.

Registration also failed, at first, but eventually a new Systray icon showed up, indicating that my Pico was now providing my security. However, when I disabled Vista Security Center's firewall and other functions, Vista didn't seem to recognize the Pico as providing those services and complained that my computer was not secure against threats.

The Pico performed well on some basic matters. The logs showed that it had blocked various attempts to scan my machine. Indeed, moving to a public network (at Starbucks), I saw a significant increase in scans, none of which made it past the Pico. However, I ran into problems getting the firewall to work. Each time I disabled the Pico's driver, I lost my connection to *Lord of the Rings Online*, the game I develop. I had to fiddle with the custom configuration options, but even after opening the necessary User Datagram Protocol and Transmission Control Protocol ports, the game would still disconnect after about a minute. Yoggie's tech support pros suggested (two days after I e-mailed them) that I turn off the Intrusion Detection System/Intrusion Prevention System. That, too, failed to solve my problem—which was just as well. I am uncomfortable with having to turn off so many systems to make something work.

I have to say the Pico isn't for gamers like me, not only because of the firewall problem but also because the promised performance edge isn't real. You don't want to entrust all your security to a device with so many rough edges; you'll still want to run the standard, cycle-sapping security software as well. However, if you're a business user who needs security when traveling across random networks in hotel rooms and coffee shops, the Pico could be just the thing for you.

Off-loading security to a separate, tiny computer on a USB drive makes sense. We can expect to see more products of this type, and to see this one improve. I, for one, will wait for the market to mature. —Harry Teasley

"The brain is complex enough to conjure fantasies of technotranscendence and also to foil their fulfillment"

—John Horgan

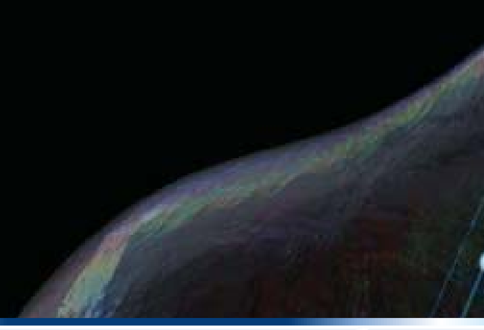

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

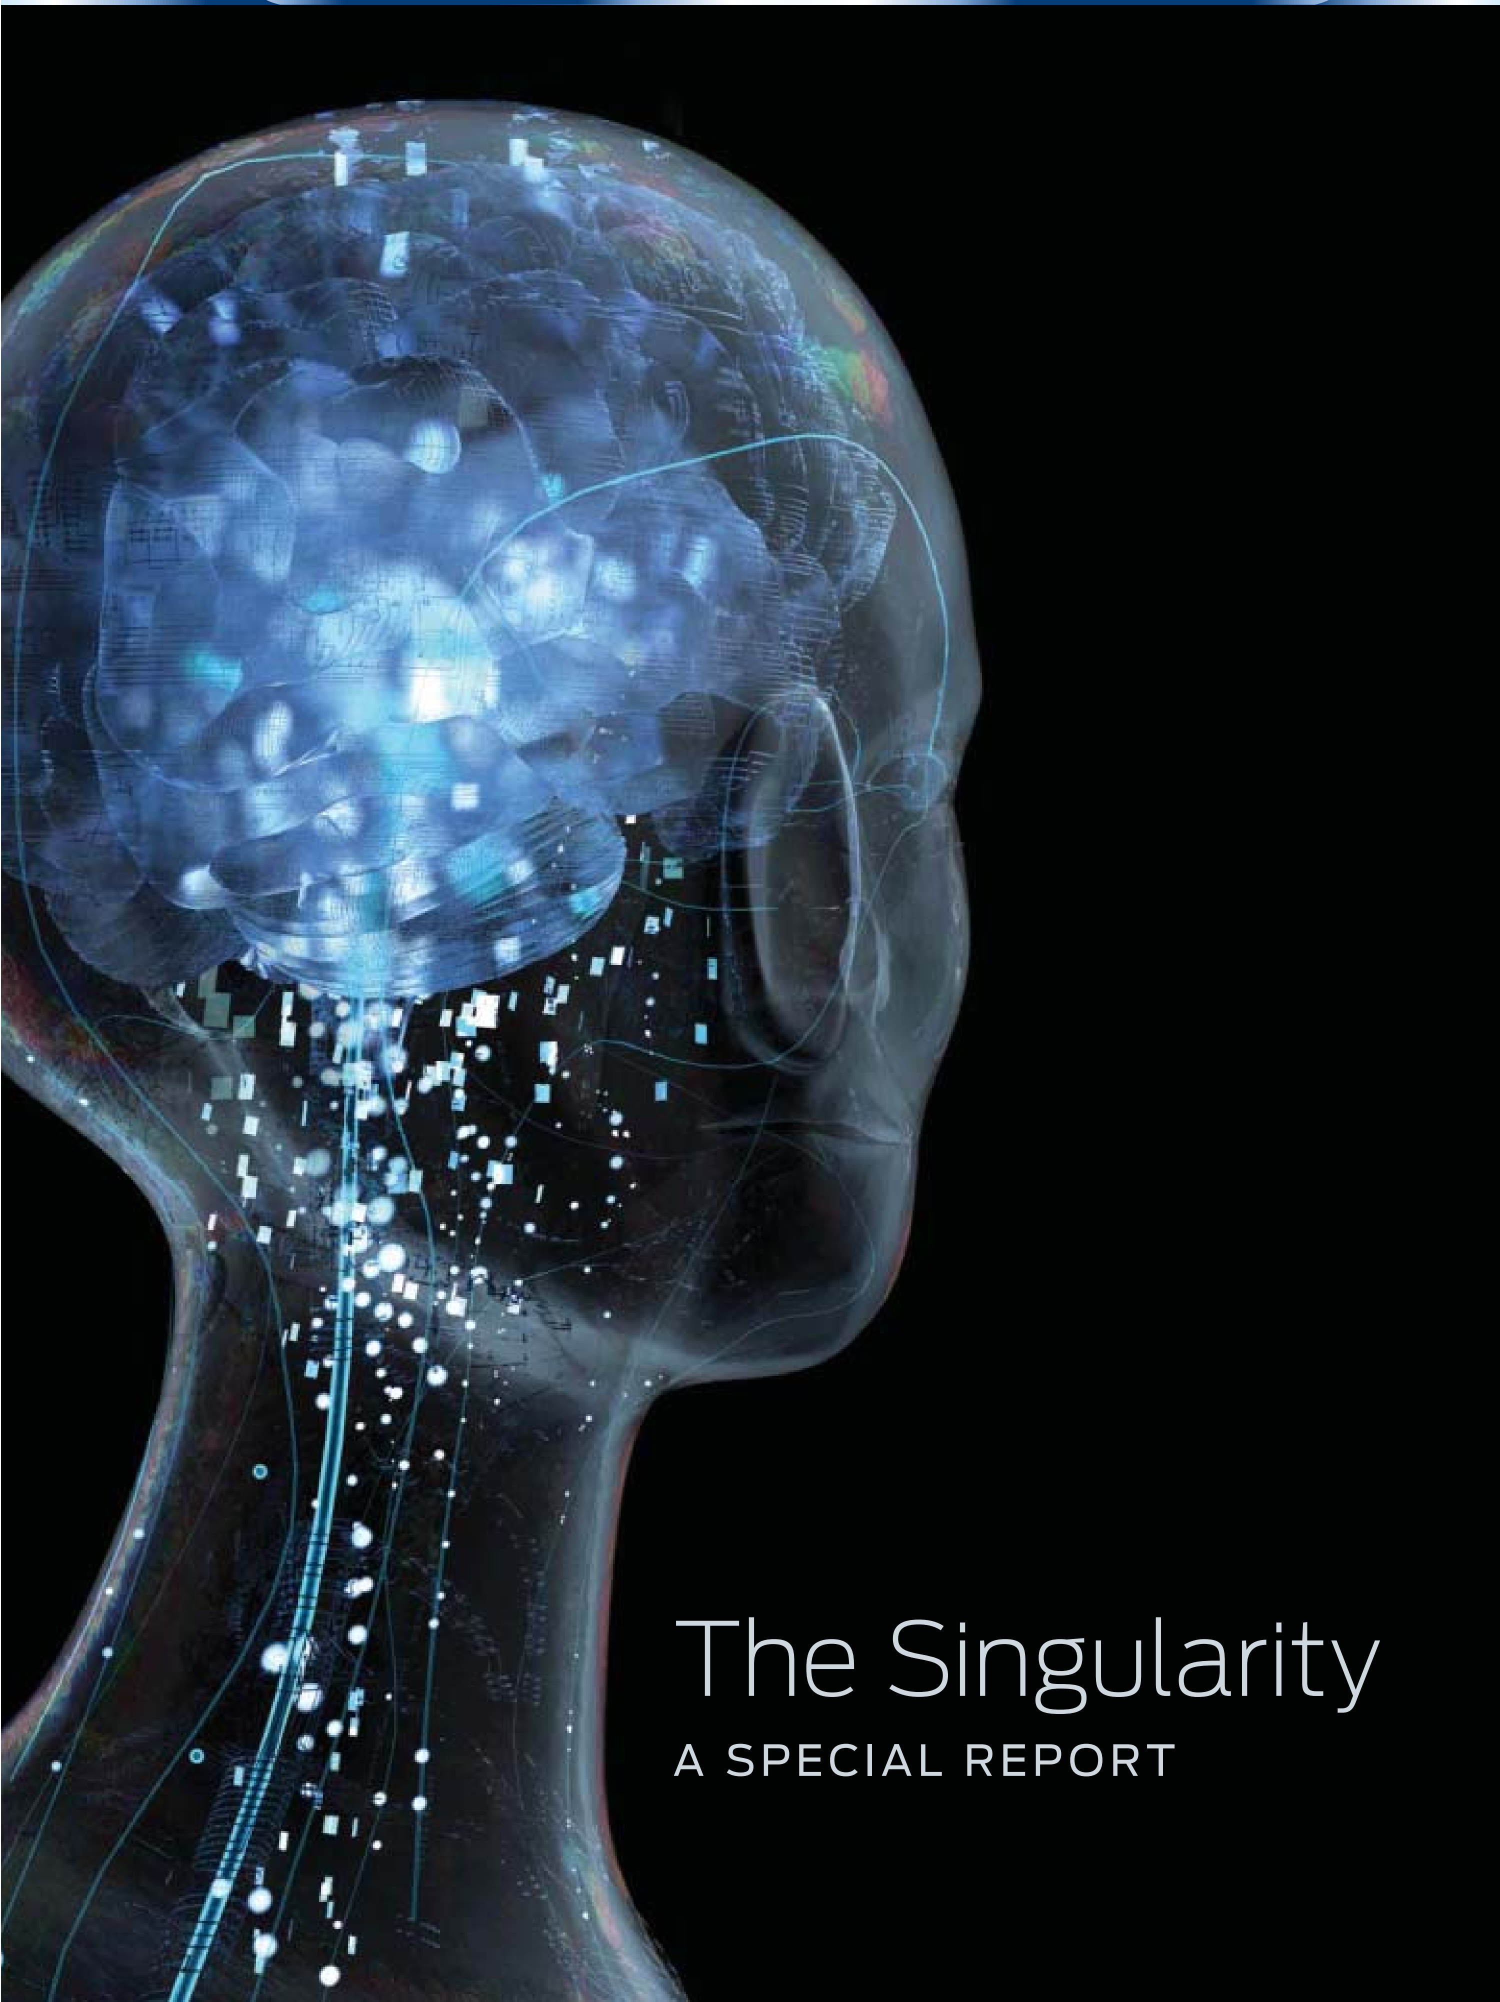

The Singularity

A SPECIAL REPORT

THE SINGULARITY | SPECIAL REPORT

WAITING FOR

TECHNOLOGICAL CONVERGENCE WILL CHANGE OUR LIVES BUT WON'T MAKE THEM INDEFINITELY LONG

BY GLENN ZORPETTE

ACROSS CULTURES, CLASSES, AND AEONS, PEOPLE HAVE YEARNED TO TRANSCEND DEATH. ¶ Bear that history in mind as you consider the creed of the singularitarians. Many of them fervently believe that in the next several decades we'll have computers into which you'll be able to upload your consciousness—the mysterious thing that makes you *you*. Then, with your consciousness able to go from mechanical body to mechanical body, or virtual paradise to virtual paradise, you'll never need to face death, illness, bad food, or poor cellphone reception. ¶ Now you know why the singularity has also been called the rapture of the geeks.

The singularity is supposed to begin shortly after engineers build the first computer with greater-than-human intelligence. That achievement will trigger a series of cycles in which superintelligent machines beget even smarter machine progeny, going from generation to generation in weeks or days rather than decades or years. The availability of all that cheap, mass-produced brilliance will spark explosive economic growth, an unending, hypersonic, techno-industrial rampage that by comparison will make the Industrial Revolution look like a bingo game.

At that point, we will have been sucked well beyond the event horizon of the singularity. It might be nice there, on the other side—by definition, you can't know for sure. Sci-fi writers, though, have served up lots of scenarios in which humankind becomes the prey, rather than the privileged beneficiaries, of synthetic savants.

But the singularity is much more than a sci-fi subgenre. A lot of smart people buy into it in one form or another—there are versions that dispense with the life-everlasting stuff. There are academic gatherings and an annual conference at Stanford. There are best-selling books, audiotapes, and videos. Scheduled for release this summer is a motion picture, *The Singularity Is Near*, starring the actress Pauley Perrette and a gaggle of aging boffins who've never acted in a movie. (Without any apparent irony, the picture's producers call it "a true story about the future.")

There's also a drumbeat of respectful and essentially credulous articles in the science press. Unlike stories about UFOs or zero-pollution energy sources, singularity stories don't exact from editors a steep payment in self-respect. That's because of the impressive attainments—albeit usually in fields unrelated to neuroscience or biology—of some of the people who chirp about mind uploading and nanomachine organ repair. The leading spokesman for the life-everlasting version of the singularity is the entrepreneur and inventor Ray Kurzweil, who's also behind the movie *The Singularity Is Near* and a recent book of the same title.

Why should a mere journalist question Kurzweil's conclusion that some of us alive today will live indefinitely? Because we all know it's wrong. We can sense it in the gaping, take-my-word-for-it extrapolations and the specious reasoning of those who subscribe to this form of the singularity argument. Then, too, there's the flawed grasp of neuroscience, human physiology, and philosophy. Most of all, we note the willingness of these

PREVIOUS SPREAD: BRYAN CHRISTIE DESIGN

people to predict fabulous technological advances in a period so conveniently short it offers themselves hope of life everlasting.

This has all gone on too long. The emperor isn't wearing anything, for heaven's sake.

The singularity debate is too rarely a real argument. There's too much fixation on death avoidance. That's a shame, because in the coming years, as computers become stupendously powerful—really and truly ridiculously powerful—and as electronics and other technologies begin to enhance and fuse with biology, life really is going to get more interesting.

So to produce this issue we invited articles from half a dozen people who impressed us with their achievements and writings on subjects central to the singularity idea in all its loopy glory, encompassing not just hardware and wetware but also economics, consciousness, robotics, nanotechnology, and philosophy. And with a few exceptions, we found people who are not on record as either embracing singularity dogma or rejecting it.

On consciousness, we have John Horgan, whose book *The Undiscovered Mind* describes how the mind resists explanation. We also have Christof Koch and Giulio Tononi, neuroscientists who specialize in consciousness. Rodney Brooks, of MIT's Computer Science and Artificial Intelligence Laboratory, weighs in on the future of machine intelligence. *IEEE Spectrum* journalism intern Sally Adee reports on a wildly ambitious effort, just gathering steam now, to map the human brain in enough detail to learn its secrets—and eventually re-create it. Robin Hanson, an economist, describes a future in which capitalist imperatives and technological capabilities drive each other toward a society that the word *weird* doesn't even begin to describe. Nanotechnology researcher Richard Jones, philosopher Alfred Nordmann, and semiconductor researcher Bill Arnold all consider aspects of singularitarian visions and explain where they're myopic.

For the last word in this issue, we turned to the computer scientist and science-fiction writer Vernor Vinge. It was Vinge's 1993 essay "The Coming Technological Singularity" that launched the modern singularity movement.

That movement has evolved since then into an array of competing hypotheses and scenarios [for a rundown, see "Who's Who in the Singularity," in this issue]. But central to them all is the paradoxical yet weirdly compelling idea of a conscious machine. Arguably, no other technology-related concept resonates with such intellectual and philosophical force.

Consciousness seems mystical and inextricably linked to organisms. What happens in the cerebral cortex that turns objective information into subjective experience—that turns chemical and neuronal activity in the mouth and nose into the taste of watermelon? pressure waves into the sound of an oboe? We don't know, but we will someday. No one argues that consciousness arises from anything but biological processes in the brain.

The brain is nothing more, and nothing less, than a very powerful and *very* odd computer. Evolution has honed it over millions of years to do a fantastic job at certain things, such as pattern recognition and fine control of muscles. The brain is deterministic, meaning that its reactions and responses, including the sensations and behavior of its "owner," are determined completely by how it is stimulated and by its own internal biophysics and biochemistry. Given those

THE RAPTURE

facts, most mathematical philosophers conclude that all the brain's functions, including consciousness, can be re-created in a machine. It's a matter of time.

Ah, but let's face it—time is what really matters. If you're obsessed with your own mortality, the idea of a computer blinking into consciousness 400 years from now isn't going to rock your world. You want the magic moment to come, say, 25 years from now at most. Unfortunately, that timetable grossly overestimates the speed of technical progress. And it underestimates the brain's awesome intricacy, as Horgan argues in his article. He, Koch, Tononi, and Adee all agree that everything we know about the central issue of brain research—how it creates consciousness, and therefore the universe each one of us inhabits—adds up to almost nothing.

What we do know is that the brain's complexity dwarfs anything we've managed to fully understand, let alone build. Koch, Tononi, and Brooks are all confident that consciousness will arise in a machine, but they are less sanguine about death-defying uploading, and especially about it happening in time to allow people alive now to preserve their minds in some sort of digitally created Eden.

Still, if you encounter my uploaded consciousness in a virtual paradise 50 years from now, feel free to tell me, "I told you so."

I won't mind a bit. ❑

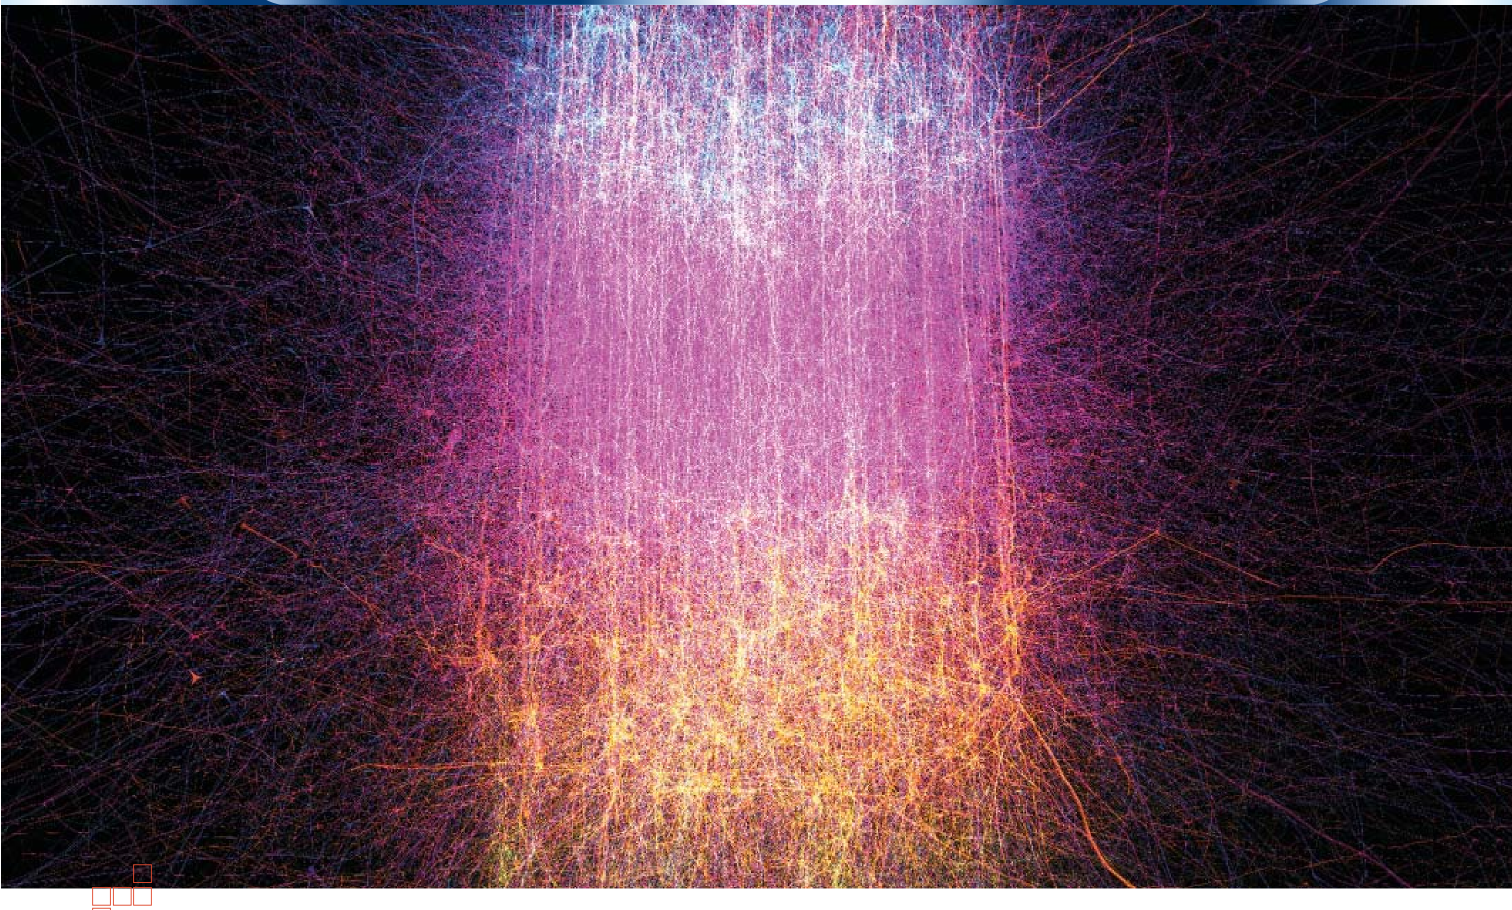

THE SINGULARITY | SPECIAL REPORT

THE WETWARE THAT GIVES RISE TO CONSCIOUSNESS IS FAR TOO COMPLEX TO BE REPLICATED IN A COMPUTER ANYTIME SOON BY JOHN HORGAN

The Consciousness Conundrum

I'M 54, WITH all that entails. Gray hair, trick knee, trickier memory. I still play a mean game of hockey, and my love life requires no pharmaceutical enhancement. But entropy looms ever larger. ¶ Suffice it to say, I would love to believe that we are rapidly approaching "the singularity." Like paradise, technological singularity comes in many versions, but most involve bionic brain boosting. At first, we'll become cyborgs, as stupendously powerful brain chips soup up our perception, memory, and intelligence and maybe even eliminate the need for annoying TV remotes. Eventually, we will abandon our flesh-and-blood selves entirely and upload our digitized psyches into computers. We will then dwell happily forever in cyberspace where, to paraphrase Woody Allen, we'll never need to look for a parking space. Sounds good to me!

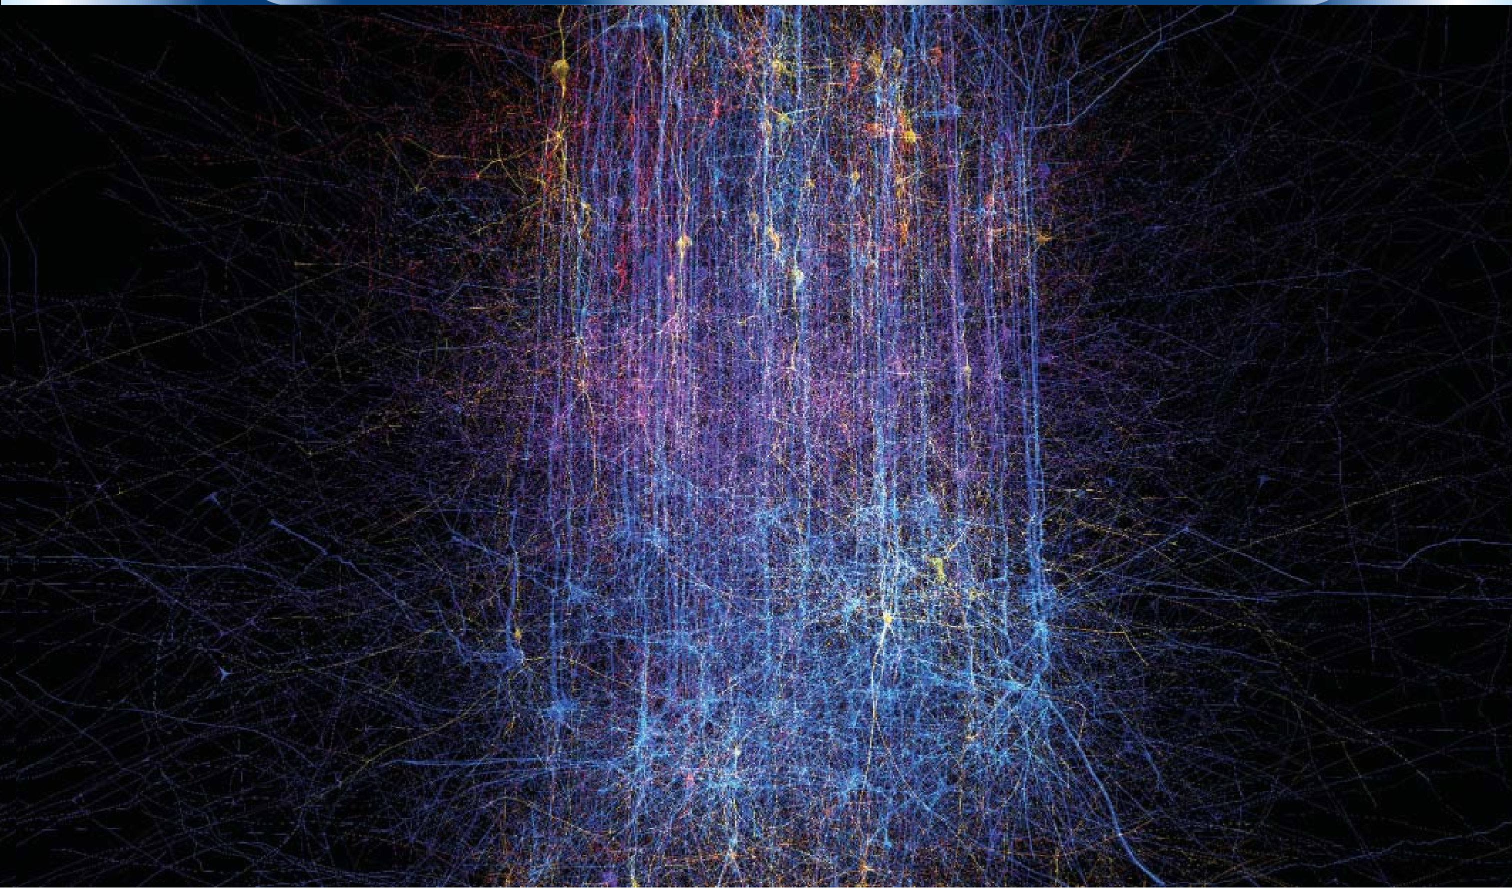

Notably, singularity enthusiasts tend to be computer specialists, such as the author and retired computer scientist Vernor Vinge, the roboticist Hans Moravec, and the entrepreneur Ray Kurzweil. Intoxicated by the explosive progress of information technologies captured by Moore's Law, such singularitarians foresee a "merger of biological and nonbiological intelligence," as Kurzweil puts it, that will culminate in "immortal software-based humans." It will happen not within a millennium, or a century, but no later than 2030, according to Vinge. These guys—and, yes, they're all men—are serious. Kurzweil says he has adopted an antiaging regimen so that he'll "live long enough to live forever."

Specialists in real rather than artificial brains find such bionic convergence scenarios naive, often laughably so. Gerald Edelman, a Nobel laureate and director of the Neurosciences Institute, in San Diego, says singularitarians vastly underestimate the brain's complexity. Not only is each brain unique, but each also constantly changes in response to new experiences. Stimulate a brain with exactly the same input, Edelman notes, and you'll never see the same signal set twice in response.

"This is a wonderful project—that we're going to have a spiritual bar mitzvah in some galaxy," Edelman says of the singularity. "But it's a very unlikely idea."

NEUROSCIENCE IS INDEED thriving. Membership in the Society for Neuroscience has surged from 500, when it was founded in Washington, D.C., in 1970, to almost 40 000 today. New brain journals seem to spring up daily, crammed with data from ever-more-powerful brain probes such as magnetic-resonance imaging and transcranial magnetic stimulation. In addition to such noninvasive methods, scientists can stick electrodes in brains to monitor and stimulate individual neurons. Researchers are also devising electrode-based "neural prostheses" to help people with nervous-system disorders such as deafness, blindness, paralysis, and memory loss.

In spite of all those advances, neuroscientists still do not understand at all how a brain (the squishy agglomeration of tissue and neurons) makes a conscious mind (the intangible entity that enables you to fall in love, find irony in a novel, and appreciate the elegance of a circuit design). "No one has the foggiest notion," says the neuroscientist Eric Kandel of Columbia University Medical Center, in New York City. "At the moment all you can get are informed, intelligent opinions." Neuroscientists lack an overarching, unifying theory to make sense of their sprawling and disjointed findings, such as Kandel's Nobel Prize–winning discovery of the chemical and genetic processes that underpin memory formation in sea slugs.

The brain, it seems, is complex enough to conjure fantasies of technotranscendence and also to foil their fulfillment.

A healthy adult brain contains about 100 billion nerve cells, or neurons. A single neuron can be linked via axons (output wires) and dendrites (input wires) across synapses (gaps between axons and dendrites) to as many as 100 000 other neurons. Crank the numbers and you find that a typical human brain has quadrillions of connections among its neurons. A quadrillion is a one followed by 15 zeroes; a stack of a quadrillion

NEOCORTICAL COLUMN of a rat, 2 millimeters high, as simulated on IBM's Blue Gene supercomputer. The image above shows about 50 kinds of neurons; the image on the opposite page, around 100. A rat's neocortex has 10 000 such columns; a human's has millions.

U.S. pennies would go from the sun out past the orbit of Jupiter.

Adding to the complexity, synaptic connections constantly form, strengthen, weaken, and dissolve. Old neurons die and—evidence now indicates, overturning decades of dogma—new ones are born.

Far from being stamped from a common mold, neurons display an astounding variety of forms and functions. Researchers have discovered scores of distinct types just in the optical system. Neurotransmitters, which carry signals across the synapse between two neurons, also come in many different varieties. In addition to neurotransmitters, neural-growth factors, hormones, and other chemicals ebb and flow through the brain, modulating cognition in ways both profound and subtle.

Indeed, the more you learn about brains, the more you may wonder how the damn things work. And in fact, sometimes they don't. They succumb to schizophrenia, bipolar disorder, depression, Alzheimer's disease, and many other disorders that resist explanation and treatment.

Nevertheless, the brain is a computer, singularitarians insist. It just has an extremely messy wiring diagram. According to this perspective, neurons resemble transistors, absorbing, processing, and reemitting the electrochemical pulses known as action potentials. With an amplitude of one-tenth of a volt and a duration of one millisecond, action potentials are remarkably uniform, and they do not dissipate even when zipping down axons a meter long (yes, a full meter). Also called spikes, to reflect their appearance on oscilloscopes, action potentials supposedly serve as the brain's basic units of information.

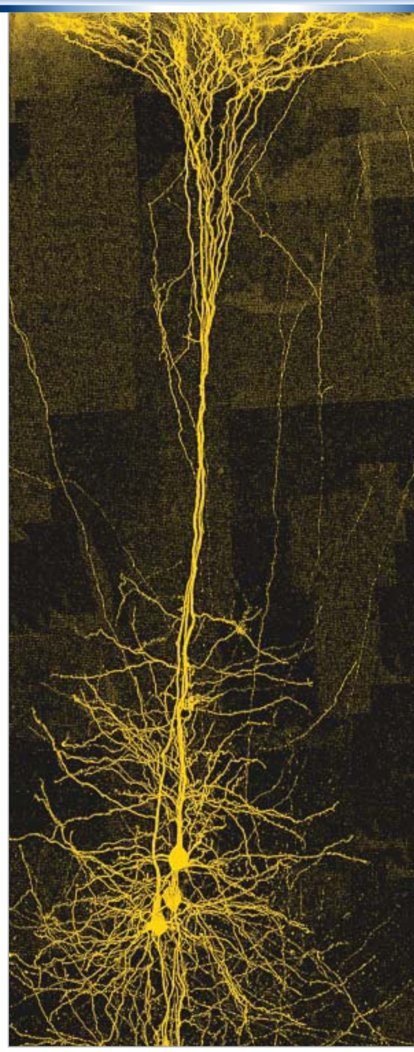

WHAT IT ALL COMES DOWN TO: In this image, three pyramidal neurons extend a couple of millimeters from their bulbous, dendrite-shrouded nuclei upward through their axons and axon endings at the very top.
PHOTO: IBM/EPFL

WITHIN A DECADE or so, computers will surpass the computational power of brains, many singularitarians say. They base this claim on the assumption that those spikes represent the brain's total computational capacity. If the brain contains one quadrillion synapses processing on average 10 action potentials per second, then the brain performs 10 quadrillion operations per second. At some point in the near future, some singularitarians say, computers will surpass that processing rate and leave us in their cognitive dust unless we embrace them through bionic convergence or uploading.

We've heard such prophesies before. A half century ago, artificial-intelligence pioneers such as Marvin Minsky of MIT and Herbert Simon of Carnegie Mellon University predicted that computers would exceed human intelligence within a generation. Their prophesies inspired sci-fi writers like Arthur C. Clarke—creator of the cybervillain HAL—as well as younger AI visionaries like Kurzweil, Moravec, and Vinge.

But even Minsky admits that computers are still idiot savants. "I wish I could tell you that we have intelligent machines, but we don't," he says. The world's most powerful computers, he acknowledges, lack the common sense of a toddler; they can't even distinguish cats from dogs unless they are explicitly and painstakingly programmed to do so.

Nevertheless, singularitarians are quite right that, if current trends continue, supercomputers will exceed 10 quadrillion operations per second within a decade. IBM's Blue Gene/P supercomputer, introduced nearly a year ago, can be configured to process up to 3 quadrillion operations per second, although no customer has yet ordered one with the full complement of 884 736 processors that would be needed to get that kind of a processing rate. Argonne National Laboratory, in Illinois, is now completing the upgrade of a Blue Gene/P that should be good for around half a quadrillion operations per second.

So would a fully configured Blue Gene/P be cognitive, perhaps like a monkey or a tree frog, if not like us? Of course not. As any singularitarian would agree, intelligence requires software at least as much as hardware. And that software will soon be available, the singularitarians say, because scientists will in the next couple of decades reverse engineer the brain's software, yielding all sorts of benefits. First, the brain's programming tricks will be transferred to computers to make them smarter. Moreover, given the right interface, our brains and computers will communicate as readily as Macs and PCs. And eventually, of course, our personal software will be extracted from our frail flesh and blood and uploaded into advanced robots or computers. (Don't forget to back yourself up on a hard drive!) We'll walk the earth in impervious titanium-boned bodies. Or we'll inhabit impossibly lush virtual paradises specifically created to please and stimulate our disembodied, digital psyches.

Many neuroscientists do assume that, just as computers operate according to a machine code, the brain's performance must depend on a "neural code," a set of rules or algorithms that transforms those spikes into perceptions, memories, meanings, sensations, and intentions. If such a neural code exists, however, neuroscientists still have no idea what that code is. Or, more accurately, like voters in a U.S. presidential primary, researchers have a surfeit of candidates, each seriously flawed.

THE FIRST NEURAL CODE was discovered more than 70 years ago by the British electrophysiologist Edgar Adrian, who found that when he increased the pressure on neurons involved in the sense of touch, they fired at an increased rate. That so-called rate code has now been demonstrated in many different animals, including *Homo sapiens*. But a rate code is a crude, inefficient way to convey information; imagine trying to communicate solely by humming at different pitches.

Neuroscientists have long suspected that the brain employs subtler codes. One of them might be a temporal code, in which information is represented not just in a cell's rate of firing but also in the precise timing between spikes. For example, a rate code would treat the spike sequences 010101 and 100011 as identical because they have the same number of 0 and 1 bits. But a temporal code would assign different meanings to the two strings because the bit sequences are different. That's a vital distinction: the biophysicist William Bialek of Princeton University calculates that temporal coding would boost the brain's information-processing capacity close to the Shannon limit, the theoretical maximum that information theory allows for a given physical system.

Some neuroscientists suspect that temporal codes predominate in the prefrontal cortex and other brain structures associated with "higher" cognitive functions, such as decision making. In these regions, neurons tend to fire on average only one or two times per second, compared with the 100 or more times of sensory and motor neurons.

Other neural-coding theories abound. On a more macro level, researchers are seeking "population codes" involving the correlated firing of many neurons. Edelman, at the Neurosciences Institute, has advocated a scheme called neural Darwinism, in which our recognition of, say, an animal emerges from competition between large populations of neurons representing different memories: Dog? Cat? Weasel? Rat? The brain quickly settles on the population that most closely matches the incoming stimulus. Perhaps because Edelman has cloaked it in impenetrable jargon, neural Darwinism has not caught on.

Wolf Singer of the Max Planck Institute for Brain Research, in Frankfurt, has won more support for a code involving many neurons firing at the same rate and time. Do such synchronous oscillations play a crucial role in cognition and perhaps even underpin consciousness? Singer thinks they might.

Consciousness is not easy to define, let alone create in a machine. The psychologist William James described it succinctly as attention plus short-term memory. It's what you possess right now as you read this article, and what you lack when you are asleep and between dreams, or under anesthesia.

In 1990, the late Nobel laureate Francis Crick and his colleague Christof Koch proposed that the 40-hertz synchronized oscillations found a year earlier by Singer and his collaborator were one of the neuronal signatures of consciousness. But Singer says the brain probably employs many different codes in addition to oscillations. He also emphasizes that researchers are "only at the beginning of understanding" how neural processes "bring forth higher cognitive and executive functions." And bear in mind that it's still a very long way from grasping those functions to understanding how they give rise to consciousness. And yet without that understanding, it's hard to imagine how anyone could build an artificial brain sophisticated enough to sustain and nurture an individual human consciousness indefinitely.

Given our ignorance about the brain, Singer calls the idea of an imminent singularity "science fiction."

Koch shares Singer's skepticism [see Koch's article in this issue]. A neuroscientist at Caltech, Koch was a close friend and collaborator of Crick, who together with James Watson unraveled the structure of DNA in 1953. During the following decade or so, Crick and other researchers established that the double helix mediates an astonishingly simple genetic code governing the heredity of all organisms. Koch says, "It is very unlikely that the neural code will be anything as simple and as universal as the genetic code."

Neural codes seem to vary in different species, Koch notes, and even in different sensory modes within the same species. "The code for hearing is not the same as that for smelling," he explains, "in part because the phonemes that make up words change within a tiny fraction of a second, while smells wax and wane much more slowly."

Evidence from research on neural prostheses suggests that brains even devise entirely new codes in response to new experiences. "There may be no universal principle" governing neural-information processing, Koch says, "above and beyond the insight that brains are amazingly adaptive and can extract every bit of information possible, inventing new codes as necessary."

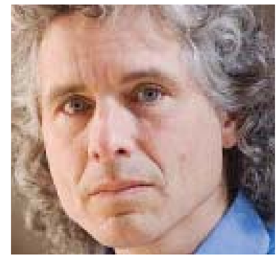

REBECCA GOLDSTEIN

EXPERT VIEW:
Steven Pinker

WHO HE IS
Professor of psychology at Harvard; previously taught in the department of Brain and Cognitive Sciences at MIT, with much of his research addressing language development. Writes best sellers about the way the brain works, including *The Blank Slate* (2002) and *The Stuff of Thought* (2007).

SINGULARITY WILL OCCUR
Never, ever

MACHINE CONSCIOUSNESS WILL OCCUR
"In one sense—information routing—they already have. In the other sense—first-person experience—we'll never know."

MOORE'S LAW WILL CONTINUE FOR
10 more years

THOUGHTS
"There is not the slightest reason to believe in a coming singularity. The fact that you can visualize a future in your imagination is not evidence that it is likely or even possible. Look at domed cities, jet-pack commuting, underwater cities, mile-high buildings, and nuclear-powered automobiles—all staples of futuristic fantasies when I was a child that have never arrived. Sheer processing power is not a pixie dust that magically solves all your problems."

Nothing New About Singularity Fantasies

"THE SINGULARITY" IS JUST the latest manifestation of our hopes, and fears, of techno-transcendence. The industrial revolution inspired dark sci-fi novels such as Mary Shelley's *Frankenstein*, in 1818, and Samuel Butler's *Erewhon*, in 1872; 20th-century thinkers offered more upbeat visions. Some notable examples:

- In his 1929 essay "The World, the Flesh, and the Devil," the Irish scientist and Marxist J.D. Bernal predicts that we will soon start improving our minds and bodies by tinkering with our genes. Eventually, we will abandon our fleshy substrates entirely and evolve into "masses of atoms in space communicating by radiation, and ultimately perhaps resolving [ourselves] entirely into light."
- In *The Future of Man*, a collection of essays published posthumously in 1959, the Jesuit philosopher Pierre Teilhard de Chardin speculates that our minds are becoming increasingly interconnected and will eventually fuse into a collective consciousness. This process will culminate in an "omega point," when we shed our physical selves and converge with the divine consciousness. Teilhard de Chardin is a bit vague on how exactly that will happen.
- In a 1978 lecture, the physicist Freeman Dyson sought to allay concerns that in an eternally expanding, "open" universe, human consciousness will eventually succumb to the "heat death" implicit within the second law of thermodynamics. Dyson calculates that through shrewd conservation of energy, intelligence can persist forever, perhaps in the form of a cloud of "dust grains carrying positive and negative charges, organizing itself and communicating with itself by means of electromagnetic forces."
- In their 1986 book, *The Anthropic Cosmological Principle*, the physicists Frank Tipler and John Barrow offer a rosier picture of our prospects in a closed universe. By the time the universe begins collapsing, they predict, we will have evolved into superintelligent machines that transform the cosmos into one big computer. As the cosmic computer shrinks toward an infinitely compressed omega point, its information-processing capacity spikes toward infinity, and the computer becomes omniscient and omnipotent. In Tipler's 1994 book, *The Physics of Immortality*, he proposes that this God-like computer will resurrect all of us within a virtual paradise, in which all our desires, including sexual ones, are fulfilled. Count me in! —*J.H.*

THEORETICAL QUIBBLES notwithstanding, singularitarians insist that neural prostheses are already leading us toward bionic convergence. By far the most successful prosthesis is the cochlear implant. During the past few decades, about 100 000 hearing-impaired people around the world have been equipped with the devices, which restore hearing by feeding signals from an external microphone to the auditory nerve via electrodes. But as the deaf memoirist Michael Chorost points out, cochlear implants are far from perfect.

In his 2005 book, *Rebuilt: How Becoming Part Computer Made Me More Human*, Chorost recounts how he learned to live with an implant after losing his hearing in 2001. Although thrilled by the device, which restored his social life, he also recognizes its limitations. Because a cochlear implant provides a crude simulacrum of our innate auditory system, it generally requires a breaking-in period, during which technicians tweak the device's settings to optimize its performance. With that assistance, the brain—perhaps by devising a brand-new coding scheme—learns how to exploit the peculiar, artificial signals. Even then, the sound quality is often poor, especially in noisy settings. Chorost says he still occasionally relies on lip reading and contextual guessing to decipher what someone is saying to him. Cochlear implants do not work at all in some people, for reasons that are not well understood.

By far the most ambitious neural-prosthesis program involves computer chips that can restore or augment memory. Researchers at the University of Southern California, in Los Angeles, have designed chips that mimic the firing patterns of tissue in the hippocampus, a minute seahorse-shaped neural structure thought to underpin memory. Biomedical engineering professor Theodore Berger, a leader of the USC program, has suggested that one day brain chips might allow us to instantly upload expertise. But the memory chips are years away from testing. In rats.

DISCUSSIONS OF memory chips leave Andrew Schwartz cold. A neural-prosthesis researcher at the University of Pittsburgh, Schwartz has shown that monkeys can learn to control robotic arms by means of chips embedded in the brain's motor cortex. But no one has any idea how memories are encoded, Schwartz says. "We know so little about the higher functions of the brain that it seems ridiculous to talk about enhancing things like intelligence and memory," he says. Moreover, he says, downloading complex knowledge directly into the brain would require not just stimulating millions of specific neurons but also altering synaptic connections throughout the brain.

That brings us to the interface problem, the most practical obstacle to bionic convergence and uploading. For now, electrodes implanted into the brain remain the only way to precisely observe and fiddle with neurons. It is a much messier, more difficult, and more dangerous interface than most people realize. The electrodes must be inserted into the brain through holes drilled in the skull, posing the risk of infection and brain damage. They often lose contact with neurons; at any one moment an array of 100 electrodes might make contact with only half that many cells. Scar tissue or blood can encrust the electrode, cells around it might shift their position or die, and electrodes have been known to corrode.

Researchers are testing various strategies for improving contact between neurons and electronics. They are making electrodes out of conducting polymers, which are more compatible with neural tissue than silicon or metal; coating electrodes with naturally occurring glues, called cell-adhesion molecules, which helps cells in the brain and elsewhere stick together; and designing electrode arrays that automatically adjust the position of the electrodes to maximize the reception of neural signals.

At Caltech and elsewhere, engineers have designed hollow electrodes that can inject fluids into the surrounding tissue. The fluids could consist of nerve-growth factors, neurotransmitters, and other substances. The nerve-growth factors encourage cells to grow around electrodes, while the neurotransmitters enhance or supplement electrical-stimulation treatment. Neuroscientists are also testing optical devices that can monitor and stimulate neurons, as well as genetic switches that turn neurons on or off.

To be sure, it's promising work. Terry Sejnowski, a neuroscientist at the Salk Institute for Biological Studies, in San Diego, says the new technologies will make it possible "to selectively activate and inactivate specific types of neurons and synapses as well as record from all the neurons in a volume of tissue." That, in turn, might make it possible to build more effective and reliable neural prostheses.

But again, it's a fantastically long way from there to consciousness uploading. Even singularitarians concede that no existing interface can provide what is required for bionic convergence and uploading: the precise, targeted communication, command, and control of billions of neurons. So they sidestep the issue, predicting that all current interfaces will soon yield to very small robots, or "nanobots." Remember the 1966 motion picture *Fantastic Voyage*? That's the basic idea. But try to imagine, in place of Raquel Welch in a formfitting wet suit, robotic submarines the size of blood cells. They infiltrate the entire brain, then record all neural activity and manipulate it by zapping neurons, tinkering with synaptic links, and so on. The nanobots will be equipped with some sort of Wi-Fi so that they can communicate with one another as well as with electronic systems inside and outside the body.

Nanobots have inspired some terrific "X-Files" episodes as well as the Michael Crichton novel *Prey*. But they have as much basis in current research as fairy dust [see "Rupturing the Nanotech Rapture," in this issue].

STEVEN ROSE HAS nothing against technoenhancement. The neurobiologist at England's Open University wears eyeglasses and is proud of his titanium knee and dental implants. He says a lot can be done to improve the brain's performance through improved drugs, neural prostheses, and perhaps genetic engineering. But he calls the claims about imminent consciousness uploading "pretty much crap."

Rose disputes the singularitarians' contention that computers will soon surpass the brain's computational capacity. He suspects that computation occurs at scales above and below the level of individual neurons and synapses, via genetic, hormonal, and other processes. So the brain's total computational power may be many orders of magnitude greater than what singularitarians profess.

Rose also rejects the basic premise of uploading, that our psyches consist of nothing more than algorithms that can be transferred from our bodies to entirely different substrates, whether silicon or glass fibers or as-yet-unimaginable quantum computers. The information processing that constitutes our selves, Rose asserts, evolved within—and may not work in any medium other than—a social, crafty, emotional, sex-obsessed flesh-and-blood primate.

To dramatize that point, Rose poses a thought experiment involving a "cerebroscope," which can record everything that happens in a brain, at micro and macro levels, in real time. Let's say the cerebroscope (hey, maybe it's based on nanobots!) records all of Rose's neural activity as he watches a red bus coming down a street. Could the cerebroscope reconstruct Rose's perception? No, he says, because his neural response to even that simple stimulus grows out of his brain's entire previous history, including the incident in his childhood when a bus almost ran him over.

To interpret the neural activity corresponding to any moment, Rose elaborates, scientists would need "access to my entire neural and hormonal life history" as well as to all his corresponding experiences. Scientists would also need detailed knowledge of the changing social context within which Rose has lived; his attitude toward buses would be different if terrorists recently had attacked one. The implication of his thought experiment is that our psyches will never be totally reducible, computable, predictable, and explainable. Or, disappointingly enough, downloadable into everlasting new containers.

Perhaps the old joke is right after all: If the brain were simple enough for us to understand, we wouldn't be smart enough to understand it.

LET'S FACE IT. The singularity is a religious rather than a scientific vision. The science-fiction writer Ken MacLeod has dubbed it "the rapture for nerds," an allusion to the end-time, when Jesus whisks the faithful to heaven and leaves us sinners behind.

Such yearning for transcendence, whether spiritual or technological, is all too understandable. Both as individuals and as a species, we face deadly serious problems, including terrorism, nuclear proliferation, overpopulation, poverty, famine, environmental degradation, climate change, resource depletion, and AIDS. Engineers and scientists should be helping us face the world's problems and find solutions to them, rather than indulging in escapist, pseudoscientific fantasies like the singularity. ❏

NUMENTA

EXPERT VIEW:
Jeff Hawkins

WHO HE IS
Cofounder of Numenta, in Menlo Park, Calif., a company developing a computer memory system based on the human neocortex. Also founded Palm Computing, Handspring, and the Redwood Center for Theoretical Neuroscience.

SINGULARITY WILL OCCUR
"If you define the singularity as a point in time when intelligent machines are designing intelligent machines in such a way that machines get extremely intelligent very quickly—an exponential increase in intelligence—then it will never happen."

MACHINE CONSCIOUSNESS WILL OCCUR
"If you think dogs and other mammals are conscious, then you will probably think some machines are conscious. If you think consciousness is a purely human phenomenon, then you won't think machines are conscious."

THOUGHTS
"Belief in this idea is based on a naive understanding of what intelligence is. As an analogy, imagine we had a computer that could design new computers (chips, systems, and software) faster than itself. Would such a computer lead to infinitely fast computers or even computers that were faster than anything humans could ever build? No. It might accelerate the rate of improvements for a while, but in the end there are limits to how big and fast computers can run. We would end up in the same place; we'd just get there a bit faster. There would be no singularity."

Who's Who In The Singularity

COMPUTER SCIENTIST and science-fiction writer Vernor Vinge sketched out a seminal modern singularity hypothesis 15 years ago in his essay "The Coming Technological Singularity." Since then lots of other people have taken up the cause, so now you can either anticipate or scoff at many different kinds of singularities. Perhaps the most inclusive variant is the one we label **"technotopia,"** the scenario in which computers reach and exceed human levels of intelligence, essentially resulting in the resolution of all human problems, including illness and limited life span. The **"event horizon"** idea corresponds most closely to Vinge's original: a confluence of technological changes, most likely including machine intelligence, beyond which we

WHO	WHY WE SHOULD LISTEN	KIND OF SINGULARITY	BASIC ARGUMENT	WHAT THEY SAY	YOU MIGHT ALSO WANT TO KNOW
Raymond Kurzweil Entrepreneur, inventor, author of popular books on the singularity 	Created several enormously successful inventions, including ones for pattern, character, and speech recognition; electronic musical instruments; and medical education.	Technotopia—accelerating change	The basic paradigm of Moore's Law—exponentially increasing improvement—will not only hold true indefinitely for logic circuits but will also apply to countless other technologies. It will lead to a singularity that will enable us to upload our consciousness into machines and, in effect, live indefinitely. This singularity will occur in about 15 years.	"I regard the freeing of the human mind from its severe physical limitations of scope and duration as the necessary next step in evolution." —"The Law of Accelerating Returns" (http://www.kurzweilai.net/articles/art0134.html?printable=1) 	Plans to release a movie later this year based on his book *The Singularity Is Near*. To maximize his chances of living until his consciousness can be downloaded, he reportedly takes upward of 200 vitamin pills a day.
Hans Moravec Adjunct professor, Carnegie-Mellon Robotics Institute 	Wrote *Mind Children* and other books on artificial intelligence and robotics. 	Machine intelligence	The processes of the brain that give rise to consciousness arise inevitably and uniquely from chains of physical events and in accordance with physical principles. They are therefore reproducible with a sufficiently powerful computer.	"The 1500-cubic-centimeter human brain is about 100 000 times as large as the retina, suggesting that matching overall human behavior will take about 100 million MIPS of computer power." —"When Will Computer Hardware Match the Human Brain?" *Journal of Evolution and Technology*, 1998 (http://www.jetpress.org/volume1/moravec.htm)	Now works full-time at a small machine-vision start-up.
Nick Bostrom Director, Future of Humanity Institute, Oxford University 	Cofounded World Transhumanist Association; writes extensively on machine intelligence, life extension, and philosophical issues surrounding the singularity.	Technotopia	Assuming that we don't render ourselves extinct, technological progress should lead to superintelligence and indefinitely extended life spans; once the singularity comes near, we will all be kicking ourselves for not having brought it about sooner.	"I would...assign less than a 50% probability to superintelligence being developed by 2033." —"How Long Before Superintelligence?" (http://www.nickbostrom.com/superintelligence.html)	Had a short career as a stand-up comic. 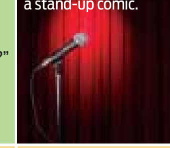
Vernor Vinge Science-fiction author; professor of computer science at San Diego State University for 28 years (retired) 	Wrote the critically acclaimed novels *A Fire Upon the Deep* (Hugo Award, 1993) and *A Deepness in the Sky* (1993); tackled issues that included human confrontation with superintelligent beings. 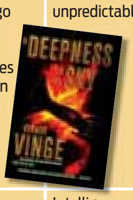	Event horizon—fundamentally unpredictable	In a seminal 1993 essay, he wrote of the period following the development of machine intelligence: "From the human point of view this change will be a throwing away of all the previous rules, perhaps in the blink of an eye, an exponential runaway beyond any hope of control. Developments that before were thought might only happen in 'a million years' (if ever) will likely happen in the next century."	"In my 1993 essay, 'The Coming Technological Singularity,' I said I'd be surprised if the singularity had not happened by 2030. I'll stand by that claim." (p. 69 in this issue)	A chapter of his recent novel, *Rainbows End*, grew out of a story he published in the July 2004 issue of *IEEE Spectrum.*
Eliezer Yudkowsky Research fellow, Singularity Institute for Artificial Intelligence 	Has developed a theory for building "friendly" artificial intelligences whose goals will not change as they themselves evolve.	Intelligence explosion	Each generation of intelligent human-machine collaborations uses its increased intelligence to design the next generation. The more intelligence, the faster the cycle goes, until intelligences appear with capacities far beyond those of unaided humans. Eventually, essentially all the work will be done on the machine side.	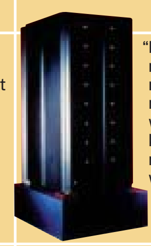 "Deep Blue's engineers might say, 'We have no idea what chess moves the machine will make, but we know they'll be great moves.'" (Interview with author)	Working on the development of machine metaethics so that superintelligences can develop moral reasoning beyond the current level while still retaining predictable characteristics.
Christof Koch Professor of cognitive and behavioral biology, Caltech 	Studies the neural basis of consciousness.	Machine intelligence	To create thinking machines, we have to understand what it is—both biologically and philosophically—that makes humans conscious beings.	"Consciousness does not seem to require many of the things we associate most deeply with being human." (p. 48 in this issue) 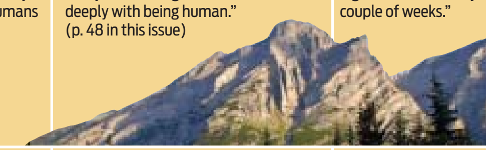	Runs up a 1700-meter-high mountain "every couple of weeks."
Kevin Kelly Senior maverick, *Wired* magazine 	Wrote *Out of Control: The New Biology of Machines, Social Systems, and the Economic World* (1995) and other works on technology and society.	Phase change	Singularities are pervasive changes in the state of the world that are often recognizable only in retrospect. As a result, the singularity is always near, but whether the current wave of technological progress constitutes a singularity is impossible to tell.	"The word *technology* was only coined in 1829, but they'd been doing it for centuries by then." (Interview with author)	Popularized the "Maes-Garreau point," which observes that most predictions of positive technological revolution fall just within the life span of the person doing the predicting.
Bill Joy Venture capitalist 	Cofounded Sun Microsystems.	Event horizon	As computer science, biotech, and nanotechnology advance, it will become easier and easier for small groups or even individuals to create incredibly destructive things. Unless measures are taken, eventually this will happen.	"The future doesn't need us." —"Why the Future Doesn't Need Us," *Wired*, April 2000 (http://www.wired.com/wired/archive/8.04/joy.html)	Has called for voluntary renunciation or at least a delay of research into fields that offer easy extinction of humanity.

ALL PORTRAITS COURTESY OF SUBJECTS—EXCEPT KURZWEIL: HELENE DELILLO/KURZWEIL TECHNOLOGIES; JOY: LYNN GOLDSMITH/CORBIS; MINSKY: INGO WAGNER/DPA/LANDOV; BROOKS: IROBOT CORP.; PENROSE: JERRY BAUER cannot see into the future to predict anything. **"Machine intelligence"** represents the prediction that computers will achieve human-level intelligence and either consciousness or something indistinguishable from it. **"Uploading,"** which is also part of the technotopian vision, refers to the prediction that individual human brains (and most probably their state of consciousness) will be replicable in computerized form, leading to (among other things) virtual immortality. Kevin Kelly's **"phase change"** hypothesis likens the singularity to the changes between the solid and liquid or liquid and gaseous states of matter: the world becomes fundamentally different but in ways that may not be immediately perceptible to individuals immersed in it at the time. The **"intelligence explosion"** hypothesis, initially developed by Turing associate I.J. Goode, postulates that as machines (initially under the direction of humans) design the next generation of machines, eventually superintelligent levels will be reached. The effect is similar to the accelerating rate of technological change anticipated by technotopians, but the path for getting there is different. *—Paul Wallich*

WHO	WHY WE SHOULD LISTEN	KIND OF SINGULARITY	BASIC ARGUMENT	WHAT THEY SAY	YOU MIGHT ALSO WANT TO KNOW
Marvin Minsky Professor of media arts and sciences and of electrical engineering and computer science, MIT 	Built the first neural network simulator in 1951. Received the Turing Award in 1969. Developed a theory of mind involving small subunits of expertise.	Machine intelligence/ uploading	Our increasing knowledge of the brain and increasing computing power will eventually intersect. Currently, however, the resources being devoted to such work are negligible compared with the amount of work to be done. 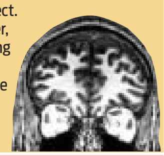	"I asked the audience how many people wanted to live for 200 years, and no one raised their hand." —"It's 2001: Where Is HAL?" *Information Week* podcast, 1 March 2007 (http://www.informationweek.com; search on "minsky podcast")	Thinks that belief in a singularity could inspire people in the way that World War II inspired his generation.
Daniel Dennett Codirector of the Center for Cognitive Studies and professor of philosophy, Tufts University 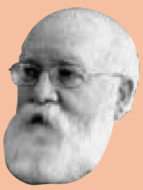	Has written many books about the nature of consciousness and intelligence, including *Brainstorms: Philosophical Essays on Mind and Psychology* (1978), *The Intentional Stance* (1989), *Consciousness Explained* (1991), and *Brainchildren—Essays on Designing Minds* (1998).	Machine intelligence	Human-level AI may be inevitable, but don't expect it anytime soon. "I don't deny the possibility a priori; I just think it is vanishingly unlikely in the foreseeable future."	"The best reason for believing that robots might someday become conscious is that we human beings are conscious, and we are a sort of robot ourselves." (From Dennett's book *Cognition, Computation, and Consciousness*, 1994)	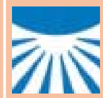 Leading member of the Brights, a group dedicated to promoting "the civic understanding and acknowledgment of the naturalistic worldview, which is free of supernatural and mystical elements."
Rodney Brooks Professor of robotics, MIT 	Has built many robots that emulate the behavior of simple animals. One crucial aspect of his work is demonstrating how little cognition, understanding, or memory is required for apparently complex behaviors.	Technotopia	Any extrapolation of technological trends, particularly exponential ones, decades into the future isn't likely to be accurate.	"I am a machine. So are you." (p. 63 in this issue) 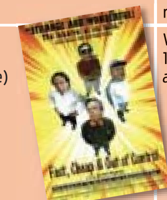	Was featured in the 1997 movie *Fast, Cheap, and Out of Control*.
Jaron Lanier Interdisciplinary scholar-in-residence at the Center for Entrepreneurship & Technology, University of California, Berkeley 	Coined the term *virtual reality*; was CEO of VPL Research, a pioneering virtual-reality start-up. 	Machine intelligence	There's no real evidence that computer representation of a brain will actually produce a mind. Furthermore, true AI can't arise from current patterns of software development because those patterns are fatally flawed. Meanwhile, accelerating technological change could still make the world as we know it nearly unrecognizable.	"If computers are to become smart enough to design their own successors, initiating a process that will lead to God-like omniscience after a number of ever swifter passages from one generation of computers to the next, someone is going to have to write the software that gets the process going, and humans have given absolutely no evidence of being able to write such software." —"One-Half of a Manifesto," *Wired*, December 2000 (http://www.wired.com/wired/archive/8.12/lanier.html)	Wrote an entertaining 1995 essay in the *Journal of Consciousness Studies* titled "You Can't Argue With a Zombie." It zinged Daniel Dennett, Daniel Dennett's critics, Dartmouth students, and philosopher David Chalmers, among others. Sample line: "Arrogance is always a bad strategy in science. In philosophy I suppose it's fine."
John Holland Professor of computer science and engineering and professor of psychology, University of Michigan 	Invented genetic programming, one of the key technologies the singularitarians have claimed will lead to superintelligent machines.	Machine intelligence	In order to combine building blocks and selectively "breed" them to form ever more intelligent programs, you have to have the building blocks in the first place. 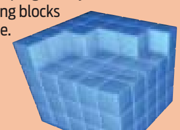	Uncritical believers in the singularity "think evolution is like monkeys at the typewriter, and if you just type fast enough you'll get somewhere." (Interview with author)	Believes that the more people know about the technologies that will supposedly bring about the singularity, the more aware they will be of the limitations of those technologies.
John Searle Professor of philosophy, University of California, Berkeley 	Has written many books on the brain and consciousness, including *Minds, Brains, and Science* (1985) and *The Mystery of Consciousness* (1997). 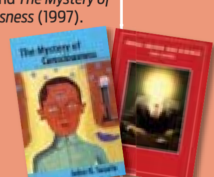	—	His "Chinese Room" thought experiment explains why a bare CPU cannot understand the intent of the program it executes. Only organisms embodied in the real world, with real experiences to draw on, can become truly conscious and intelligent. All else is merely symbol manipulation.	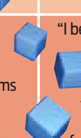 "I believe that there is no objection in principle to constructing an artificial hardware system that would duplicate the powers of the brain to cause consciousness using some chemistry different from neurons. But to produce consciousness any such system would have to duplicate the actual causal powers of the brain." —"I Married a Computer," review of Kurzweil's *The Age of Spiritual Machines*, *The New York Review of Books*, 8 April 1999	Was a member of the UC Berkeley Free Speech movement in the early 1960s.
Roger Penrose Physicist and professor of mathematics, Oxford University 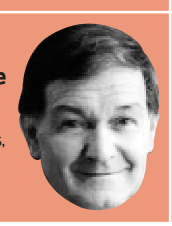	Expounded on Albert Einstein's general theory of relativity. Wrote *The Nature of Space and Time* with Stephen Hawking and *The Emperor's New Mind*, about computation and consciousness.	—	Consciousness cannot be duplicated in computational machines, because it depends on "noncomputational physical processes." Does not know what these might be but suggests it emerges from "large-scale" quantum-mechanical phenomena in microtubules in the brain's neurons.	"I'm not saying that consciousness is beyond physics...although I'm saying that it's beyond the physics we know now." (From Penrose's book *The Third Culture: Beyond the Scientific Revolution*, 1995)	Invented Penrose tiles, a shape capable of covering an infinite plane in a nonperiodic fashion. 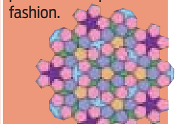

Singularity color code: True believer: thinks it will happen within 30 years · Yes, but... · Maybe someday · No way

ALL OTHER PHOTOS: ISTOCKPHOTO—EXCEPT DEEP BLUE: IBM; VIRTUAL REALITY: NASA

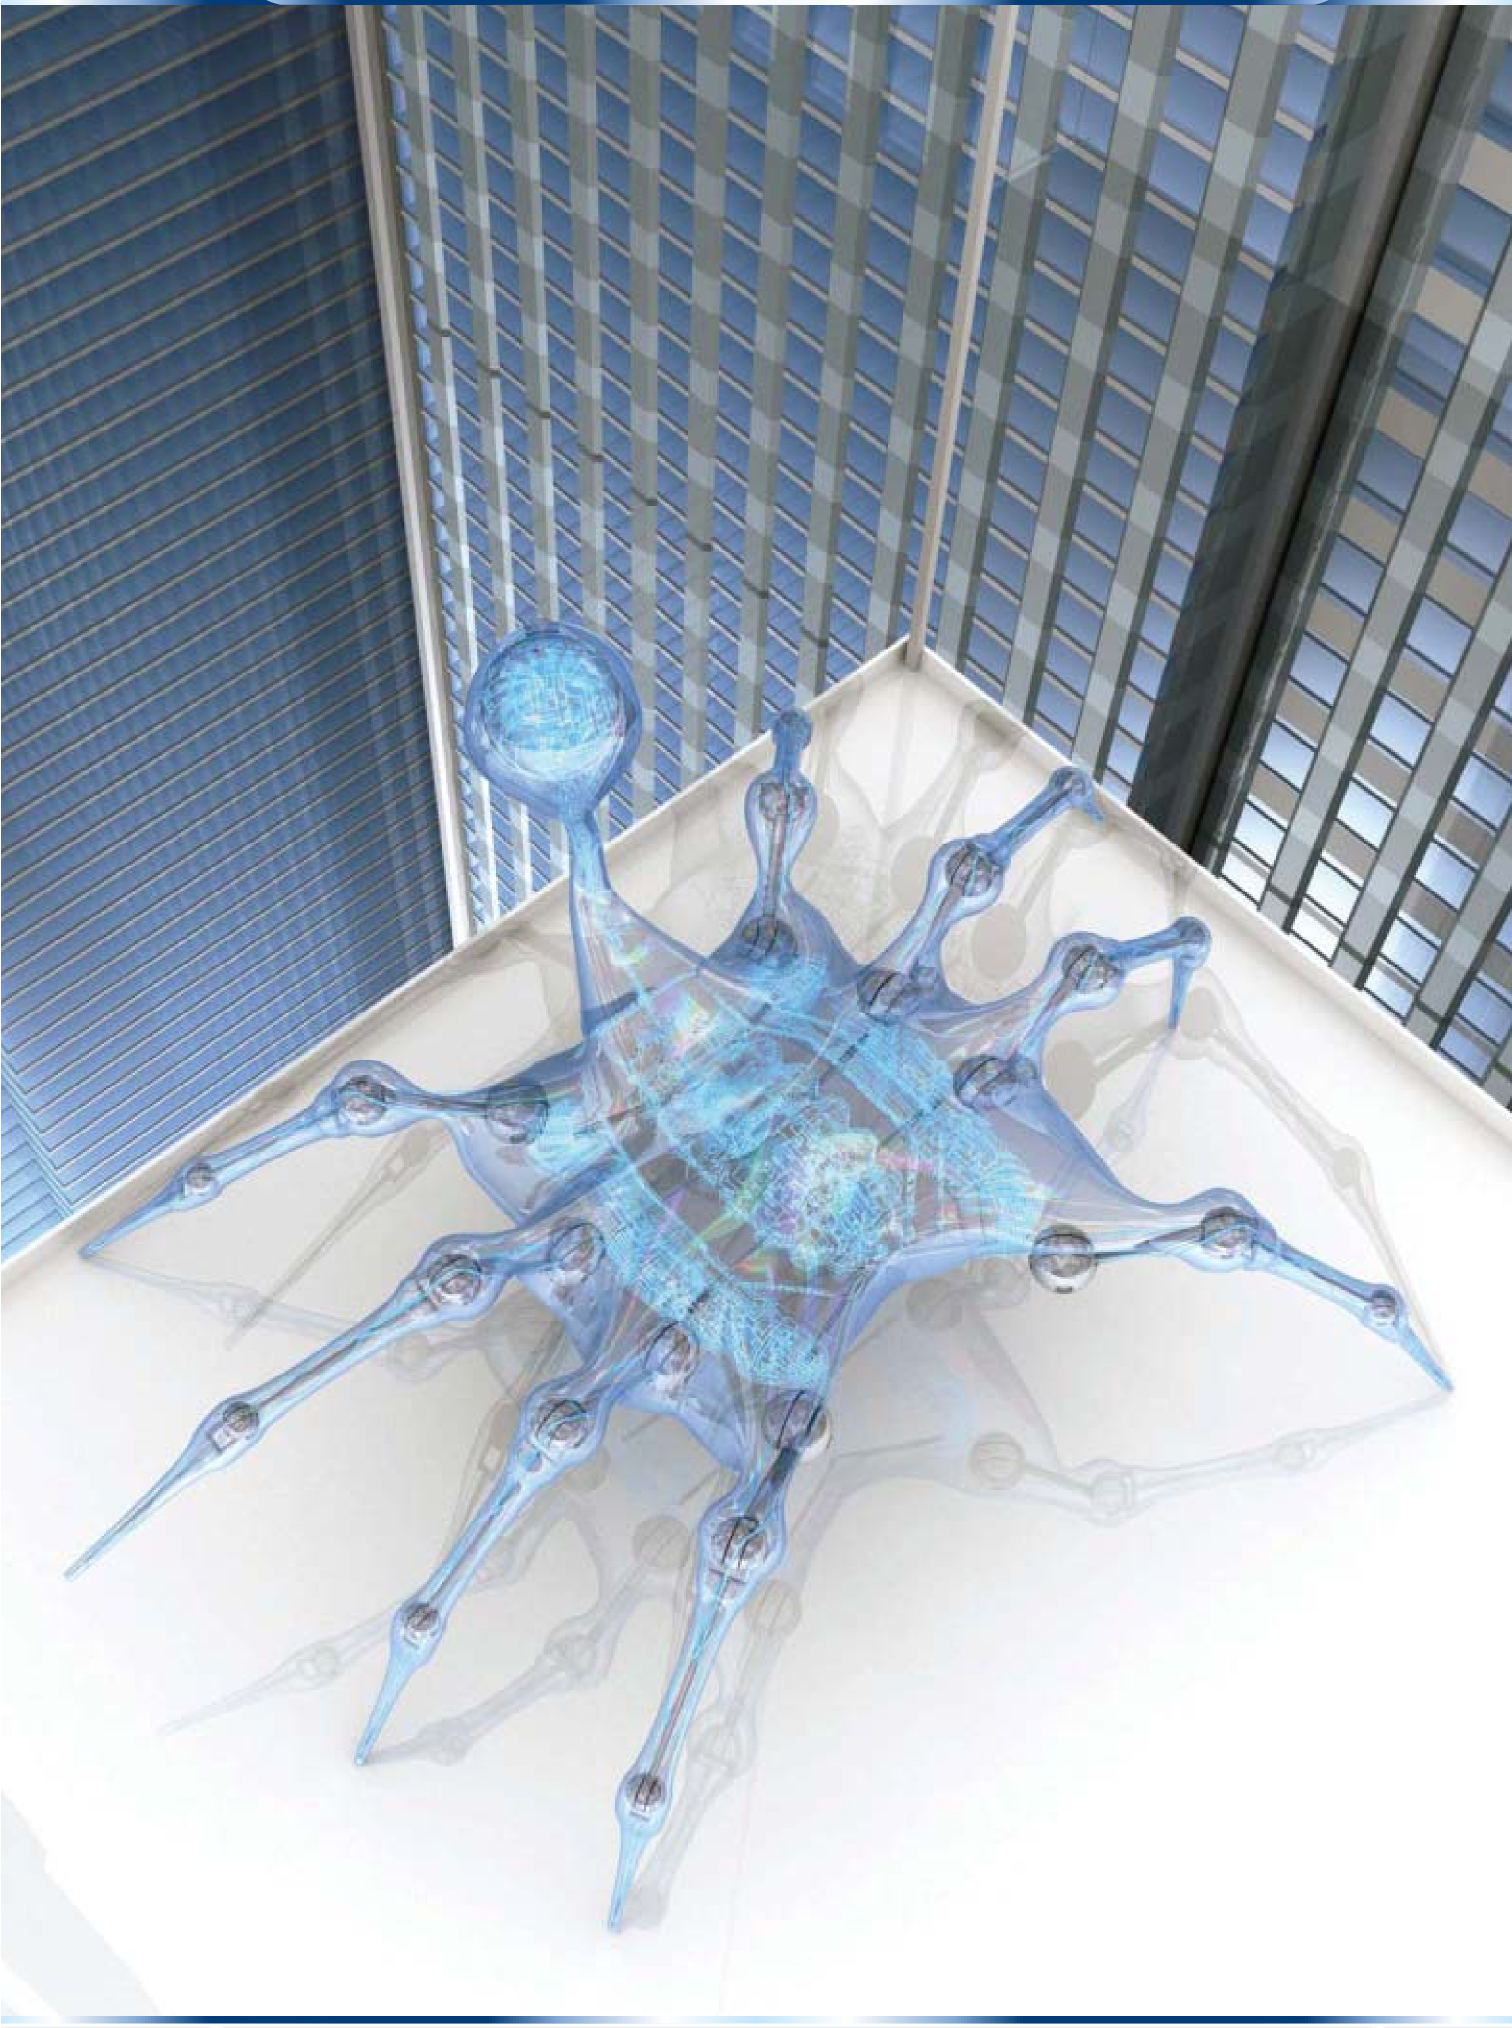

STUFFED INTO SKYSCRAPERS BY THE BILLION, BRAINY BUGBOTS WILL BE THE KNOWLEDGE WORKERS OF THE FUTURE BY ROBIN HANSON

Economics Of The Singularity

OUR GLOBAL ECONOMY would stupefy a Roman merchant as much as the Roman economy would have confounded a caveman. But we would be similarly amazed to see the economy that awaits our grandchildren, for I expect it to follow a societal discontinuity more dramatic than those brought on by the agricultural and industrial revolutions. The key, of course, is technology. A revolutionary speedup in economic growth requires an unprecedented and remarkable enabling tool. Machine intelligence on a human level, if not higher, would do nicely. Its arrival could produce a singularity—an overwhelming departure from prior trends, with uneven and dizzyingly rapid change thereafter. A future shock to end future shocks.

Yes, this theory of mine is a social and economic one, and therefore not as unfailingly accurate or testable as one in the physical sciences. Nevertheless, social scientists routinely make short-term forecasts that hit the mark, and economists often offer insightful forecasts about unprecedented situations.

So indulge me as I outline how we economists view technological change. In so doing, I hope to explain why it's reasonable to view past history as a series of abrupt, seemingly unheralded transitions from one economic era to another, transitions marked by the sudden and drastic increase in the rate of economic growth. I will then show why another singularity is perhaps just around the corner. Finally, I will outline its possible consequences.

A complex device, like a tractor or a building, can have thousands of parts, and each part can rely on dozens of technologies. Yet in most cases even a spectacular gain in the quality of one part bestows at best only a small improvement on the whole. Keep improving a part in successive increments of equal degree and you'll get ever smaller gains to the whole. This is the law of diminishing returns, and it applies not only to devices and organizations but to entire industries. Consider your personal computer: every couple of years its power-to-cost ratio has doubled, and yet as you go from one generation to the next, you probably notice only

BRYAN CHRISTIE DESIGN

How Many Singularities Have There Been?

THE TWO SOLIDLY DEMONSTRATED singularities, the agricultural and industrial revolutions, came with little warning.

Were there any singularities before farming and industry? If we look back further in time, we can find even slower modes of growth that made sudden transitions to faster modes. For example, human hunter-gatherers vastly expanded their niche and spread throughout the world in a biologically short period of time. That transition apparently was made possible by special innovations in the unusually large protohuman brain. Before that transition, and after the emergence of animals some 500 million years earlier, the largest animal brains doubled in size roughly every 30 million years—less than 1 percent of the growth rate of human brains.

Looking further back, it is difficult to find long-term trends that may have paved the way for the emergence of animals. Still, it is interesting to note that the volume of our nearly 14 billion-year-old universe has been expanding exponentially due to a mysterious "dark energy," with a doubling time of 3 billion years—about 1 percent the growth rate of animal brain size.

Of course, because we have no accepted theory saying why various growth modes and transitions should be related, any similarities between them may be pure coincidence. But they do constitute precedents, for they show that vast changes can appear seemingly overnight. —*R.H.*

The world economy, which now doubles in 15 years, would soon double in a week to a month

a small improvement as you plug away on your word processors and spreadsheets.

It turns out that most of these small, innovative gains come not from research labs but from hands-on builders and users. So the more a thing gets used, the more it tends to improve. It doesn't matter whether that thing is a physical device, such as a car, or a social organization, such as a corporation.

If any large system of interacting parts tends to improve by smooth gradations, then we should expect systems of systems, with their larger number of components and interactions, to improve even more smoothly. By this reasoning, the world economy should improve most smoothly of all. The world economy consists of the largest number of interacting parts of any man-made system, and everyone not stranded on an uncharted island contributes to the improvements in all those parts by using them. Finally, in each economic era the question of whether growth speeds up or slows down depends on two competing factors. Deceleration typically ensues as innovators exhaust the easy ideas—the low-hanging fruit. But acceleration also ensues as the economy, by getting larger, enables its members to explore an ever-increasing number of innovations.

We have the tools to measure the world's economic product not only for today—it's about US $50 trillion per year—but also for times long past. A few years ago Angus Maddison, an economic historian at the University of Groningen, in the Netherlands, plotted a graph of world economic product—basically everything of value produced globally: bananas, submarines, magazine articles, you name it. It shows that from 1950 to 2003, growth was relatively steady. During that time, despite enormous technical change, no particular technology left much of a fingerprint on the data; no short-term accelerations in growth could be attributed to this or that technological development. Also, Maddison's data offer little support for the idea that innovation and growth have been accelerating recently.

Now look at the data for world product over the past 7,000 years, estimated by Bradford DeLong, an economic historian at the University of California, Berkeley. The data here tell a somewhat different story. For most of that time, growth proceeded at a relatively steady exponential rate, with a doubling of output about every 900 years. But within the past few centuries, something dramatic happened: output began doubling faster and faster, approaching a new steady doubling time of about 15 years. That's about 60 times as fast as it had been in the previous seven millennia.

WE CALL THIS transition the Industrial Revolution, but that does not mean we understand it well or even know precisely how and why it arose. But whatever the Industrial Revolution was, clearly it was an event worthy of the name "singularity."

If we look further back, we see what appears to be at least one previous singularity—the transition to an economy based on agriculture. And slow as economic growth during the agricultural era may seem in the aftermath of the Industrial Revolution, it was actually lightning fast compared with that of the economic era that came before, which was based on hunting and gathering.

In the roughly 2 million years our ancestors lived as hunters and gatherers, the population rose from about 10 000 protohumans to about 4 million modern humans. If, as we believe, the growth pattern during this era was fairly steady, then the population must have doubled about every quarter million years, on average. Then, beginning about 10 000 years ago, a few of those 4 million humans began to settle down and live as farmers. The resulting communities grew so fast that they quickly accounted for most of the world population. From that time on, the farming population doubled about every 900 years—some 250 times as fast as before.

Our understanding of the existence, nature, and relevance of these transitions clearly becomes more speculative the further back we look in time [see sidebar, "How Many Singularities Have There Been?"]. There may well have been two earlier singularities that started eras of this sort, although our ability to identify them and weigh their relevance is very speculative. I suggest an era defined by the growth of the brain from the emergence of animal life to the first protohumans and perhaps an earlier era defined by the growth of the universe from a time shortly after the big bang to the first animals.

So we have perhaps five eras during which the thing whose growth is at issue—the universe, brains, the hunting economy, the farming economy, and the industrial economy—doubled in size at fixed intervals. Each era of growth before now, however, has eventually switched suddenly to a new era having a growth rate that was between 60 and 250 times as fast. Each switch was completed in much less time than it had taken the previous regime to double—from a few millennia for the agricultural revolution to a few centuries for the industrial one. These switches constituted singularities.

Whatever may have been the key innovations behind these transitions, it is clear that they were far more potent than such familiar textbook examples of great innovations as fire, writing, computers, or plastics. Most innovations happen within a given growth era and do not change its basic nature, including its basic growth rate. A few exceedingly rare innovations, however, do suddenly change everything. One such innovation led to agriculture; another led to industry.

Therefore, we must admit that another singularity—at least the third one, and perhaps the fifth, depending on how you count—could lie ahead. Furthermore, data on these previous apparently similar singularities are some of the few concrete guides available to what such a transition might look like. We would be fools if we confidently expected all patterns to continue. But it strikes me as pretty foolish to ignore the patterns we see.

If a new transition were to show the same pattern as the past two, then growth would quickly speed up by between 60- and 250-fold. The world economy, which now doubles in 15 years or so, would soon double in somewhere from a week to a month. If the new transition were as gradual (in power-law terms) as the Industrial Revolution was, then within three years of a noticeable departure from typical fluctuations, it would begin to double annually, and within two more years, it might grow a million-fold. If the new transition were as rapid as the agricultural revolution seems to have been, change would be even more sudden.

Though such growth may seem preposterous, consider that in the era of hunting and gathering, the economy doubled nine times; in the era of farming, it doubled seven times; and in the current era of industry, it has so far doubled 10 times. If, for some as yet unknown reason, the number of doublings is similar across these three eras, then we seem already overdue for another transition. If we instead compare our era with the era of brain growth, which doubled 16 times before humans appeared, we would expect the next transition by around 2075.

What innovation could possibly induce so fabulous a speedup in economic growth? It is easier to say what could not. Because of diminishing returns, no change that improved just one small sector of the economy could do the trick. In advanced countries today, farming, mining, energy, communications, transportation, and construction each account for only a small percentage of economic activity. Even so extraordinary an innovation as radical nanotechnology would do no more than dramatically lower the cost of capital for manufacturing, which now makes up less than 10 percent of U.S. GDP.

No, the next radical jump in economic growth seems more likely to come from something that has a profound effect on everything, because it addresses the one permanent shortage in our entire economy: human time and attention. They are by far the most productive components of today's economy. About two-thirds of all income in the rich countries is paid directly for wages, and much of the remaining third represents indirect costs of labor. (For example, corporate income largely reflects earlier efforts by entrepreneurs.) So any innovation that could replace or dramatically improve human labor would be a very big deal.

One of the pillars of the modern singularity hypothesis in its many forms is that intelligence is a general elixir, able to cure many if not all economic ailments. Typically, this belief is expressed in the form of an argument that the arrival of very intelligent machines will produce the next singularity. Some people hope this arrival will follow a new Einstein, who will discover a powerful general theory of intelligence applicable to those machines. Others envision an "intelligence explosion" via a series of powerful design innovations, beginning with one that would make machines smart enough to help us quickly find a second innovation, allowing even smarter machines, and so on. A few even imagine innovations so unprecedentedly potent that a single machine embodying the first innovation could go through the entire innovation series by itself, unnoticed, within a week, and then take over the world.

There are many views on how intelligence might arise in a machine. One argument holds that hardware is the critical limiting factor and predicts that human-level machine intelligence will come soon after we have computer hardware whose performance is comparable with that of the human brain.

Another argument focuses on knowledge as the true limiting factor. This view is behind several huge artificial-intelligence database projects, including Cyc, under construction for 23 years

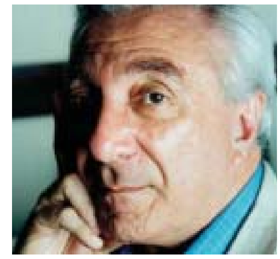

JOHN CASTI

EXPERT VIEW:
John Casti

WHO HE IS
Senior Research Scholar, the International Institute for Applied Systems Analysis, in Laxenburg, Austria, and cofounder of the Kenos Circle, a Vienna-based society for exploration of the future. Builds computer simulations of complex human systems, like the stock market, highway traffic, and the insurance industry. Author of popular books about science, both fiction and nonfiction, including *The Cambridge Quintet*, a fictional account of a dinner-party conversation about the creation of a thinking machine.

SINGULARITY WILL OCCUR
Within 70 years

MACHINE CONSCIOUSNESS WILL OCCUR
Questionable

MOORE'S LAW WILL CONTINUE FOR
20 more years with current technology

THOUGHTS
"I think it's scientifically and philosophically on sound footing. The only real issue for me is the time frame over which the singularity will unfold. [The singularity represents] the end of the supremacy of *Homo sapiens* as the dominant species on planet Earth. At that point a new species appears, and humans and machines will go their separate ways, not merge one with the other. I do not believe this necessarily implies a malevolent machine takeover; rather, machines will become increasingly uninterested in human affairs just as we are uninterested in the affairs of ants or bees. But in my view it's more likely than not that the two species will comfortably and more or less peacefully coexist—unless human interests start to interfere with those of the machines."

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

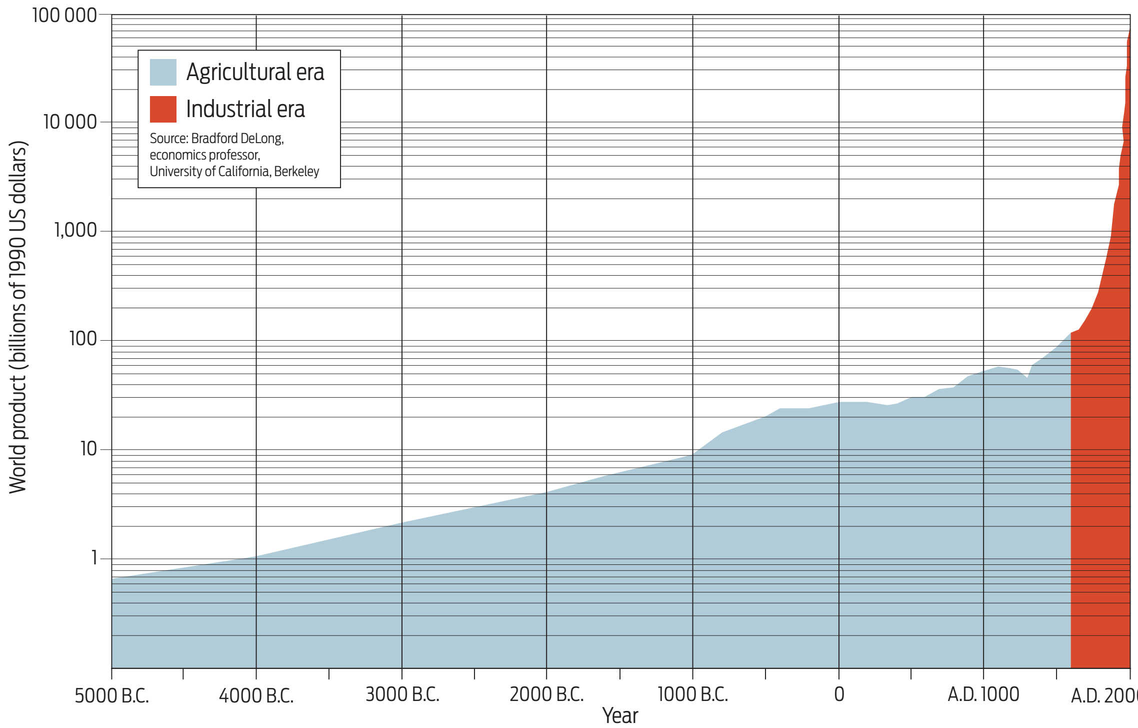

SPOT THE TRANSITION: After millennia of slow growth tied to rising population, the world economy took off. For the first time ever, output per capita rose above mere subsistence levels. It's been rising ever since.

and now at Cycorp in Austin, Texas. Cyc now possesses millions of pieces of commonsense knowledge, added mostly by hand. Eventually, Cyc may know enough to begin to read and assimilate all written knowledge, and the more it knows, the faster it should be able to learn. So it is possible, though hardly inevitable, that Cyc will eventually undergo a rapid knowledge explosion.

I find those scenarios interesting but unlikely to come to pass anytime soon. Regarding advanced machine intelligence, my guess is that our best chance of achieving it within a century is to put aside the attempt to understand the mind, at least for now, and instead simply focus on copying the brain.

This approach, known as whole brain emulation, starts with a real human brain, scanned in enough detail to see the exact location and type of each part of each neuron, such as dendrites, axons, and synapses. Then, using models of how each of these neuronal components turns input signals into output signals, you would construct a computer model of this specific brain. With accurate enough models and scans, the final simulation should have the same input-output behavior as the original brain. It would, in a sense, be the "uploaded mind" of whoever served as the template. Whether the emulation indeed constitutes a person and whether that person has rights is another story, to which I will return later.

If current trends continue, we should have computer hardware and brain scans fast and cheap enough to support this scenario in a few decades. What may well take longer are input-output models in sufficient detail for every relevant type of human neuron part. But I think those details will accrue in time. We already have sufficient models for some types of neuronal components, gathered after only a modest effort. And we have no reason to expect the other types to be harder.

An emulation of a brain could merely do what that brain can already do, although if done in sufficiently powerful hardware, the cognition might occur faster. Still, even if all we were able to achieve was a computer with the mental powers of a particular human, that would be more than just interesting—it would also be incredibly useful.

Though it might cost many billions of dollars to build one such machine, the first copy might cost only millions and the millionth copy perhaps thousands or less. Mass production could then supply what has so far been the one factor of production that has remained critically scarce throughout human history: intelligent, highly trained labor.

> Brain emulation would simulate the "uploaded mind" of whoever served as the template

Project Blue Brain, a joint effort by IBM and the Ecole Polytechnique Fédérale de Lausanne, in Switzerland, has made some impressive progress: in December 2006, the project finished mapping and modeling the 10 000-odd neurons and 30 million synapses in a rat's neocortical column. Similarly impressive, in 2004 a Stockholm University team observed realistic behavior in a simulation of 8 million neurons and 4 billion synapses. But we still have far to go.

Okay, so might these machines be conscious, with wills of their own, and if so, could they be selfish, even malevolent? Yes, yes, yes, and yes. More on that later; for now, let's get back to the economic argument.

Creating human-level intellect in a machine would be an astounding achievement, but it is not immediately obvious that it would launch a new era of much faster growth, with doubling times measured in months or less. After all, more and more capable machines have been replacing and aiding humans for centuries without sparking such an explosion. To answer that objection, we've got to start with the fundamentals: what economic theory says about growth rates.

To keep a modern economy thriving, we must accomplish many mental tasks. Some people (we call them engineers) have to design new products, systems, and services. Other people have to build, market, transport, distribute, and maintain them, and so on. These myriad tasks are mostly complements, so that doing one task better increases the value of doing other tasks well. But for each task, humans and machines may also be substitutes; it can be a wasted effort to have them both do the same task.

The relative advantages of humans and machines vary from one task to the next. Imagine a chart resembling a topographic cross section, with the tasks that are "most human" forming a human advantage curve on the higher ground. Here you find chores best done by humans, like gourmet cooking or elite hairdressing. Then there is a "shore" consisting of tasks that humans and machines are equally able to perform and, beyond them an "ocean" of tasks best done by machines. When machines get cheaper or smarter or both, the water level rises, as it were, and the shore moves inland.

This sea change has two effects. First, machines will substitute for humans by taking over newly "flooded" tasks. Second, doing machine tasks better complements human tasks, raising the value of doing them well. Human wages may rise or fall, depending on which effect is stronger.

For example, in the 1920s, when the mass-produced automobile came along, it was produced largely by machines, with human help. So machines dominated that function—the assembly of cars. The resulting proliferation of machine-assembled cars raised the value of related human tasks, such as designing those cars, because the financial stakes were now much higher. Sure enough, automobiles raised the wages of machinists and designers—in these cases, the complementary effect dominated. At the same time, the automobile industry lowered the pay of saddle makers and stable hands, an example of the substitution effect.

So far, machines have displaced relatively few human workers, and when they have done so, they have in most cases greatly raised the incomes of other workers. That is, the complementary effect has outweighed the substitution effect—but this trend need not continue.

In our graph of machines and humans, imagine that the ocean of machine tasks reached a wide plateau. This would happen if, for instance, machines were almost capable enough to take on a vast array of human jobs. For example, it might occur if machines were on the very cusp of human-level cognition. In this situation, a small additional rise in sea level would flood that plateau and push the shoreline so far inland that a huge number of important tasks formerly in the human realm were now achievable with machines. We'd expect such a wide plateau if the cheapest smart machines were whole-brain emulations whose relative abilities on most tasks should be close to those of human beings.

In such a scenario, the economy would start growing much faster, for three reasons. First, we could create capable machines in much less time than it takes to breed, rear, and educate new human workers. Being able to make and retire machine workers as fast as needed could easily double or quadruple growth rates.

Second, the cost of computing has long been falling much faster than the economy has been growing. When the workforce is largely composed of computers, the cost of making workers will therefore fall at that faster rate, with all that this entails for economic growth.

Third, as the economy begins growing faster, computer usage and the resources devoted to developing computers will also grow faster. And because innovation is faster when more people use and study something, we should expect computer performance to improve even faster than in the past.

Together these effects seem quite capable of producing economic doubling times much shorter than anything the world has ever seen. And note that this forecast does not depend on the rate at which we achieve machine intelligence capabilities or the rate at which the intelligence of machines increases. Merely having computer-like machines able to do most important mental tasks as well as humans do seems sufficient to produce very rapid growth.

CYPRESS SEMICONDUCTOR

EXPERT VIEW:
T.J. Rodgers

WHO HE IS
Founder and CEO of Cypress Semiconductor Corp., in San Jose, Calif., known for his brash opinions about the business world and politics. Owner of the Clos de la Tech winery and vineyards, in California, where he's trying to make the best American pinot noir.

SINGULARITY WILL OCCUR
Never

THOUGHTS
"I don't believe in technological singularities. It's like extraterrestrial life—if it were there, we would have seen it by now. However, I do believe in something that is more powerful because it is real—namely, exponential learning. An exponential function has the property that its slope is proportional to its value. The more we know, the faster we can learn.

"Technological transitions are required to maintain an exponential rate of learning. The first airplanes were certainly not as good as well-appointed trains in moving masses comfortably, but the transition later proved essential to maintaining our progress in human mobility. Gene splicing is a breakthrough technology but has not yet done (or been allowed to do) a lot for mankind. That will change.

"I don't believe in the good old days. We will be freer, more well-educated and even smarter in the future—but exponentially so, not as a result of some singularity."

SPECTRUM | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

LIFE IN A robot economy would not be merely a sped-up version of our lives today. When I apply basic economic theory and some common sense to this scenario, I conclude that humans would probably be neither the immortal, all-powerful gods that some hope for nor the hated and hunted prey that some fear.

Yes, robot-human wars would be possible, but it is important to remember that few differences between humans ever lead to war. We do not fear that the short will conspire to murder the tall in their sleep, nor that the right-handed will exterminate the left-handed. Short, tall, left-handed, and right-handed people all trade with, befriend, and marry one another with abandon, making such wars almost unthinkable. Instead, wars today happen between largely separate nations and ethnic groups. Similarly, robots well-integrated into our economy would be unlikely to exterminate us.

Wages could fall so far that most humans could not live on them

Would robots be slaves? Laws could conceivably ban robots or only allow robots "born" with enough wealth to afford a life of leisure. But without global and draconian enforcement of such laws, the vast wealth that cheap robots offer would quickly induce a sprawling, unruly black market. Realistically, since modest enforcement could maintain only modest restrictions, huge numbers of cheap (and thus poor) robots would probably exist; only their legal status would be in question. Depending on local politics, cheap robots could be "undocumented" illegals, legal slaves of their creators or owners, "free" minds renting their bodies and services and subject to "eviction" for nonpayment, or free minds saddled with debts and subject to "repossession" for nonpayment.

The following conclusions do not much depend on which of these cases is more common. For example, in any of these cases human wages would rise or fall rapidly, depending on the shape of the human advantage landscape. After the flood of the plateau, there might still be some mountain peaks of human tasks left. Some rich people might still want to be served and entertained by real human beings. So for those jobs, human wages could rise. But if in the end the machine ocean completely inundated all of Task-Land, then wages would fall so far that most humans would not, through their labor alone, be able to live on them, though they might work for other reasons.

In either case, human labor would no longer earn most income. Owners of real estate or of businesses that build, maintain, or supply machines would see their wealth grow at a fabulous rate—about as fast as the economy grows. Interest rates would be similarly great. Any small part of this wealth should allow humans to live comfortably somewhere, even if not as all-powerful gods.

Because copying a machine mind would be cheap, training and education would cost no more than a software update. Instead of long years to train each worker, a few machines would be trained intensely, and then many copies would be made of the very best trainees. Presumably, strong security would prevent bootleg copies.

Organizational decision cycles would shorten, favoring streamlined, decentralized processes run by fast machine minds in key positions of authority. Fast minds could be whole-brain emulations sped up relative to human brains. This scenario would marginalize slow bureaucratic human committees, regulators, and the like. Fast growth rates would likely discourage slow long-distance transport and encourage local production.

Some robots responsible for administration, research, law, and other cognitive work might live and work entirely in virtual environments. For others, crude calculations suggest that tiny bodies a few millimeters tall, with sped-up minds to match their faster body motions, might allow insectlike urban densities, with many billions living in the volume of a current skyscraper, paying astronomical rents that would exclude most humans.

As emulations of humans, these creatures would do the same sorts of things in their virtual realities and skyscrapers that humans have done for hundreds of thousands of years: form communities and coalitions, fall in love, gossip, argue, make art, commit crimes, get work done, innovate, and have fun. Just as farming was more alien to our human nature than hunting and gathering, and industry was more alien still, their world would be even more distant from human origins. But human nature seems flexible enough to accommodate such changes.

The population of smart machines would explode even faster than the economy. So even though total wealth would increase very rapidly, wealth per machine would fall rapidly. If these smart machines are considered "people," then most people would be machines, and per-person wealth and wages would quickly fall to machine-subsistence levels, which would be far below human-subsistence levels. Salaries would probably be just high enough to cover the rent on a tiny body, a few cubic centimeters of space, the odd spare part, a few watts of energy and heat dumping, and a Net connection.

While copying would make robot immortality feasible in principle, few robots would be able to afford it. And when reproduction via copying dominates, few robots would be able to afford robot versions of human children.

While whole-brain-emulation robots would be copies of particular humans, we should expect vast inequality in copy rates. Investors who paid the high costs for scanning a human brain would carefully select the few humans most likely to be flexible, cooperative, and productive workers, even while living a short, hardscrabble, childless, and alien life in robotic bodies or virtual offices. Investors who paid for copying existing machine minds would select robots with a track record of achieving this ideal. As a result, there would be large first-mover advantages and winner-take-all effects. For example, if docile minds turned out to be the most productive, then the robot world might consist mainly of trillions of copies each of a few very docile human minds.

In this case, the meek would indeed inherit the Earth. ❑

TO PROBE FURTHER For additional resources on reconstructing the deep economic past, speculations on a rapid intelligence explosion, and the likely effects of machine intelligence on economic growth, see http://spectrum.ieee.org/jun08/singularityprobe.

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

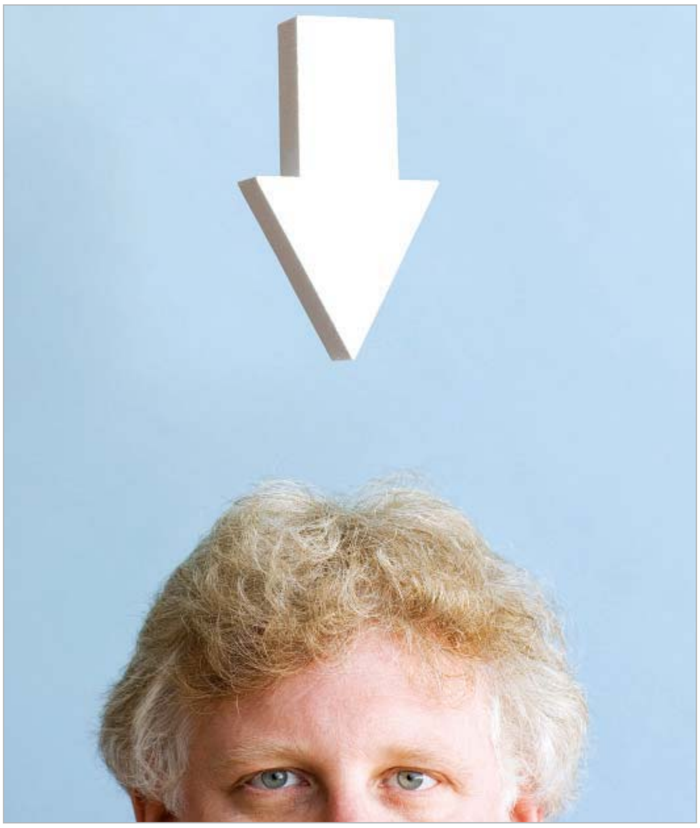

DAVID ADLER DREAMS OF A GOOGLE MAP FOR THE HUMAN BRAIN BY SALLY ADEE

Reverse Engineering The Brain

WHAT DO FRUIT-FLY brains have in common with microchips? That's not the setup for a bad joke; it's David Adler's life. Under Adler's ultrasophisticated electron beam microscopes, advanced microprocessors with transistors far smaller than red blood cells have been reduced to their wiring diagrams. Now the noggin of the humble *Drosophila melanogaster* is next, as Adler is being courted by researchers at a neurobiology wing of the Howard Hughes Medical Institute to help them reverse engineer the human brain. They're starting small, with the fruit fly.

TIMOTHY ARCHIBALD

Located in the green, rolling hills of Ashburn, in northern Virginia, the campus, known as Janelia Farm, has been described as a kind of Bell Labs for neurobiology. Its task is solving what Adler calls the most important question in science: How exactly does the human brain do what it does? Lots of people are trying to answer this question, and there's a growing impetus toward using high-definition brain scans to find out how the brain works.

"In a hundred years I'd like to know how human consciousness works," says Janelia director Gerry Rubin. "The 10- or 20-year goal is to understand the fruit-fly brain." It's this difference between consciousness and brain that has neuroscience researchers stymied. The simplest system stores and processes information the same way the most complex system does; a primitive computer from 1986 works a lot like a supercomputer. Similarly, Rubin suspects that the human brain and the fruit-fly brain are separated only by degrees of complexity: "Just because it's much more advanced doesn't mean the basic wiring rules are different." Right now, Janelia is working on a circuit diagram of the fruit-fly brain.

To that end, Rubin has stocked the Janelia campus with a collection of neuroscientists, biologists, physicists, engineers, and computer scientists. The process resembles that of reverse engineering a microprocessor. It starts with a full-scale, three-dimensional wiring diagram of the fly's brain, in which the density of neurons is substantially higher—"but not infinitely higher," insists Adler—than the wiring in a high-end IC. "If we can get a circuit diagram of the human brain," says Adler, "then we can understand what causes a lot of neurological disorders—depression, epilepsy, maybe even Alzheimer's."

Like an IC, the fruit-fly brain is subjected to logic and optical testing to derive its circuit diagram. With one approach, called neuronal electrophysiology, researchers can record the electrical activity of neurons. "But the fly brain is even more complicated than an integrated circuit," says engineer and group leader Eric Betzig. "With an IC, you know that every transistor fires the same—it's either on or off. But the neurons in the brain don't necessarily do that—they fire sometimes 20, sometimes 80, sometimes 100 percent." So in addition to logic testing, the researchers also need to do imaging, and that's where Adler and his amazing microscope come in.

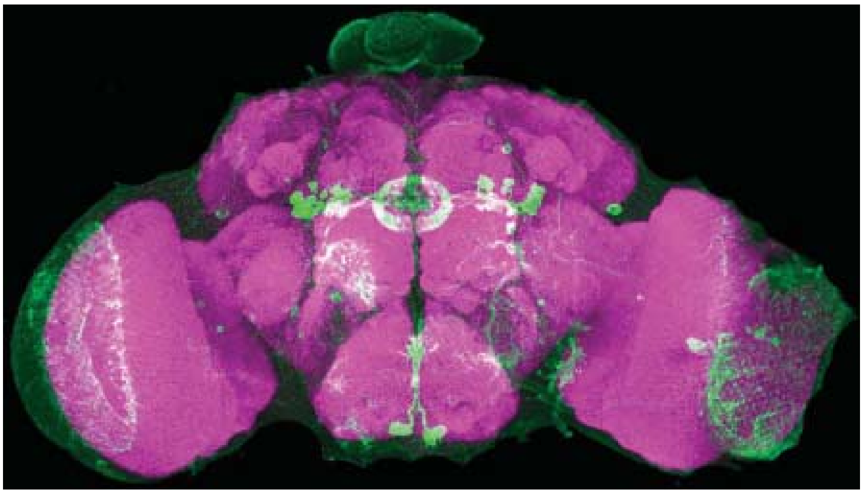

SLICED BRAINS: Confocal microscopy imagery of fluorescently labeled neurons is one of the tools used to develop circuit diagrams of the fruit-fly brain. PHOTO: JULIE SIMPSON & PHUONG CHUNG/JANELIA FARM

A standard scanning electron microscope (SEM) images at about 10 million pixels per second. For comparison, a high-definition TV screen runs at 30 million pixels per second. In 2005, the Pentagon gave funding to California-based KLA-Tencor Corp., where Adler was then working, to invent a microscope that could operate at 1 billion pixels per second to verify circuit patterns on defense chips. Shortly after that, Janelia lured Harald Hess, a former colleague of Adler's, to the campus to direct its applied physics and instrumentation group. When the Janelia team started looking into imaging, Hess called Adler. When he found out what Janelia was working on, Adler says, "it blew my mind." Hess wasn't interested in a microscope that could image at a paltry billion pixels per second. He wanted one that could process 10 billion per second.

To image the fruit-fly brain, the researchers use what they all refer to, gruesomely, as a "deli slicer"—the machine shaves 50-nanometer slices off the top of the infinitesimal fruit-fly brain "like slices of prosciutto," says Betzig. (The same technique is used to reverse engineer microchips.) Then an electron microscope takes images of the brain slices, and these images are stacked carefully to form a 3-D virtual wiring diagram.

Slice, image, slice, image. Easy, right? Wrong.

COMPARED WITH AN IC, even a tiny fruit-fly brain is a mess. One major bottleneck is the sample preparation: the brains must be sliced into perfectly even slivers before they're imaged. Right now, that slice-and-image routine takes a whopping 10 months. The real time-waster isn't the actual imaging—it's the time it takes for each slice of brain to settle into place. Any movement, however slight, will make that hard-won image blurry. The fly brain is only about 300 micrometers on a side, but imaging one, even at 10 billion pixels per second, would take a whole day. You're trying to image everything down to about 5 nm—about one-hundredth the size of what a regular lab microscope can resolve.

The storage requirements for the raw data alone are staggering: Adler estimates that scientists could rack up about a petabyte—that's 1000 terabytes—of data for every day of imaging. Bear in mind that 1000 terabytes is for one fruit fly, with its sorry speck of a brain, and the biggest hard drive you can buy from a commercial vendor today holds only one terabyte of data. To get any good data, you'd have to compare hundreds of fruit-fly brains. Imaging hundreds of them at the speed and resolution of Adler's technology would require a warehouse. "If nothing else," he says, "you're going to run out of space."

The brain is the ultimate micromachine. The fact that it's made out of meat is a red herring"

Anyone over 30 remembers when a gigabyte of storage in one place was laughably sci-fi. It won't be long before a 10-PB hard drive is as boring as today's 100-GB hard drives. But this project doesn't have as its goal merely collecting data; it is trying to establish the exact connections among the neurons and synapses of the tiny creature's brain. And therein lies the big challenge. Each slice holds billions of pixels, and once every slice has been imaged, scientists have to piece them all back together to generate a 3-D wiring diagram. Adler compares the scale of the undertaking to trying to put together a real-time traffic map of North America from high-resolution satellite photos. "Now imagine that the United States is paved coast-to-coast as densely as New York City," he says. At the resolution necessary to see individual synapses, the data glut is crippling. "You have to turn that data glut into a wiring diagram that doesn't take up 1000 hard drives," says Adler.

He's hoping that machine learning will compensate for the data glut. When American adventurer Steve Fossett disappeared in the Nevada desert last year, a virtual worldwide hunt ensued. People combed obsessively through Google Earth images for signs of the man and his plane. While telescopes and microscopes can image incredibly fine details, they still lack the all-important ability to interpret these images and throw away unnecessary data. Adler estimates that processing all the data from a 10-billion-pixel-per-second representation of the fruit-fly brain could take five years. So Mitya Chklovskii, Janelia's resident theoretical neurobiologist, is trying to teach his computers to discriminate neurons from synapses, and synapses from axons. A computer that could store an enormous image for 5 minutes while it decides which data is relevant, Janelia director Rubin says, is a far more elegant solution than a bigger hard drive. "This would solve the data problem," he says.

Let's say all the engineering problems can be solved in the next five or 10 years. Could researchers then actually reverse engineer the human brain, creating its functional duplicate in silicon? Would consciousness and all its attendant joy, pain, insanity, and genius be freed from biological containment? Adler sees no reason why not. "The brain is the ultimate micromachine," he insists. "The fact that it's made out of meat is a red herring."

His vision is a Google map of the human brain that incorporates not just

Janelia's circuit diagrams but also other work in neuroscience. Adler cites the work of Stanford neuroscientist Stephen Smith as "the first steps to finding the soul." At Harvard's Center for Brain Science, neuroscientist Jeff Lichtman mapped mouse neurons by "painting" them with fluorescent proteins. Rubin believes he'll live long enough to see an MRI-like device that measures function with such high-resolution output that neurons in fruit flies, mice, or even humans can be observed taking in and processing information in real time.

How would all these different systems work together to show us how the brain does what it does? With his 10-billion-pixel-per-second microscope, Adler is confident he'll be able to produce brain-topography images like Google's satellite views, resolving fine details in sharp focus. Smith's cartography, on the other hand, he compares with Google's map views, including street names. Rubin's fMRI data would be like real-time traffic data. Layering these different maps atop each other, says Adler, could lead to a hybrid comparable to a Google map.

Such a Google-mapped brain, Adler says, could do more than let us understand and cure disease: it could lead to a map of human consciousness. And he believes that understanding the wiring of the brain could lead to transformative technologies. What are memories, he asks, but rewired patterns in our brains? "If you can understand how memories are formed," he says, "you can create memories." Just as today's sophisticated circuit-editing tools can modify microchips after they've been manufactured and packaged, a brain-editing tool could perhaps one day modify the brain. Adler jokes about an application straight out of *Total Recall*: buying fond memories of a vacation instead of taking the actual trip.

IN THIS HEADY context, the leap from reverse engineering the human brain to building a thinking machine doesn't seem ridiculous. To Adler, the existence of human beings is proof enough that humans can be engineered. "When we study biology, we're just studying a different version of nanotechnology—only it's a more advanced nanotechnology." But he quickly qualifies that statement: silicon is the wrong material, he adds. The nanotechnology we use today is static; we can move electrons around but not atoms, which means the chip doesn't change when you use it. "We may not ever be able to get there using the silicon technology of moving electrons," he says. "But someone could come along tomorrow and invent a different way of making a circuit that's closer to what the brain does. Then, within 50 to 100 years, we'll have something that can do what the brain does."

But there's nothing like a little healthy competition to speed up this timetable. Janelia isn't the only player in the high-speed brain-imaging arena: both Harvard and the Max Planck Institute for Medical Research, in Heidelberg, Germany (where the 3-D SEM method of brain reconstruction was actually invented), are also working on the brain problem, and they compete heavily for milestones. The Harvard team may have solved the image-settling problem: they plan to adapt a conveyor-belt device used in the semiconductor industry as a continuously moving stage that allows an uninterrupted panoramic image, eliminating the need for time-wasting, steadying pauses.

Adler also consults for Harvard, helping its team push the limits of its existing SEMs by "supercharging them" to hit their full potential. Before a 10-billion-pixel-per-second microscope can be useful, he says, many other roadblocks have to be negotiated. So in the meantime, he takes these souped-up SEMs to the limits imposed on them by physics, not factory settings. That means, for instance, that a lab microscope with a default rate of 10 million pixels per second can jump to 100 million pixels per second after Adler is finished tweaking it.

Despite all the obstacles, the good news, Adler says, is that the fundamental physics of the superhigh-throughput electron microscope has been resolved. It's no longer a science problem, he says: now it's an engineering problem. Hess agrees. "Finding that one 65-nm shorted-wire defect in a Pentium chip and that one miswired neuron in a fruit-fly brain," he says, are fundamentally similar problems. They're counting on the inexorable climb of Moore's curve to aid them in their process. Rubin describes the phenomenon in terms of his Ph.D. work sequencing a single yeast gene. "Thirtyish years later, DNA-sequencing machines are at the point where students are doing 100 of my Ph.D.s per second," he says, laughing. "We're at millisecond data acquisition. These are the kinds of advances we'll need to make a map of the human mind."

In Rubin's mind, solving the fruit-fly brain is a 20-year problem. "After we solve this, I'd say we're one-fifth of the way to understanding the human mind." ❑

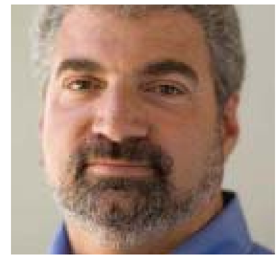

TIMOTHY ARCHIBALD

EXPERT VIEW:
Eric Hahn

WHO HE IS
Serial entrepreneur and early-stage investor who founded Collabra Software (sold to Netscape) and Lookout Software (sold to Microsoft) and backed Red Hat, Loudcloud, and Zimbra. CTO of Netscape during the browser wars.

SINGULARITY WILL OCCUR
Within 70 years

MACHINE CONSCIOUSNESS WILL OCCUR
"Yes, in that they eventually pass the Turing Test for 'Is it thinking?'"

MOORE'S LAW WILL CONTINUE FOR
30 more years

THOUGHTS
"I think that machine intelligence is one of the most exciting remaining 'great problems' left in computer science. For all its promise, however, it pales compared with the advances we could make in the next few decades in improving the health and education of the existing human intelligences already on the planet. I believe the first thing a tabula rasa intelligence (machine or otherwise) would conclude is that humans are very poor stewards of their own condition."

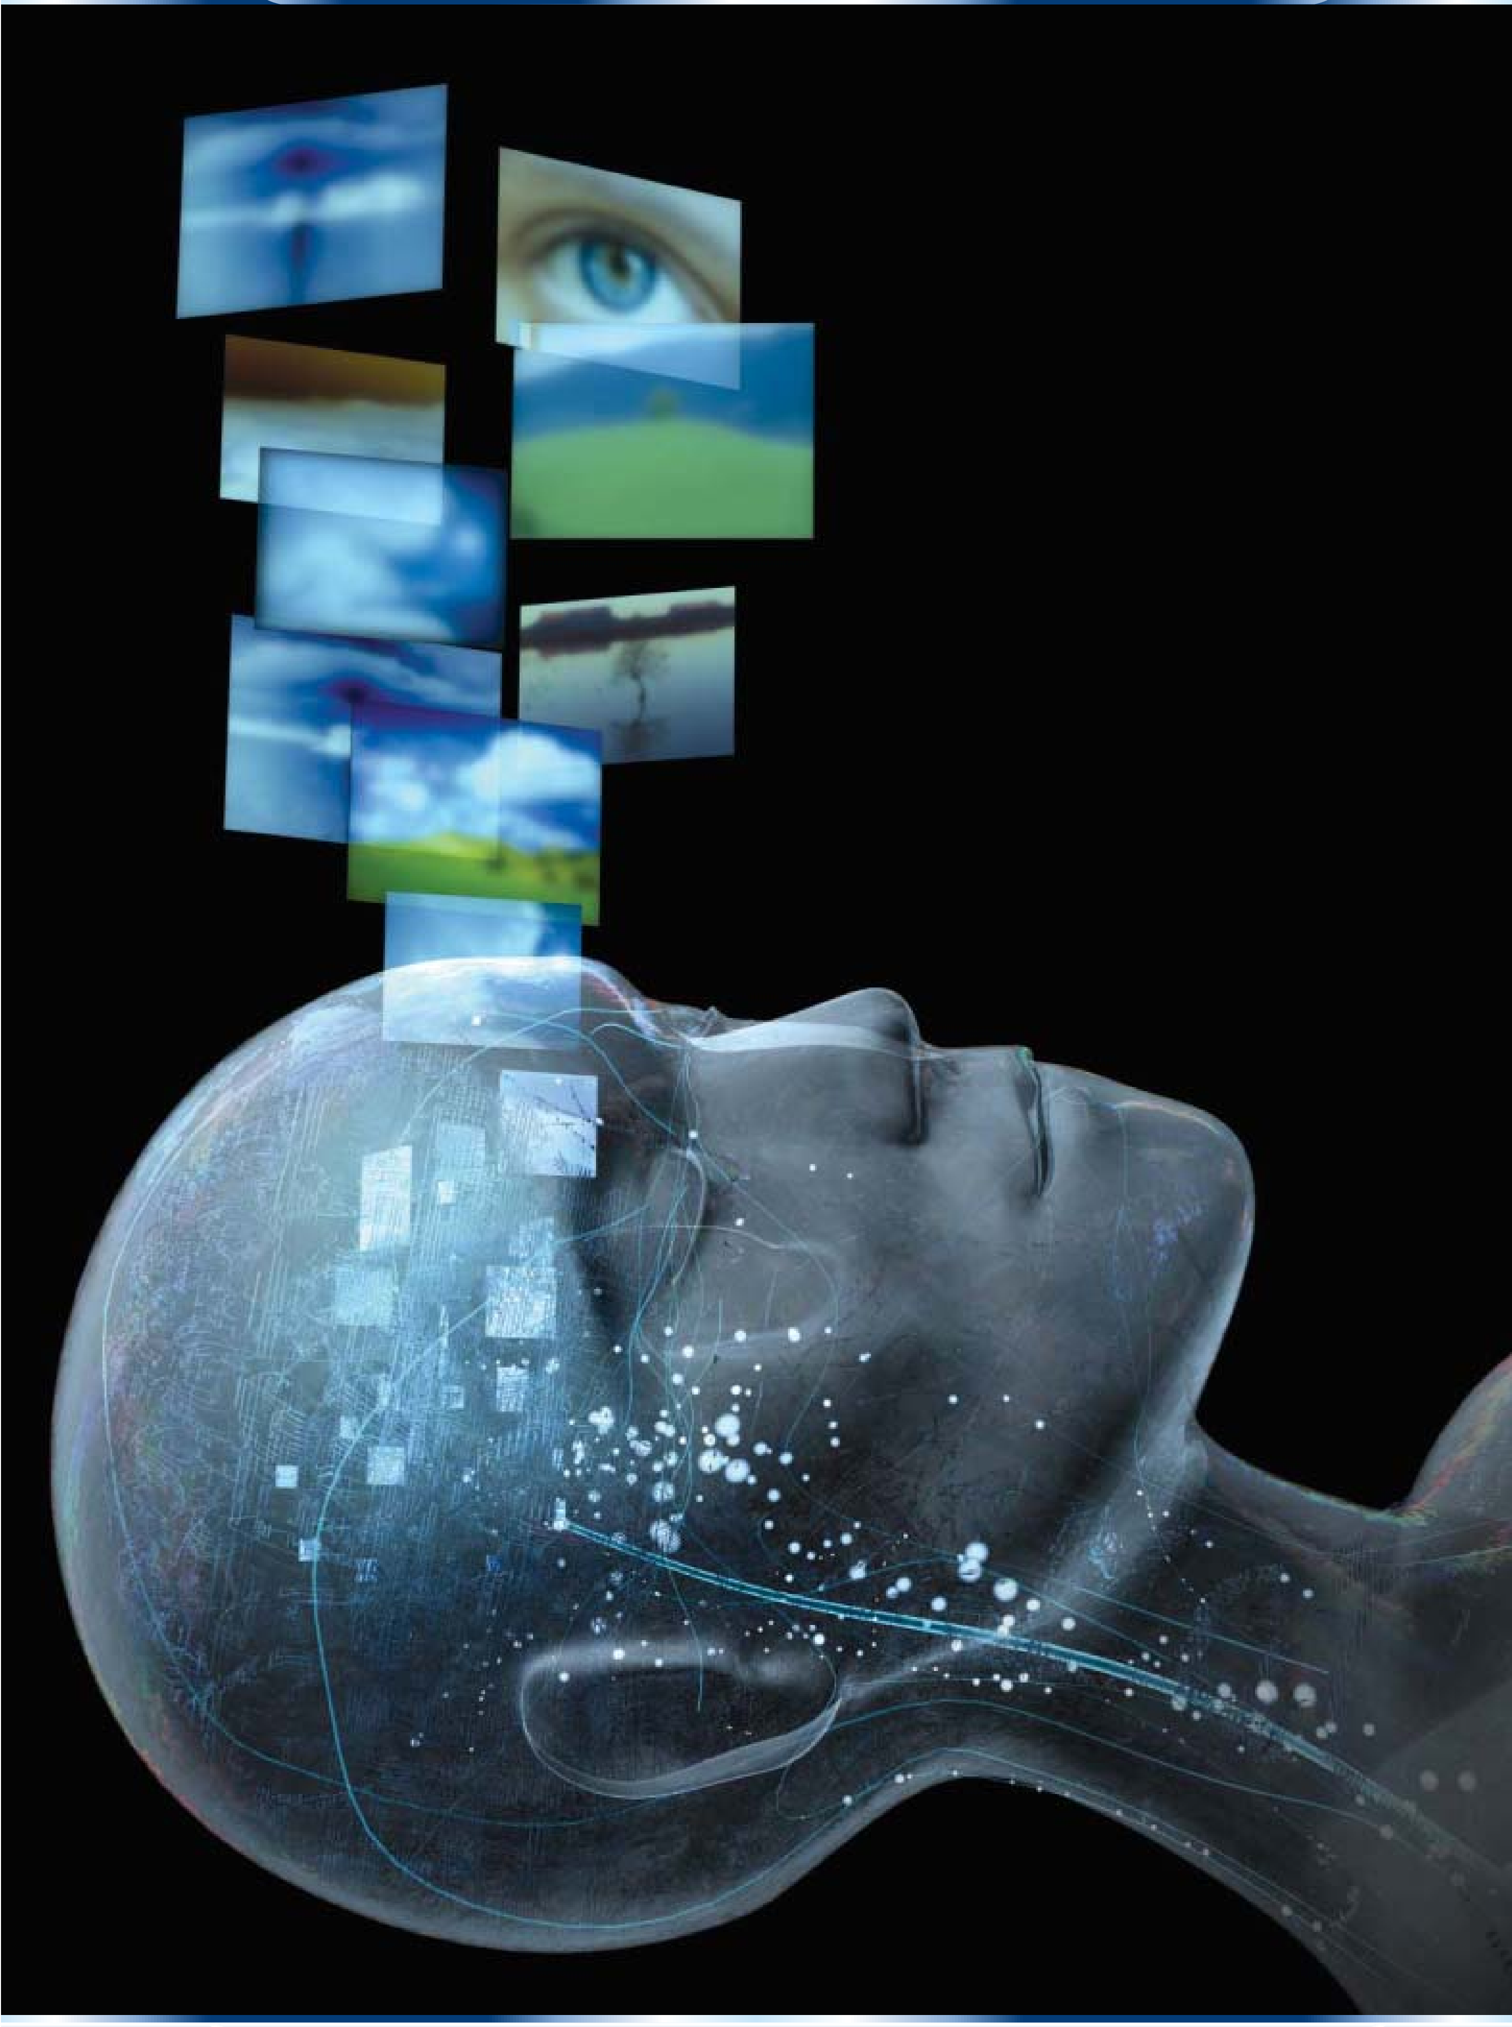

Can Machines Be Conscious?

YES—AND A NEW TURING TEST MIGHT PROVE IT

BY CHRISTOF KOCH AND GIULIO TONONI

WOULD YOU SELL your soul on eBay? Right now, of course, you can't. But in some quarters it is taken for granted that within a generation, human beings—including you, if you can hang on for another 30 years or so—will have an alternative to death: being a ghost in a machine. You'll be able to upload your mind—your thoughts, memories, and personality—to a computer. And once you've reduced your consciousness to patterns of electrons, others will be able to copy it, edit it, sell it, or pirate it. It might be bundled with other electronic minds. And, of course, it could be deleted.

That's quite a scenario, considering that at the moment, nobody really knows exactly what consciousness is. Pressed for a pithy definition, we might call it the ineffable and enigmatic inner life of the mind. But that hardly captures the whirl of thought and sensation that blossoms when you see a loved one after a long absence, hear an exquisite violin solo, or relish an incredible meal. Some of the most brilliant minds in human history have pondered consciousness, and after a few thousand years we still can't say for sure if it is an intangible phenomenon or maybe even a kind of substance different from matter. We know it arises in the brain, but we don't know how or where in the brain. We don't even know if it requires specialized brain cells (or neurons) or some sort of special circuit arrangement of them.

Nevertheless, some in the singularity crowd are confident that we are within a few decades of building a computer, a simulacrum, that can experience the color red, savor the smell of a rose, feel pain and pleasure, and fall in love. It might be a robot with a "body." Or it might just be software—a huge, ever-changing cloud of bits that inhabit an immensely complicated and elaborately constructed virtual domain.

We are among the few neuroscientists who have devoted a substantial part of their careers to studying consciousness. Our work has given us a unique perspective on what is arguably the most momentous issue in all of technology: whether consciousness will ever be artificially created.

We think it will—eventually. But perhaps not in the way that the most popular scenarios have envisioned it.

CONSCIOUSNESS IS PART of the natural world. It depends, we believe, only on mathe-

BRYAN CHRISTIE DESIGN

matics and logic and on the imperfectly known laws of physics, chemistry, and biology; it does not arise from some magical or otherworldly quality. That's good news, because it means there's no reason why consciousness can't be reproduced in a machine—in theory, anyway.

In humans and animals, we know that the specific content of any conscious experience—the deep blue of an alpine sky, say, or the fragrance of jasmine redolent in the night air—is furnished by parts of the cerebral cortex, the outer layer of gray matter associated with thought, action, and other higher brain functions. If a sector of the cortex is destroyed by stroke or some other calamity, the person will no longer be conscious of whatever aspect of the world that part of the brain represents. For instance, a person whose visual cortex is partially damaged may be unable to recognize faces, even though he can still see eyes, mouths, ears, and other discrete facial features. Consciousness can be lost entirely if injuries permanently damage most of the cerebral cortex, as seen in patients like Terri Schiavo, who suffered from persistent vegetative state. Lesions of the cortical white matter, containing the fibers through which parts of the brain communicate, also cause unconsciousness. And small lesions deep within the brain along the midline of the thalamus and the midbrain can inactivate the cerebral cortex and indirectly lead to a coma—and a lack of consciousness.

To be conscious also requires the cortex and thalamus—the corticothalamic system—to be constantly suffused in a bath of substances known as neuromodulators, which aid or inhibit the transmission of nerve impulses. Finally, whatever the mechanisms necessary for consciousness, we know they must exist in both cortical hemispheres independently.

A BETTER TURING TEST: Shown this frame from the cult classic *Repo Man* [top], a conscious machine should be able to home in on the key elements [bottom]—a man with a gun, another man with raised arms, bottles on shelves—and conclude that it depicts a liquor-store robbery.
PHOTO: EDGE CITY/UNIVERSAL/THE KOBAL COLLECTION

Much of what goes on in the brain has nothing to do with being conscious, however. Widespread damage to the cerebellum, the small structure at the base of the brain, has no effect on consciousness, despite the fact that more neurons reside there than in any other part of the brain. Neural activity obviously plays some essential role in consciousness but in itself is not enough to sustain a conscious state. We know that at the beginning of a deep sleep, consciousness fades, even though the neurons in the corticothalamic system continue to fire at a level of activity similar to that of quiet wakefulness.

Data from clinical studies and from basic research laboratories, made possible by the use of sophisticated instruments that detect and record neuronal activity, have given us a complex if still rudimentary understanding of the myriad processes that give rise to consciousness. We are still a *very* long way from being able to use this knowledge to build a conscious machine. Yet we can already take the first step in that long journey: we can list some aspects of consciousness that are not strictly necessary for building such an artifact.

Remarkably, consciousness does not seem to require many of the things we associate most deeply with being human: emotions, memory, self-reflection, language, sensing the world, and acting in it. Let's start with sensory input and motor output: *being conscious requires neither*. We humans are generally aware of what goes on around us and occasionally of what goes on within our own bodies. It's only natural to infer that consciousness is linked to our interaction with the world and with ourselves.

Yet when we dream, for instance, we are virtually disconnected from the environment—we acknowledge almost nothing of what happens around us, and our muscles are largely paralyzed. Nevertheless, we are conscious, sometimes vividly and grippingly so. This mental activity is reflected in electrical recordings of the dreaming brain showing that the corticothalamic system, intimately involved with sensory perception, continues to function more or less as it does in wakefulness.

Neurological evidence points to the same conclusion. People who have lost their eyesight can both imagine and dream in images, provided they had sight earlier in their lives. Patients with locked-in syndrome, which renders them almost completely paralyzed, are just as conscious as healthy subjects. Following a debilitating stroke, the French editor Jean-Dominique Bauby dictated his memoir, *The Diving Bell and the Butterfly*, by blinking his left eye. Stephen Hawking is a world-renowned physicist, best-selling author, and occasional guest star on "The Simpsons," despite being immobilized from a degenerative neurological disorder.

So although being conscious depends on brain activity, it does not require any interaction with the environment. Whether the *development* of consciousness requires such interactions in early childhood, though, is a different matter.

How about emotions? Does a conscious being need to feel and display them? No: *being conscious does not require emotion*. People who've suffered damage to the frontal area of the brain, for instance, may exhibit a flat, emotionless affect; they are as dispassionate about their own predicament as they are about the problems of people around them. But even though their behavior is impaired and their judgment may be unsound, they still experience the sights and sounds of the world much the way normal people do.

Primal emotions like anger, fear, surprise, and joy are useful and perhaps even essential for the survival of a conscious organism. Likewise, a conscious machine might rely on emotions to make choices and deal with the complexities of the world. But it could be just a cold, calculating engine—and yet still be conscious.

Psychologists argue that consciousness requires selective attention—that is, the ability to focus on a given object, thought, or activity. Some have even argued that consciousness *is* selective attention. After all, when you pay attention to something, you become conscious of that thing and its properties; when your attention shifts, the object fades from consciousness.

Nevertheless, recent evidence favors the idea that a person can consciously perceive an event or object without paying attention to it. When you're focused on a riveting movie, your surroundings aren't reduced to a tunnel. You may not hear the phone ringing or your spouse calling your name, but you remain aware of certain aspects of the world around you. And here's a surprise: the converse is also true. People can attend to events or objects—that is, their brains can preferentially process them—without consciously perceiving them. This fact suggests that *being conscious does not require attention*.

One experiment that supported this conclusion found that, as strange as it sounds, people could pay attention to an object that they never "saw." Test subjects were shown static images of male and female nudes in one eye and rapidly flashing colored squares in the other eye. The flashing color rendered the nudes invisible—the subjects couldn't even say where the nudes were in the image. Yet the psychologists showed that subjects nevertheless registered the unseen image if it was of the opposite sex.

What of memory? Most of us vividly remember our first kiss, our first car, or the images of the crumbling Twin Towers on 9/11. This kind of episodic memory would seem to be an integral part of consciousness. But the clinic tells us otherwise: *being conscious does not require either explicit or working memory*.

In 1953, an epileptic man known to the public only as H.M. had most of his hippocampus and neighboring regions on both sides of the brain surgically removed as an experimental treatment for his condition. From that day on, he couldn't acquire any new long-term memories—not of the nurses and doctors who treated him, his room at the hospital, or any unfamiliar well-wishers who dropped by. He could recall only events that happened before his surgery. Such impairments, though, didn't turn H.M. into a zombie. He is still alive today, and even if he can't remember events from one day to the next, he is without doubt conscious.

The same holds true for the sort of working memory you need to perform any number of daily activities—to dial a phone number you just looked up or measure out the correct amount of crushed thyme given in the cookbook you just consulted. This memory is called dynamic because it lasts only as long as neuronal circuits remain active. But as with long-term memory, you don't need it to be conscious.

Self-reflection is another human trait that seems deeply linked to consciousness. To assess consciousness, psychologists and other scientists often rely on verbal reports from their subjects. They ask questions like "What did you see?" To answer, a subject conjures up an image by "looking inside" and recalling whatever it was that was just viewed. So it is only natural to suggest that consciousness arises through your ability to reflect on your perception.

As it turns out, though, *being conscious does not require self-reflection*. When we become absorbed in some intense perceptual task—such as playing a fast-paced video game, swerving on a motorcycle through moving traffic, or running along a mountain trail—we are vividly conscious of the external world, without any need for reflection or introspection.

Neuroimaging studies suggest that we can be vividly conscious even when the front of the cerebral cortex, involved in judgment and self-representation, is relatively inactive. Patients with widespread injury to the front of the brain demonstrate serious deficits in their cognitive, executive, emotional, and planning abilities. But they appear to have nearly intact perceptual abilities.

Finally, *being conscious does not require language*. We humans affirm our consciousness through speech, describing and discussing our experiences with one another. So it's natural to think that speech and consciousness are inextricably linked. They're not. There are many patients who lose the ability to understand or use words and yet remain conscious. And infants, monkeys, dogs, and mice cannot speak, but they are conscious and can report their experiences in other ways.

So what about a machine? We're going to assume that a machine does not require anything to be conscious that a naturally evolved organism—you or me, for example—doesn't require. If that's the case, then, to be conscious a machine does not need to engage with its environment, nor does it need long-term memory or working memory; it does not require attention, self-reflection, language, or emotion. Those things may help the machine survive in the real world. But to simply have subjective experience—being pleased at the sight of wispy white clouds scurrying across a perfectly blue sky—those traits are probably not necessary.

So what *is* necessary? What are the essential properties of consciousness, those without which there is no experience whatsoever?

We think the answer to that question has to do with the amount of *integrated information* that an organism, or a machine, can generate. Let's say you are facing a blank screen that is alternately on or off, and you have been instructed to say "light" when the screen turns on and "dark" when it turns off. Next to you, a photodiode—one of the very simplest of machines—is set up to beep when the screen emits light and to stay silent when the screen is dark. The first problem that consciousness poses boils down to this: both you and the photodiode can differentiate between the screen being on or off, but while you can see light or dark, the photodiode does not consciously "see" anything. It merely responds to photons.

The key difference between you and the photodiode has to do with how much information is generated when the differentiation between light and dark is made. Information is classically defined as the reduction of uncertainty that occurs when one among many possible outcomes is chosen. So when the screen turns dark, the photodiode enters one of its two possible states; here, a state corresponds to one bit of information. But when *you* see the screen turn dark, you enter one out of a huge number of states: seeing a dark screen means you aren't seeing a blue, red, or green screen, the Statue of Liberty, a picture of your child's piano recital, or any of the other uncountable things that you have ever seen or could ever see. To you, "dark" means not just the opposite of light but also, and simultaneously, something different from colors, shapes, sounds, smells, or any mixture of the above.

So when you look at the dark screen, you rule out not just "light" but countless other possibilities. You don't think of the stupefying number of possibilities, of course, but their mere existence corresponds to a huge amount of information.

Conscious experience consists of more than just differentiating among many states, however. Consider an idealized 1-megapixel digital camera. Even if each photodiode in the imager were just binary, the number of different patterns that imager could record is $2^{1\,000\,000}$. Indeed, the camera could easily enter a different state for every frame from every movie that was or could ever be produced. It's a staggering amount of information. Yet the camera is obviously not conscious. Why not?

We think that the difference between you and the camera has to do with *integrated* information. The camera can indeed be in any one of an absurdly large number of different states. However, the 1-megapixel sensor chip isn't a single integrated system but rather a collection of one million individual, completely independent photodiodes, each with a repertoire of two states. And a million photodiodes are collectively no smarter than one photodiode.

By contrast, the repertoire of states available to you cannot be subdivided. You know this from experience: when you consciously see a certain image, you experience that image as an integrated whole. No matter how hard you try, you cannot divvy it up into smaller thumbprint images, and you cannot experience its colors independently of the shapes, or the left half of your field of view independently of the right half. Underlying this unity is a multitude of causal interactions among the relevant parts of your brain. And unlike chopping up the photodiodes in a camera sensor, disconnecting the elements of your brain that feed into consciousness would have profoundly detrimental effects.

TO BE CONSCIOUS, then, you need to be a single integrated entity with a large repertoire of states. Let's take this one step further: your *level* of consciousness has to do with how much integrated information you can generate. That's why you have a higher level of consciousness than a tree frog or a supercomputer.

It is possible to work out a theoretical framework for gauging how effective different neural architectures would be at generating integrated information and therefore attaining a conscious state. This framework, the integrated information theory of consciousness, or IIT, is grounded in the mathematics of information and complexity theory and provides a specific measure of the amount of integrated information generated by any system comprising interacting parts. We call that measure Φ and express it in bits. The larger the value of Φ, the larger the entity's conscious repertoire. (For students of information theory, Φ is an intrinsic property of the system, and so it is different from the Shannon information that can be sent through a channel.)

IIT suggests a way of assessing consciousness in a machine—a Turing Test for consciousness, if you will. Other attempts at gauging machine consciousness, or at least intelligence, have fallen short. Carrying on an engaging conversation in natural language or playing strategy games were at various times thought to be uniquely human attributes. Any machine that had those capabilities would also have a human intellect, researchers once thought. But subsequent events proved them wrong—computer programs such as the chatterbot ALICE and the chess-playing supercomputer Deep Blue, which famously bested Garry Kasparov in 1997, demonstrated that machines can display human-level performance in narrow tasks. Yet none of those inventions displayed evidence of consciousness.

Scientists have also proposed that displaying emotion, self-recognition, or purposeful behavior are suitable criteria for machine consciousness. However, as we mentioned earlier, there are people who are clearly conscious but do not exhibit those traits.

What, then, would be a better test for machine consciousness? According to IIT, consciousness implies the availability of a large repertoire of states belonging to a single integrated system. To be useful, those internal states should also be highly informative about the world.

One test would be to ask the machine to describe a scene in a way that efficiently differentiates the scene's key features from the immense range of other possible scenes. Humans are fantastically good at this: presented with a photo, a painting, or a frame from a movie, a normal adult can describe what's going on, no matter how bizarre or novel the image is.

Consider the following response to a particular image: "It's a robbery—there's a man holding a gun and pointing it at another man, maybe a store clerk." Asked to elaborate, the person could go on to say that it's probably in a liquor store, given the bottles on the shelves, and that it may be in the United States, given the English-language newspaper and signs. Note that the exercise here is not to spot as many details as one can but to discriminate the scene, as a whole, from countless others.

So this is how we can test for machine consciousness: show it a picture and ask it for a concise description [see photos, "A Better Turing Test"]. The machine should

> Consciousness does not seem to require many of the things we associate with being human

be able to extract the gist of the image (it's a liquor store) and what's happening (it's a robbery). The machine should also be able to describe which objects are in the picture and which are not (where's the getaway car?), as well as the spatial relationships among the objects (the robber is holding a gun) and the causal relationships (the other man is holding up his hands because the bad guy is pointing a gun at him).

The machine would have to do as well as any of us to be considered as conscious as we humans are—so that a human judge could not tell the difference—and not only for the robbery scene but for any and all other scenes presented to it.

No machine or program comes close to pulling off such a feat today. In fact, image understanding remains one of the great unsolved problems of artificial intelligence. Machine-vision algorithms do a reasonable job of recognizing ZIP codes on envelopes or signatures on checks and at picking out pedestrians in street scenes. But deviate slightly from these well-constrained tasks and the algorithms fail utterly.

Very soon, computer scientists will no doubt create a program that can automatically label thousands of common objects in an image—a person, a building, a gun. But that software will still be far from conscious. Unless the program is explicitly written to conclude that the combination of man, gun, building, and terrified customer implies "robbery," the program won't realize that something dangerous is going on. And even if it were so written, it might sound a false alarm if a 5-year-old boy walked into view holding a toy pistol. A sufficiently conscious machine would not make such a mistake.

WHAT IS THE best way to build a conscious machine? Two complementary strategies come to mind: either copying the mammalian brain or evolving a machine. Research groups worldwide are already pursuing both strategies, though not necessarily with the explicit goal of creating machine consciousness.

Though both of us work with detailed biophysical computer simulations of the cortex, we are not optimistic that modeling the brain will provide the insights needed to construct a conscious machine in the next few decades. Consider this sobering lesson: the roundworm *Caenorhabditis elegans* is a tiny creature whose brain has 302 nerve cells. Back in 1986, scientists used electron microscopy to painstakingly map its roughly 6000 chemical synapses and its complete wiring diagram. Yet more than two decades later, there is still no working model of how this minimal nervous system functions.

Now scale that up to a human brain with its 100 billion or so neurons and a couple hundred trillion synapses. Tracing all those synapses one by one is close to impossible, and it is not even clear whether it would be particularly useful, because the brain is astoundingly plastic, and the connection strengths of synapses are in constant flux. Simulating such a gigantic neural network model in the hope of seeing consciousness emerge, with millions of parameters whose values are only vaguely known, will not happen in the foreseeable future.

A more plausible alternative is to start with a suitably abstracted mammal-like architecture and evolve it into a conscious entity. Sony's robotic dog, Aibo, and its humanoid, Qrio, were rudimentary attempts; they operated under a large number of fixed but flexible rules. Those rules yielded some impressive, lifelike behavior—chasing balls, dancing, climbing stairs—but such robots have no chance of passing our consciousness test.

So let's try another tack. At MIT, computational neuroscientist Tomaso Poggio has shown that vision systems based on hierarchical, multilayered maps of neuronlike elements perform admirably at learning to categorize real-world images. In fact, they rival the performance of state-of-the-art machine-vision systems. Yet such systems are still very brittle. Move the test setup from cloudy New England to the brighter skies of Southern California and the system's performance suffers. To begin to approach human behavior, such systems must become vastly more robust; likewise, the range of what they can recognize must increase considerably to encompass essentially all possible scenes.

Contemplating how to build such a machine will inevitably shed light on scientists' understanding of our own consciousness. And just as we ourselves have evolved to experience and appreciate the infinite richness of the world, so too will we evolve constructs that share with us and other sentient animals the most ineffable, the most subjective of all features of life: consciousness itself. ❑

TO PROBE FURTHER For more on the integrated information theory of consciousness, go to http://spectrum.ieee.org/jun08/consciousmachines.

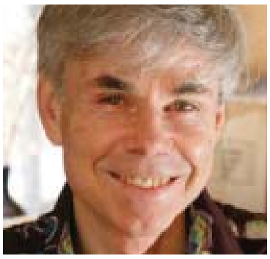

CHRIS MEYER/INDIANA UNIVERSITY

EXPERT VIEW:
Douglas Hofstadter

WHO HE IS
Pioneer in computer modeling of mental processes; director of the Center for Research on Concepts and Cognition at Indiana University, Bloomington; winner of the 1980 Pulitzer Prize for general nonfiction.

SINGULARITY WILL OCCUR
Someday in the distant future

MACHINE CONSCIOUSNESS WILL OCCUR
Yes

MOORE'S LAW WILL CONTINUE FOR
20 more years

THOUGHTS
"It might happen someday, but I think life and intelligence are far more complex than the current singularitarians seem to believe, so I doubt it will happen in the next couple of centuries. [The ramifications] will be enormous, since the highest form of sentient beings on the planet will no longer be human. Perhaps these machines—our 'children'—will be vaguely like us and will have culture similar to ours, but most likely not. In that case, we humans may well go the way of the dinosaurs."

Age of Invention

Technological optimists maintain that the impact of innovation on our lives is increasing, but the evidence goes the other way. The author's grandmother [see photo] lived from the 1880s through the 1960s and witnessed the adoption of electricity, phonographs, telephones, radio, television, airplanes, antibiotics, vacuum tubes, transistors, and the automobile. In 1924 she became one of the first in her neighborhood to own a car. The author contends that the inventions unveiled in his own lifetime have made a far smaller difference.

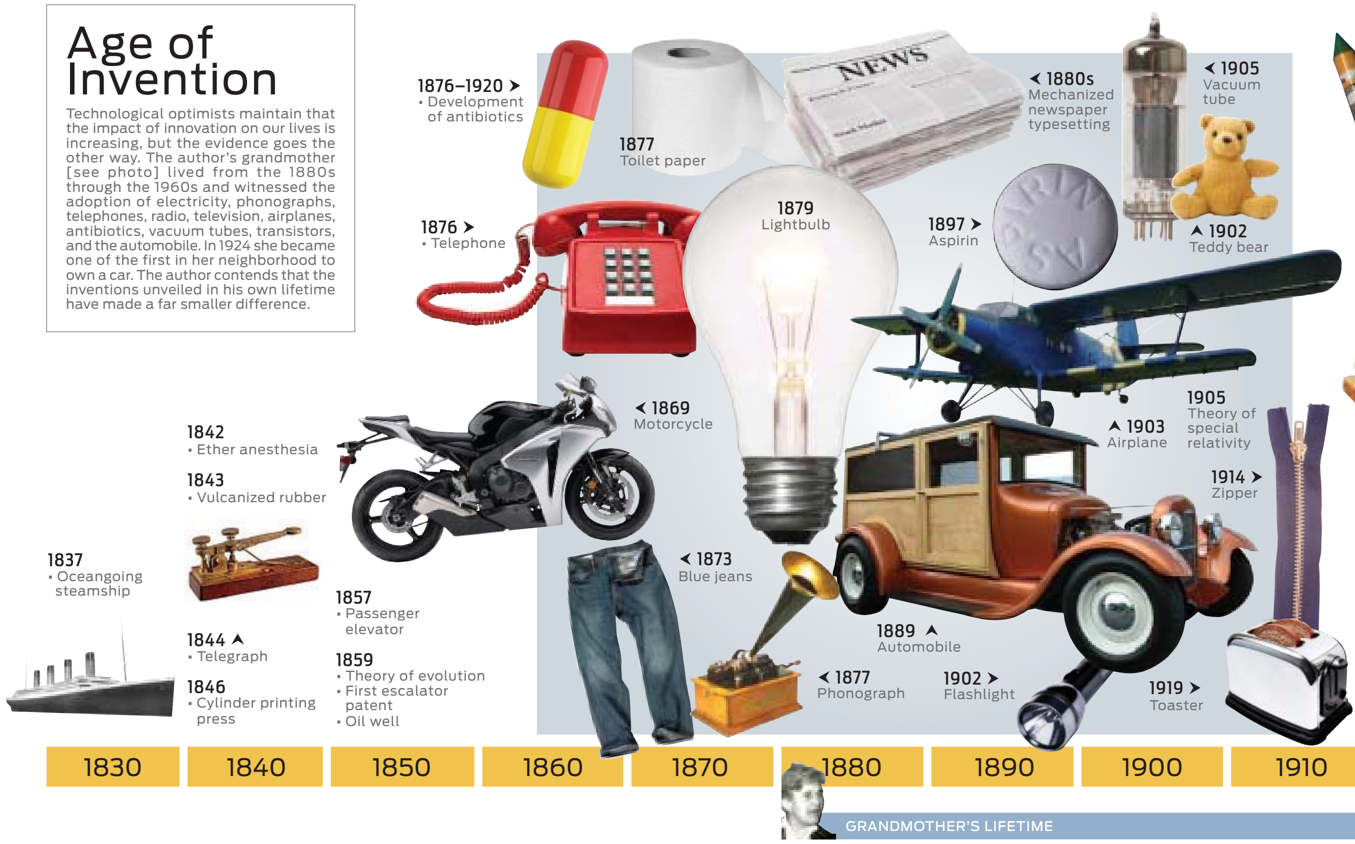

THE STORY OF THE SINGULARITY IS SWEEPING, DRAMATIC, SIMPLE—AND WRONG BY ALFRED NORDMANN

Singular Simplicity

TAKE THE IDEA of exponential technological growth, work it through to its logical conclusion, and there you have the singularity. Its bold incredibility pushes aside incredulity, as it challenges us to confront all the things we thought could never come true—the creation of superintelligent, conscious organisms, nanorobots that can swim in our bloodstreams and fix what ails us, and direct communication from mind to mind. And the pièce de résistance: a posthuman existence of disembodied uploaded minds, living on indefinitely without fear, sickness, or want in a virtual paradise ingeniously designed to delight, thrill, and stimulate.

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

PHOTO CREDITS: THE NATIONAL ARCHIVES; HONDA; NATIONAL MUSEUM OF AMERICAN HISTORY; RADIOMUSEUM KOELN; ESA; U.S. AIR FORCE; UNIVERSITY OF MARY WASHINGTON; NASA; NATIONAL INSTITUTES OF HEALTH; APPLE; ISTOCKPHOTO

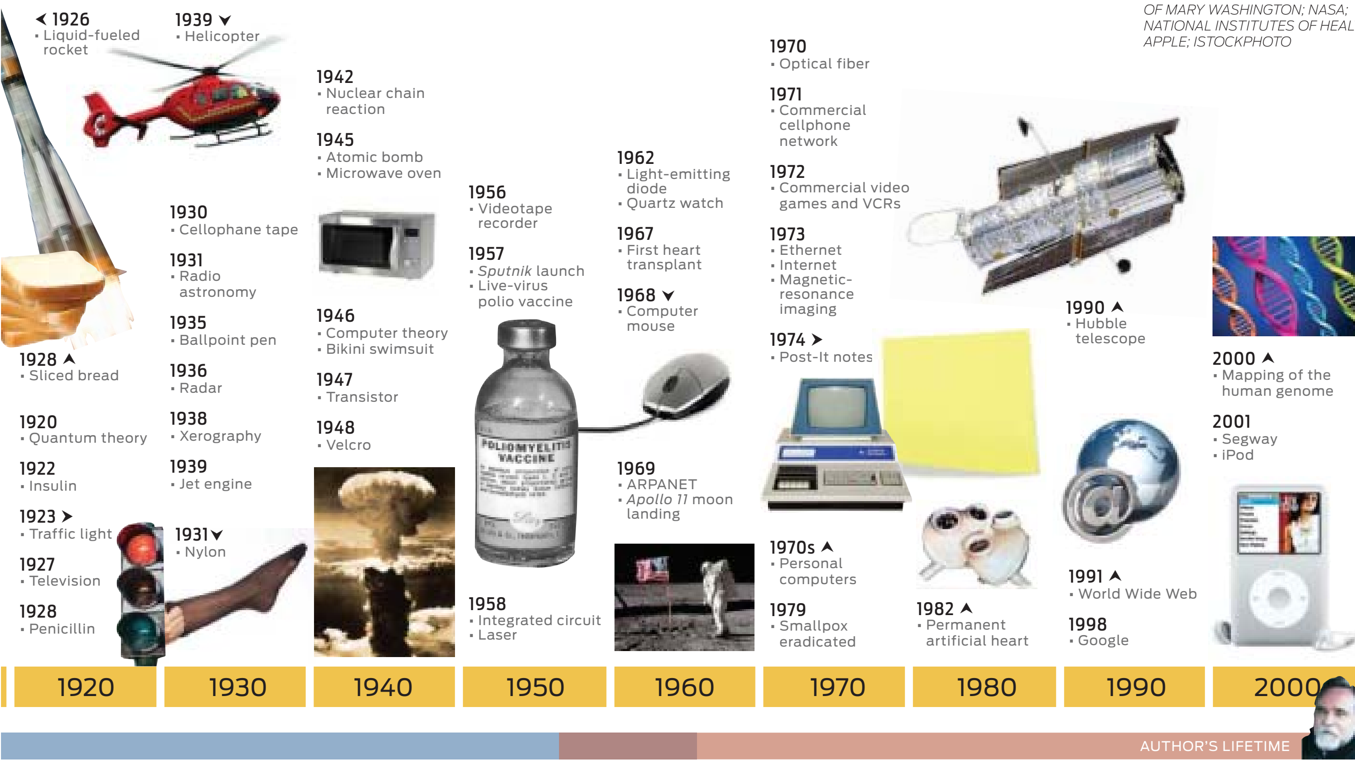

This vision argues that machines will become conscious and then perfect themselves, as described elsewhere in this issue. Yet for all its show of tough-minded audacity, the argument is shot through with sloppy reasoning, wishful thinking, and irresponsibility. Infatuated with statistics and seduced by the power of extrapolation, singularitarians abduct the moral imagination into a speculative no-man's-land. To be sure, they are hardly the first to spread fanciful technological prophecies, but among enthusiasts and doomsayers alike their proposition enjoys an inexplicable popularity. Perhaps the real question is how they have gotten away with it.

THE TROUBLE BEGINS with the singularitarians' assumption that technological advances have accelerated. I'd argue that I have seen less technological progress than my parents did, let alone my grandparents. Born in 1956, I can testify primarily to the development of the information age, fueled by the doubling of computing power every 18 to 24 months, as described by Moore's Law. The birth-control pill and other reproductive technologies have had an equally profound impact, on the culture if not the economy, but they are not developing at an accelerating speed. Beyond that, I saw men walk on the moon, with little to come of it, and I am surrounded by bio- and nanotechnologies that so far haven't affected my life at all. Medical research has developed treatments that make a difference in our lives, particularly at the end of them. But despite daily announcements of one breakthrough or another, morbidity and mortality from cancer and stroke continue practically unabated, even in developed countries.

Now consider the life of someone who was born in the 1880s and died in the 1960s—my grandmother, for instance. She witnessed the introduction of electric light and telephones, of automobiles and airplanes, the atomic bomb and nuclear power, vacuum electronics and semiconductor electronics, plastics and the computer, most vaccines and all antibiotics. All of those things mattered greatly in human terms, as can be seen in a single statistic: child mortality in industrialized countries dropped by 80 percent in those years.

So on what do intelligent people base the idea that technological progress is moving faster than ever before? It's simple: a chart of productivity from the dawn of humanity to the present day. It shows a line that inclines very gradually until around 1750, when it suddenly shoots almost straight up.

But that's hardly surprising. Since around 1750 the world has witnessed the spread of an economic system, by the name of capitalism, that is predicated on economic growth. And how the economy has grown since then! But surely the creation of new markets and the increasingly fine division of labor cannot be equated with technological progress, as every consumer knows.

Even if we were to accept, for the sake of argument, that technological innovation has truly accelerated, the line leading to the singularity would still be nothing but the simple-minded extrapolation of an existing pattern. Moore's Law has been remarkably successful at describing and predicting

MICHAEL CALLOPY/THE SKOLL FOUNDATION

EXPERT VIEW:
Jim Fruchterman

WHO HE IS
Founder and CEO of the Benetech Initiative, in Palo Alto, Calif., one of the first companies to focus on social entrepreneurship. Former rocket scientist and optical-character-recognition pioneer. Winner of a 2006 MacArthur Fellowship, the so-called genius grant.

SINGULARITY WILL OCCUR
Within 70 years

MACHINE CONSCIOUSNESS WILL OCCUR
Yes

MOORE'S LAW WILL CONTINUE FOR
30 more years

THOUGHTS
"I believe the singularity theory is plausible in that there will be a major shift in the rate of technology change. I am less convinced by projections of what it will mean to humans and humanity, such as human downloading in our lifetimes.

"Two things that rarely come up are the bug and algorithm questions. Douglas Hofstadter has more or less proved that perfect programs are not practically possible. And algorithms don't scale as nicely as processing power does: $n \log(n)$ is not our friend. A Linux system that needs rebooting only every three years is a modern technological marvel. But do you want to reboot your brain regularly?"

the development of semiconductors, in part because it has molded that development, ever since the semiconductor manufacturing industry adopted it as its road map and began spending vast sums on R&D to meet its requirements. Yet researchers and developers in the semiconductor industry have never denied that Moore's Law will finally come up against physical limits—indeed, many fear that the day of reckoning is nigh—whereas singularitarians happily extrapolate the law indefinitely into the future. And just as the semiconductor industry wonders nervously whether nanotechnology really can give Moore's Law another lease on life, singularitarians accept that this will occur as a given and then appropriate the exponential growth curve of Moore's Law not only to all the nano- and biotechnologies but to the cognitive sciences as well.

A typical example is the therapeutic development of brain-machine interfaces. In 2002, people were able to transmit 2 bits per minute to a computer. Four years later that figure had risen to 40 bits—that is, five letters—per minute. If this rate of progress continues, the argument goes, then by 2020 brain communication with computers will be as fast as speech. This isn't just the breathless cant of a true believer. The idea that an enhanced communication of thoughts will exceed speech can also be found in the 2002 report "Converging Technologies for Improving Human Performance," issued by the U.S. National Science Foundation and the Department of Commerce. It says that such methods "could complement verbal communication, sometimes replacing spoken language when speed is a priority or enhancing speech when needed to exploit maximum mental capabilities." Presumably, the singularity will be reached soon afterward, when transmission rates exceed the speed of thought itself, allowing the computer to transmit our thoughts before we think them.

This fantastic vision works only by ignoring the critical limit, which is the great concentration you have to muster to send the bits. It is a procedure far more tedious than speech. To ease that requirement—to make a brain-machine interface into a true mind-machine interface—we'd have to know a lot more than we do about the relation between specific thoughts and corresponding physical processes in the brain.

The seductive power of extrapolation has also been applied in ways less spectacular but no less foolish. The "lab on a chip" and other technologies for biochemical analysis have significantly increased the number of measurements—blood lipids, for instance—that can be obtained from a single drop of blood. It's a fine achieve-

Semiconductors and the Singularity

A COMMON BELIEF in singularity circles is that Moore's Law will not only continue indefinitely but will also apply to other areas of technology, including some yet to be invented. First expressed in 1965 by Intel cofounder Gordon E. Moore, Moore's Law describes the periodic doubling of the density of components on semiconductors. Sustained by enormous investments in R&D, the law still holds true, 43 years later, much to the surprise of Moore himself.

In his interesting book *The Singularity Is Near*, inventor and entrepreneur Ray Kurzweil accurately reviews the history of Moore's Law and describes some of the near-term challenges. Kurzweil then declares his belief that the industry is on the verge of a new paradigm that again increases the exponential growth rate of computational power, this time through the introduction of three-dimensional integrated circuits.

He may be right. That idea is very much in line with the views of industry leaders, who foresee not only stacked memories in the 2010-to-2013 time frame but also hybrid devices that incorporate logic, sensors, and microelectromechanical devices in the same 3-D package. An international research effort known as More Than Moore is now looking for ways to build extremely high-performance ICs for many different applications by combining existing technologies on silicon or in the same package. Nanowires in 3-D chips will help keep Moore's Law going by letting designers pack in transistors vertically as well as horizontally, in much the same way that skyscrapers in New York City pack in more people by housing them above and below one another.

Kurzweil goes on to speculate that the advanced 3-D molecular-computing devices of the future will more likely be grown in a test tube than made in a silicon fab. I find this forecast much less plausible. There is no such process that shows any promise of coming anywhere near silicon's manufacturing capabilities and economies of scale. There are research efforts that aim to build biological structures on silicon

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

ment, no doubt, but visionaries stretch the imagination when they assume that a second Moore's Law is about to produce astounding success stories and a transformation of all medical diagnostics.

Yet that assumption, which extrapolates an extrapolation—Moore's Law—to another field, is precisely what lies behind the now commonly expressed fear that increasing diagnostic powers are creating ethical problems in medicine. Physicians, we are told, will routinely inform patients of impending diseases for which they can offer no cure.

Yet in fact the path is very long from quicker blood analysis to instantaneous detection of the near certainty of a dread disease in a patient's future. A lab on a chip may provide mountains of data, but without great advances in many other fields—notably systems biology, pathology, and physiology—no one will be able to do much with it. Doctors already have more physiological information than they can profitably use.

Both examples of mindless extrapolation constitute wishful thinking. And in both cases, public debate is diverted from the real moral issues and quandaries that technology raises.

Rather than dream about how technology will soon effect an almost magical transformation of human life, societies need to debate the many real problems connected with technological changes that are already under way. These problems belong to the here and now.

WHY, THEN, ARE so many people captivated by the simple story of exponential growth that culminates in a life-altering singularity? Part of the appeal lies in simplicity itself, part in technological optimism—yet both of these tendencies are very old. What's new, though, is the changing role of technical expertise.

Plainly put, it is getting harder than ever to know whom to believe. Policy makers and members of the public have always had to put a degree of trust in experts. But now, when considering highly complex phenomena—in cellular processes, in chips containing billions of transistors, or in programs numbering hundreds of thousands of lines of code—even the experts must take a great deal on trust. That is because they have no choice but to study such phenomena using a cross-disciplinary approach.

These experts greet extraordinary claims made from within their own disciplines with skepticism and even indignation. But they can find it very hard to maintain such methodological vigilance in the hothouse atmosphere of a high-stakes collaboration in which researchers want desperately to believe that their own contributions can have wonderfully synergistic effects when combined with those of experts in other fields. And so, modest researchers recruit one another into immodest funding schemes.

The electronics engineer and the physiologist, the cognitive scientist and the physicist, the economist and the manufacturing specialist—all must take one another's statements on trust. They must trust in the contributions from other disciplines, trust in the power of visions to motivate the cooperation, trust in techniques and instruments that remain somewhat opaque to their users, trust in the trajectories of technical development.

Where trust has become a virtue even for scientists, there is little incentive to challenge outrageous claims or to hold singularitarians accountable. They describe the progressive realization of technical possibility, after all, and their story has a pleasant ring to it. Indeed, there is nothing wrong with the singular simplicity of the singularitarian myth—unless you have something against sloppy reasoning, wishful thinking, and an invitation to irresponsibility. ❑

MICROSOFT

By Bill Arnold, *chief scientist, ASML*

using photolithography to define the lattice on which the structures would grow, like vines clinging to a lattice. But these are in a very primitive state today.

Self-replication of superintelligent machines is another pillar of the singularity hypothesis. That kind of reproduction would require the machines to be able to fabricate superadvanced chips because those chips would form the basis of the machines' brains. It's possible that a lithography scanner itself could be the prototype of a self-replicating intelligent robot. I should add here, though, that I have no idea whether truly human-level artificial intelligence is possible, even if the device density and parallelism or the number of computations per second achieved by chips comes to rival those of the human brain.

Researchers will have to meet many, many challenges to keep lithography on the astounding curve it's been on for the past four decades. I have no doubt that they will. Lithography will continue to be the core technology that drives our future. ❑

EXPERT VIEW:
Gordon Bell

WHO HE IS
Principal researcher at Microsoft Research, Silicon Valley. Led the development of or helped design a long list of time-share computers and minicomputers at Digital Equipment Corp., including the PDP-6 and the VAX. A founder of Encore Computer; Ardent Computer; the Computer Museum, in Boston; and the Computer History Museum, in Mountain View, Calif.

SINGULARITY WILL OCCUR
Someday in the distant future

MACHINE CONSCIOUSNESS WILL OCCUR
No

MOORE'S LAW WILL CONTINUE FOR
20 years

THOUGHTS
"Singularity is that point in time when computing is able to know all human and natural-systems knowledge and exceed it in problem-solving capability with the diminished need for humankind as we know it. I basically support the notion, but I have trouble seeing the specific transitions or break points that let the exponential take over and move to the next transition. [If it does,] there'll be a hierarchy of machines versus having a separate race. [But] it is unlikely to happen, because the population will destroy itself before the technology singularity."

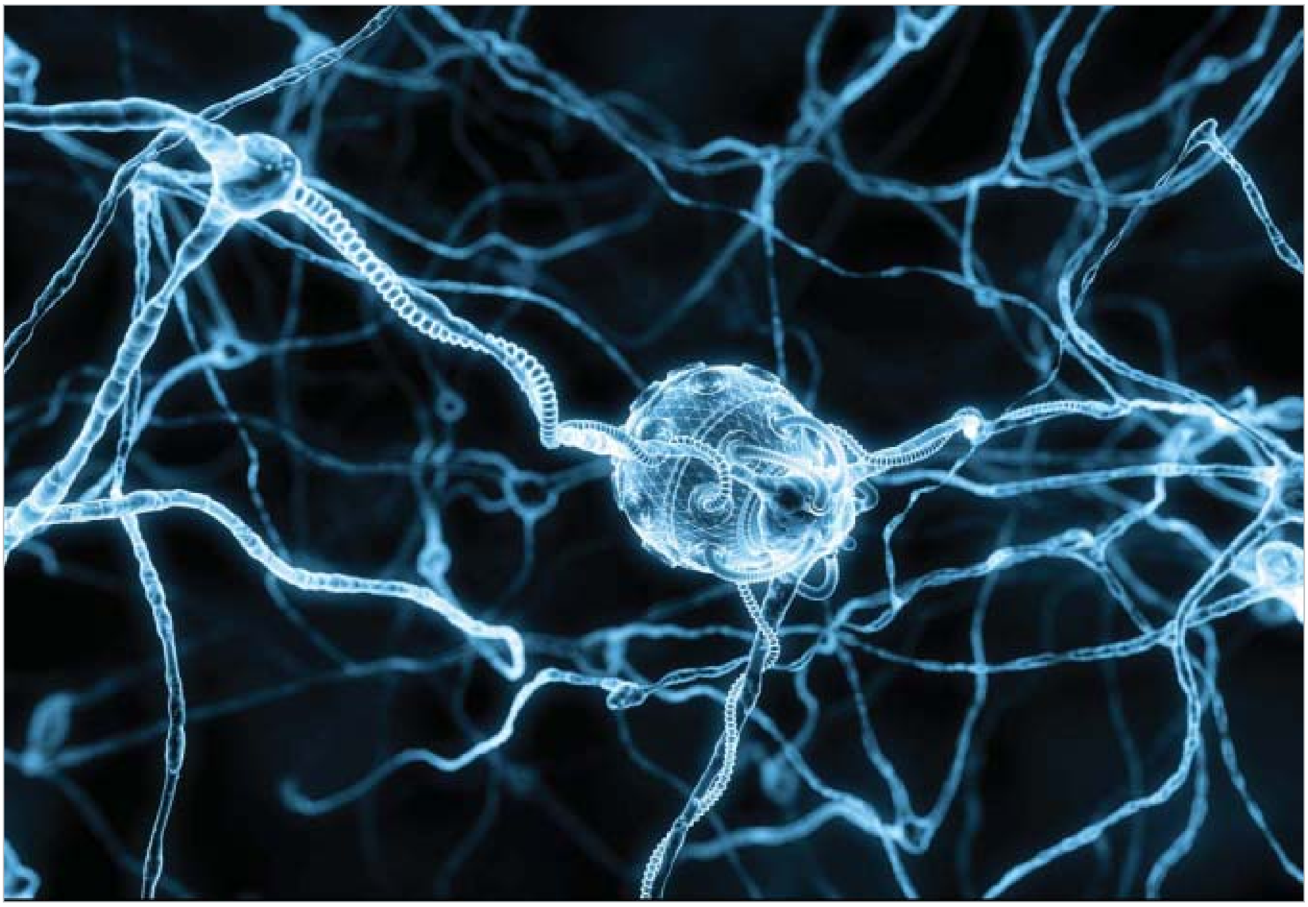

PHILIPPE VAN NEDERVELDE/E-SPACES/CG4TV

BIOLOGICAL NANOBOTS COULD REPAIR AND IMPROVE THE HUMAN BODY, BUT THEY'LL BE MORE BIO THAN BOT BY RICHARD A.L. JONES

Rupturing The Nanotech Rapture

HOW TO USHER humanity into an era of transhumanist bliss: first, end scarcity. Second, eradicate death. Third, eliminate the bungled mechanisms that introduce imperfections into the human body. The vehicle for accomplishing all three? Molecular nanotechnology—in essence, the reduction of all material things to the status of software.

To reduce the splendid complexity of our world to a list of instructions, a mere recipe, would involve harnessing the most basic components of life. Start with Earth's supply of atoms. Evolution, the laws of physics, and a big dose of chance have arranged those atoms into the objects and life-forms around us. If we could map the position and type of every atom in an object and also place atoms in specific positions, then in principle we could reproduce with absolute fidelity any material thing from its constituent parts. At a stroke, any material or artifact—a Stradivarius or a steak—could be available in abundance. We could build replacement body parts with capabilities that would hugely exceed their natural analogues. The economy, the environment, even what it means to be human, would be utterly transformed.

This vision holds wide currency among those anticipating a singularity, in which the creation of hyperintelligent, self-replicating machines triggers runaway technological advancement and economic growth, transforming human beings into cyborgs that are superhuman and maybe even immortal.

Some of these futurists are convinced that this renaissance is just a few decades away. But in academia and industry, nanotechnologists are working on a very different set of technologies. Many of these projects will almost certainly prove to be useful, lucrative, or even transformative, but none of them are likely to bring about the transhumanist rapture foreseen by singularitarians. Not in the next century, anyway.

It's not that the singularity vision is completely unrecognizable in today's work. It's just that the gulf between the two is a bit like the gap between traveling by horse and buggy and by interplanetary transport. The birth of nanotechnology is popularly taken to be 1989, when IBM Fellow Don Eigler used a scanning tunneling microscope to create the company's logo out of xenon atoms. Since then a whole field has emerged, based mainly on custom-engineered molecules that have gone into such consumer items as wrinkle-free clothes, more-effective sunscreens, and sturdier sports rackets.

However, it is a very long way indeed from a top-notch tennis racket to smart nanoscale robots capable of swarming in our bodies like infinitesimal guardian angels, recognizing and fixing damaged cells or DNA, and detecting, chasing, and destroying harmful viruses and bacteria. But the transhumanists underestimate the magnitude of that leap. They look beyond the manipulation of an atom or molecule with a scanning tunneling microscope and see swarms of manipulators that are themselves nanoscale. Under software control, these "nanofactories" would be able to arrange atoms in any pattern consistent with the laws of physics.

Rather than simply copying existing materials, the transhumanists dream of integrating into those materials almost unlimited functionality: state-of-the-art sensing and information processing could be built into the very fabric of our existence, accompanied by motors with astounding power density. Singularitarians anticipate that Moore's Law will run on indefinitely, giving us the immense computing power in tiny packages needed to control these nanofactories. These minuscule robots, or nanobots, need not be confined to protecting our bodies, either: if they can fix and purify, why not extend and enhance? Neural nanobots could allow a direct interface between our biological wetware and powerful computers with vast databases.

Maybe we could leave our bodies entirely. Only the need to preserve the contents of our memories and consciousness, our mental identities, ties us to them. Perhaps those nanobots will even be able to swim through our brains to read and upload our thoughts and memories, indeed entire personalities, to a powerful computer.

THIS EXPANSIVE VIEW of molecular nanotechnology owes as much to K. Eric Drexler as to anyone else. An MIT graduate and student of Marvin Minsky [see table, "Who's Who in the Singularity," in this issue], Drexler laid out his vision in the 1992 book *Nanosystems* (John Wiley & Sons). Those ideas have been picked up and expanded by other futurists over the past 16 years.

In his book, Drexler envisaged nanostructures built from the strongest and stiffest materials available, using the rational design principles of mechanical engineering. The fundamental building blocks of this paradigm are tiny, rigid cogs and gears, analogous to the plastic pieces of a Lego set. The gears would distribute power from nanoscale electric motors and be small enough to assist in the task of attaching molecules to one another. They would also process information. Drexler drew inspiration from a previous generation of computing devices, which used levers and gears rather than transistors, for his vision of ultrasmall mechanical computers.

Assuming that an object's structure could easily be reduced to its molecular blueprint, the first order of business is figuring out how to translate macroscale manufacturing methods into nanoscale manipulations. For example, let's say you wanted a new pancreas. Your first major challenge stems from the fact that a single human cell is composed of about 10^{14} atoms, and the pancreas you want has at least 80 billion cells, probably more. We could use a scanning tunneling microscope to position individual atoms with some precision, but to make a macroscopic object with it would take a very long time.

The theoretical solution, initially, was an idea known as exponential manufacturing. In its simplest form, this refers to a hypothetical nanoscale "assembler" that could construct objects on its own scale. For instance, it could make another assembler, and each assembler could go on to make more assemblers, resulting in a suite of assemblers that would combine forces to make a macroscopic object.

Setting aside the enormous challenges of creating and coordinating these nanoassemblers, some theorists have worried about a doomsday scenario known as the "gray goo" problem. Runaway replicators could voraciously consume resources to produce ever more stuff, a futuristic take on the old story of the sorcerer's apprentice. Not to worry, say Drexler and colleagues. In the latest vision of the nanofactory, the reproducing replicators give way to Henry Ford–style mass production, with endlessly repeated elementary operations on countless tiny production lines.

It's a seductive idea, seemingly validated by the workings of the cells of our own bodies. We're full of sophisticated nanoassemblers: delve into the inner workings of a typical cell and you'll find molecular motors that convert chemical energy into mechanical energy and membranes with active ion channels that sort molecules—two key tasks needed for basic nanoscale assembly. ATP synthase, for example, is an intricate cluster of proteins constituting a mechanism that makes adenosine triphosphate, the molecule that fuels the contraction of muscle cells and countless other cellular processes. Cell biology also exhibits software-controlled manufacturing, in the form of protein synthesis. The process starts with the ribosome, a remarkable molecular machine that can read information from a strand of messenger RNA and convert the code into a sequence of amino acids. The amino-acid sequence in turn defines the three-dimensional structure of a protein and its function. The ribosome fulfils the functions expected of an artificial assembler—proof that complex nanoassembly is possible.

If biology can produce a sophisticated nanotechnology based on soft materials like proteins and lipids, singularitarian

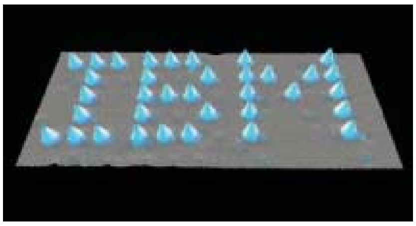

ATOMIC ARCHITECTURE: By working with xenon and nickel at very low temperatures, an IBM scientist was the first to use a tiny needle to position individual atoms on a surface. PHOTO: IBM CORP.

thinking goes, then how much more powerful our synthetic nanotechnology would be if we could use strong, stiff materials, like diamond. And if biology can produce working motors and assemblers using just the random selections of Darwinian evolution, how much more powerful the devices could be if they were rationally designed using all the insights we've learned from macroscopic engineering.

But that reasoning fails to take into account the physical environment in which cell biology takes place, which has nothing in common with the macroscopic world of bridges, engines, and transmissions. In the domain of the cell, water behaves like thick molasses, not the free-flowing liquid that we are familiar with. This is a world dominated by the fluctuations of constant Brownian motion, in which components are ceaselessly bombarded by fast-moving water molecules and flex and stretch randomly. The van der Waals force, which attracts molecules to one another, dominates, causing things in close proximity to stick together. Clingiest of all are protein molecules, whose stickiness underlies a number of undesirable phenomena, such as the rejection of medical implants. What's to protect a nanobot assailed by particles glomming onto its surface and clogging up its gears?

The watery nanoscale environment of cell biology seems so hostile to engineering that the fact that biology works at all is almost hard to believe. But biology does work—and very well at that. The lack of rigidity, excessive stickiness, and constant random motion may seem like huge obstacles to be worked around, but biology is aided by its own design principles, which have evolved over billions of years to exploit those characteristics. That brutal combination of strong surface forces and random Brownian motion in fact propels the self-assembly of sophisticated structures, such as the sculpting of intricately folded protein molecules. The cellular environment that at first seems annoying—filled with squishy objects and the chaotic banging around of particles—is essential in the operation of molecular motors, where a change in a protein molecule's shape provides the power stroke to convert chemical energy to mechanical energy.

In the end, rather than ratifying the "hard" nanomachine paradigm, cellular biology casts doubt on it. But even if that mechanical-engineering approach were

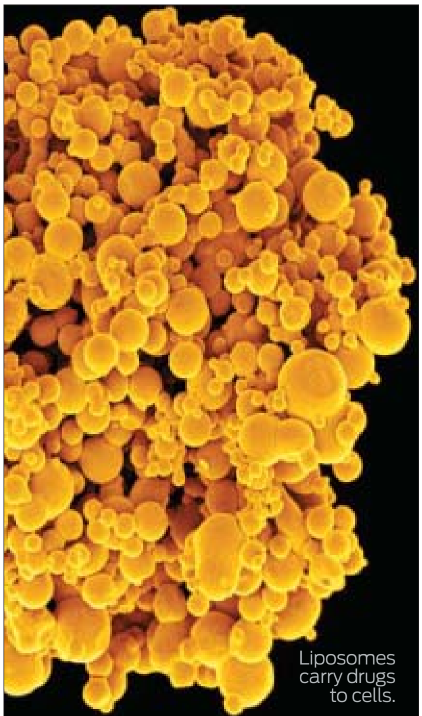

Liposomes carry drugs to cells.

The Real Nanobot

HAVING ABANDONED the strict mechanical approach to nanotechnology, what sorts of nanobots might we see instead?

IN DRUG DELIVERY Resembling viruses or bacteria more than shrunken submarines, self-assembling molecular bags called liposomes and polymersomes have been used in the clinic to wrap up drug molecules, such as anticancer compounds, for safe delivery to their targets and to introduce new genetic material into cells.

IN MEDICAL DIAGNOSTICS Crude prototypes exist of a nanolaboratory that could one day be able to read the state of a person's health by analyzing a single drop of blood for an array of disease markers.

IN OPTICS Nanotechnology researchers have sculpted intricate combinations of metals and plastics that cause light to bend in exotic ways. The aim is to make light bend backward, which could lead to invisibility cloaks and flawless lenses.

IN COMPUTING Engineers recently demonstrated a molecular machine that can do parallel processing. The nanodevice is composed of 17 molecules of a chemical called duroquinone. A scanning tunneling microscope probes a central molecule to change its state, which in turn shifts the other 16 molecules by manipulating the hydrogen bonds that connect them. —R.A.L.J.

to work in the body, there are several issues that, in my view, have been seriously underestimated by its proponents.

First, those building blocks—the cogs and gears made famous in countless simulations supporting the case for the singularity—have some questionable chemical properties. They are essentially molecular clusters with odd and special shapes, but it's far from clear that they represent stable arrangements of atoms that won't rearrange themselves spontaneously. These crystal lattices were designed using molecular modeling software, which works on the principle that if valences are satisfied and bonds aren't too distorted from their normal values, then the structures formed will be chemically stable. But this is a problematic assumption.

A regular crystal lattice is a 3-D arrangement of atoms or molecules with well-defined angles between the bonds that hold them together. To build a crystal lattice in a nonnatural shape—say, with a curved surface rather than with the flat faces characteristic of crystals—the natural distances and angles between atoms need to be distorted, severely straining those bonds. Modeling software might tell you that the bonds will hold. However, life has a way of confounding computer models. For example, if you try to make very small, spherical diamond crystals, a layer or two of carbon atoms at the surface will spontaneously rearrange themselves into a new form—not of diamond, but of graphite.

A second problem has to do with the power of surface forces and the high surface area anticipated for these nanobots. Researchers attempting to shrink existing microelectromechanical systems to the nanoscale have already discovered that the combination of friction and persistent sticking can be devastating. Nanobots are expected to operate at very high power densities, so even rather low values of friction may vaporize or burn up the minuscule machines. At the very least, this friction and sticking will play havoc with the machines' chemical stability.

Then there's the prospect of irreversible damage if reactive substances—such as water or oxygen—get caught up in a nanobot's exposed surfaces, upsetting the careful chemistry of each. To avoid those molecules, nanodevices will have to be fabricated in a fully controlled environment. No one yet knows how a medical nanobot would be protected once it is released into the warm, crowded turbu-

HYBRID MEDICAL ANIMATION/PHOTO RESEARCHERS lence of the body, perhaps the most heterogeneous environment imaginable.

Finally, there's the question of how an intricate arrangement of cogs and gears that depends on precision and rigidity to work will respond to thermal noise and Brownian bombardment at room temperature. The turbulence that nanobots will be subjected to will far exceed that inflicted on macroscopically engineered structures, and even the most rigid materials, like diamond, will bend and wobble in response. It would be like making a clock and its gears out of rubber, then watching it tumble around in a clothes dryer and wondering why it doesn't keep time. The bottom line is that we have no idea whether complex and rigid mechanical systems—even ones made from diamond—can survive in the nanoworld.

Put all these complications together and what they suggest, to me, is that the range of environments in which rigid nanomachines could operate, if they operate at all, would be quite limited. If, for example, such devices can function only at low temperatures and in a vacuum, their impact and economic importance would be virtually nil.

IN 15 YEARS of intense nanotechnology research, we have not even come close to experiencing the exponentially accelerating technological progress toward the goals set out by singularitarians. Impressive advances are emerging from the labs of real-world nanotechnologists, but these have little to do with the Drexlerian vision, which seems to be accumulating obstacles faster than it can overcome them. Given these facts, I can't take seriously the predictions that life-altering molecular nanotechnology will arrive within 15 or 20 years and hasten the arrival of a technological singularity before 2050.

Rather than try to defy or resist nature, I say we need to work with it. DNA itself can be used as a construction material. We can exploit its astounding properties of self-assembly to make programmed structures to execute new and beneficial functions [see sidebar, "The Real Nanobot"]. Chemists have already made nanoscale molecular shuttles and motors inspired directly by biology, with exciting applications in drug delivery and tissue engineering.

We will reap major medical advances by radically reengineering existing microorganisms, especially in nanodevices that perform integrated diagnosis and treatment of some disorders. But the timescales to reach the clinic are going to be long, and the goal of cell-by-cell repair is far, far beyond our incomplete grasp of biological complexity.

Much the same can be said about the singularitarian computers that are needed to generate a complete reading of a mental state and brain implants that seamlessly integrate our thought processes with a computer network. True, brain-interface systems have already been built. A state-of-the-art system can read about 128 neurons. So: 128 down, 20 billion or so to go.

Nonetheless, I'm an optimist. I think that in the near future we'll successfully apply nanotechnology to the most pressing social challenges, such as energy and the environment. For example, new polymer- and nanoparticle-based photovoltaics may soon lead to dramatic improvements in the price and production of solar cells.

What, then, of software-controlled matter? Complete control will remain an unattainable goal for generations to come. But some combination of self-assembly and directed assembly could very well lead to precisely built nanostructures that would manipulate the way light, matter, and electrons interact—an application of nanotechnology that's already leading to exciting new discoveries. We've barely scratched the surface of what we'll eventually be able to do with these custom-built nanostructures. It is altogether possible that we will finally harness the unfamiliar quantum effects of the nanoscale to implement true quantum computing and information processing. Here, I suspect, is the true killer application for the idea of software-controlled matter: devices that integrate electronics and optics, fully exploiting their quantum character in truly novel ways—a far cry from the minuscule diamond engines foreseen by the transhumanists.

We shouldn't abandon all of the more radical goals of nanotechnology, because they may instead be achieved ultimately by routes quite different from (and longer than) those foreseen by the proponents of molecular nanotechnology. Perhaps we should thank Drexler for alerting us to the general possibilities of nanotechnology, while recognizing that the trajectories of new technologies rarely run smoothly along the paths foreseen by their pioneers. ❑

JOSON

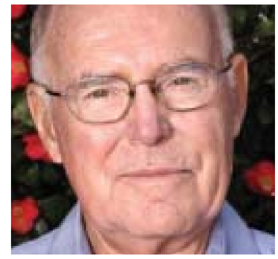

EXPERT VIEW:
Gordon E. Moore

WHO HE IS
Cofounder and chairman emeritus of Intel Corp., cofounder of Fairchild Semiconductor, winner of the 2008 IEEE Medal of Honor, chairman of the board of the Gordon and Betty Moore Foundation. Made the prediction about the increasing number of components on a semiconductor chip that came to be known as Moore's Law.

SINGULARITY WILL OCCUR
Never

THOUGHTS
"I am a skeptic. I don't believe this kind of thing is likely to happen, at least for a long time. And I don't know why I feel that way. The development of humans, what evolution has come up with, involves a lot more than just the intellectual capability. You can manipulate your fingers and other parts of your body. I don't see how machines are going to overcome that overall gap, to reach that level of complexity, even if we get them so they're intellectually more capable than humans."

"Although I do firmly believe that the brain is a machine, whether this machine is a computer is another question"

—Rodney Brooks

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

A POWERFUL ARTIFICIAL INTELLIGENCE WON'T SPRING FROM A SUDDEN TECHNOLOGICAL "BIG BANG"—IT'S ALREADY EVOLVING SYMBIOTICALLY WITH US

BY RODNEY BROOKS

I, Rodney Brooks, Am a Robot

I AM A MACHINE. So are you. ¶ Of all the hypotheses I've held during my 30-year career, this one in particular has been central to my research in robotics and artificial intelligence. I, you, our family, friends, and dogs—we all are machines. We are really sophisticated machines made up of billions and billions of biomolecules that interact according to well-defined, though not completely known, rules deriving from physics and chemistry. The biomolecular interactions taking place inside our heads give rise to our intellect, our feelings, our sense of self. ¶ Accepting this hypothesis opens up a remarkable possibility. If we really are machines and if—this is a big *if*—we learn the rules governing our brains, then in principle there's no reason why we shouldn't be able to replicate those rules in, say, silicon and steel. I believe our creation would exhibit genuine human-level intelligence, emotions, and even consciousness.

PREVIOUS SPREAD: BRYAN CHRISTIE DESIGN; THIS PAGE: SETH RESNICK/GETTY IMAGES

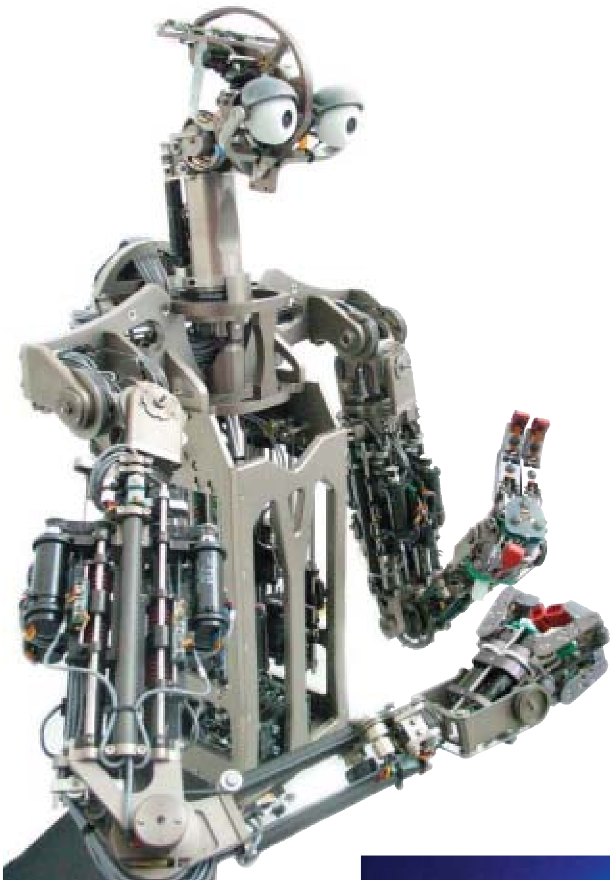

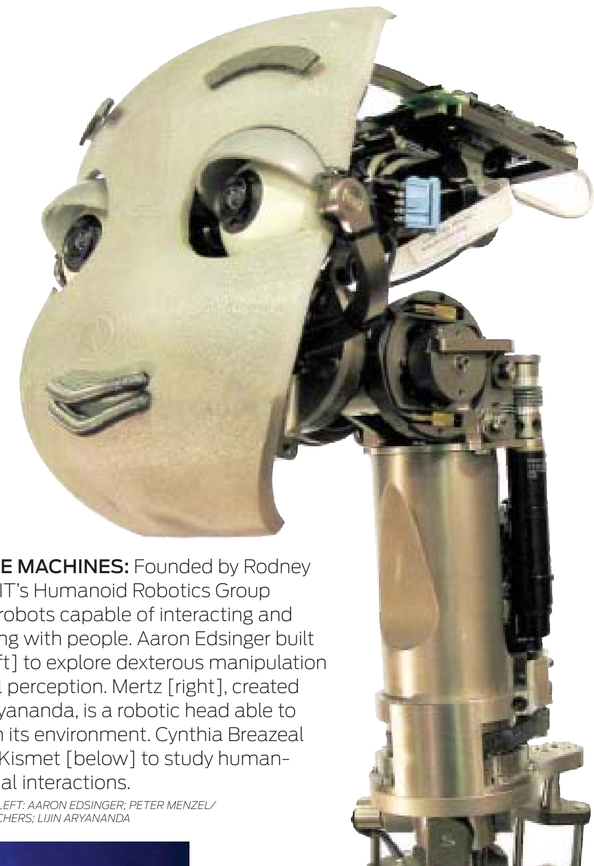

SOCIABLE MACHINES: Founded by Rodney Brooks, MIT's Humanoid Robotics Group develops robots capable of interacting and cooperating with people. Aaron Edsinger built Domo [left] to explore dexterous manipulation and visual perception. Mertz [right], created by Lijin Aryananda, is a robotic head able to learn from its environment. Cynthia Breazeal designed Kismet [below] to study human-robot social interactions.

PHOTOS, FROM LEFT: AARON EDSINGER; PETER MENZEL/PHOTO RESEARCHERS; LIJIN ARYANANDA

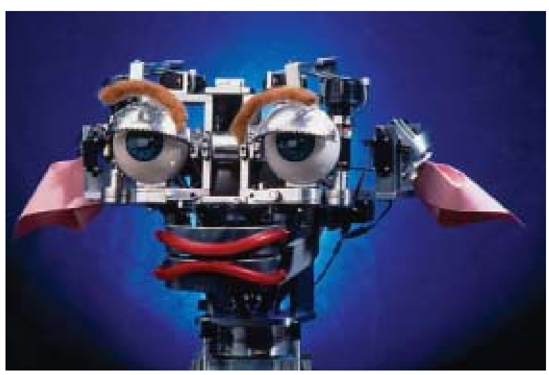

I'm far from alone in my conviction that one day we will create a human-level artificial intelligence, often called an artificial general intelligence, or AGI. But how and when we will get there, and what will happen after we do, are now the subjects of fierce debate in my circles. Some researchers believe that AGIs will undergo a positive-feedback self-enhancement until their comprehension of the universe far surpasses our own. Our world, those individuals say, will change in unfathomable ways after such superhuman intelligence comes into existence, an event they refer to as the singularity.

Perhaps the best known of the people proselytizing for this singularity—let's call them singularitarians—are acolytes of Raymond Kurzweil, author of *The Singularity Is Near: When Humans Transcend Biology* (Viking, 2005) and board member of the Singularity Institute for Artificial Intelligence, in Palo Alto, Calif. Kurzweil and his colleagues believe that this super AGI will be created either through ever-faster advances in artificial intelligence or by more biological means—"direct brain-computer interfaces, biological augmentation of the brain, genetic engineering, [and] ultrahigh-resolution scans of the brain followed by computer emulation" are some of their ideas. They don't believe this is centuries away; they think it will happen sometime in the next two or three decades.

What will the world look like then? Some singularitarians believe our world will become a kind of techno-utopia, with humans downloading their consciousnesses into machines to live a disembodied, after-death life. Others, however, anticipate a kind of technodamnation in which intelligent machines will be in conflict with humans, maybe waging war against us. The proponents of the singularity are technologically astute and as a rule do not appeal to technologies that would violate the laws of physics. They well understand the rates of progress in various technologies and how and why those rates of progress are changing. Their arguments are plausible, but plausibility is by no means certainty.

My own view is that things will unfold very differently. I do not claim that any specific assumption or extrapolation of theirs is faulty. Rather, I argue that an artificial intelligence could evolve in a much different way. In particular, I don't think there is going to be one single sudden technological "big bang" that springs an AGI into "life." Starting with the mildly intelligent systems we have today, machines will become gradually more intelligent, generation by generation. The singularity will be a period, not an event.

This period will encompass a time when we will invent, perfect, and deploy, in fits and starts, ever more capable systems, driven not by the imperative of the singularity itself but by the usual economic and sociological forces. Eventually, we will create truly artificial intelligences, with cognition and consciousness recognizably similar to our own. I have no idea how, exactly, this creation will come about. I also don't know when it will happen, although I strongly sus-

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

pect it won't happen before 2030, the year that some singularitarians predict.

But I expect the AGIs of the future—embodied, for example, as robots that will roam our homes and workplaces—to emerge gradually and symbiotically with our society. At the same time, we humans will transform ourselves. We will incorporate a wide range of advanced sensory devices and prosthetics to enhance our bodies. As our machines become more like us, we will become more like them.

And I'm an optimist. I believe we will all get along.

LIKE MANY AI researchers, I've always dreamed of building the ultimate intelligence. As a longtime fan of *Star Trek*, I have wanted to build Commander Data, a fully autonomous robot that we could work with as equals. Over the past 50 years, the field of artificial intelligence has made tremendous progress. Today you can find AI-based capabilities in things as varied as Internet search engines, voice-recognition software, adaptive fuel-injection modules, and stock-trading applications. But you can't engage in an interesting heart-to-power-source talk with any of them.

We have many very hard problems to solve before we can build anything that might qualify as an AGI. Many problems have become easier as computer power has reliably increased on its exponential and seemingly inexorable merry way. But we also need fundamental breakthroughs, which don't follow a schedule.

To appreciate the challenges ahead of us, first consider four basic capabilities that any true AGI would have to possess. I believe such capabilities are fundamental to our future work toward an AGI because they might have been the foundation for the emergence, through an evolutionary process, of higher levels of intelligence in human beings. I'll describe them in terms of what children can do.

■ The object-recognition capabilities of a 2-year-old child. A 2-year-old can observe a variety of objects of some type—different kinds of shoes, say—and successfully categorize them as shoes, even if he or she has never seen soccer cleats or suede oxfords. Today's best computer vision systems still make mistakes—both false positives and false negatives—that no child makes.

■ The language capabilities of a 4-year-old child. By age 4, children can engage in a dialogue using complete clauses and can handle irregularities, idiomatic expressions, a vast array of accents, noisy environments, incomplete utterances, and interjections, and they can even correct nonnative speakers, inferring what was really meant in an ungrammatical utterance and reformatting it. Most of these capabilities are still hard or impossible for computers.

■ The manual dexterity of a 6-year-old child. At 6 years old, children can grasp objects they have not seen before; manipulate flexible objects in tasks like tying shoelaces; pick up flat, thin objects like playing cards or pieces of paper from a tabletop; and manipulate unknown objects in their pockets or in a bag into which they can't see. Today's robots can at most do any one of these things for some very particular object.

■ The social understanding of an 8-year-old child. By the age of 8, a child can understand the difference between what he or she knows about a situation and what another person could have observed and therefore could know. The child has what is called a "theory of the mind" of the other person. For example, suppose a child sees her mother placing a chocolate bar inside a drawer. The mother walks away, and the child's brother comes and takes the chocolate. The child knows that in her mother's mind the chocolate is still in the drawer. This ability requires a level of perception across many domains that no AI system has at the moment.

But even if we solve these four problems using computers, I can't help wondering: What if there are some essential aspects of intelligence that we still do not understand and that do not lend themselves to computation? To a large extent we have all become computational bigots, believers that any problem can be solved with enough computing power. Although I do firmly believe that the brain is a machine, whether this machine is a computer is another question.

I recall that in centuries past the brain was considered a hydrodynamic machine. René Descartes could not believe that flowing liquids could produce thought, so he came up with a mind-body dualism, insisting that mental phenomena were nonphysical. When I was a child, the prevailing view was that the brain was a kind of telephone-switching network. When I was a teenager, it became an electronic computer, and later, a massively parallel digital computer. A few years ago someone asked me at a talk I was giving, "Isn't the brain just like the World Wide Web?"

We use these metaphors as the basis for our philosophical thinking and even let them pervade our understanding of what the brain truly does. None of our past metaphors for the brain has stood the test of time, and there is no reason to expect that the equivalence of current digital computing and the brain will survive. What we might need is a new conceptual framework: new ways of sorting out and piecing together the bits of knowledge we have about the brain.

Creating a machine capable of effectively performing the four capabilities above may take 10 years, or it may take 100. I really don't know. In 1966, some AI pioneers at MIT thought it would take three months—basically an undergraduate student working during the summer—to completely solve the problem of object recognition. The student failed. So did I in my Ph.D. project 15 years later. Maybe the field of AI will need several Einsteins to bring us closer to ultraintelligent machines. If you are one, get to work on your doctorate now.

I GREW UP IN a town in South Australia without much technology. In the late 1960s, as a teenager, I saw *2001: A Space Odyssey*, and it was a revelation. Like millions of others, I was enthralled by the soft-spoken computer villain HAL 9000 and wondered if we could one day get to that level of artificial intelligence. Today I believe the answer is yes. Nevertheless, in hindsight, I believe that HAL was missing a fundamental component: a body.

My early work on robotic insects showed me the importance of coupling AI systems to bodies. I spent a lot of time observing how those creatures crawled their way through complex obstacle courses, their gaits emerging from the interaction of their simple leg-control programs and the environment itself. After a decade building such insectoids, I decided to skip robotic lizards and cats and monkeys and jump straight to humanoids, to see what I could do there.

My students and I have learned a lot simply by putting people in front of a robot and asking them to talk to the machine. One of the most surprising things we've observed is that if a robot has a humanlike body, people will interact with it in a humanlike way. That's one of the reasons I came to believe that to

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

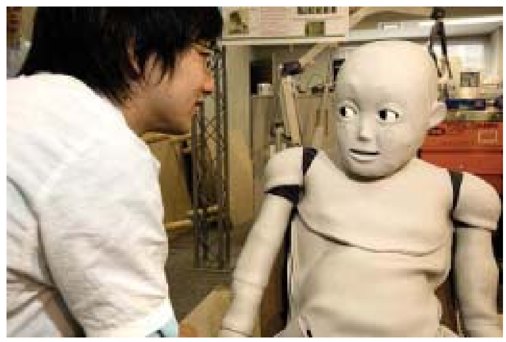

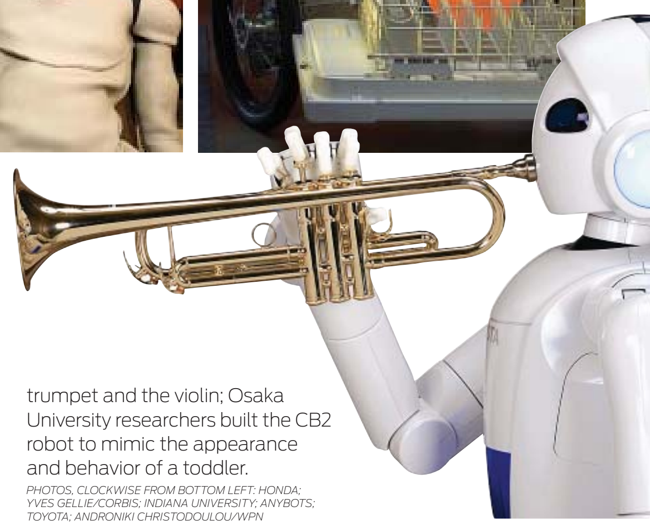

ANDROIDS ARISING: Clockwise from bottom left: Honda's Asimo can dance and climb stairs and has even conducted an orchestra; the Actroid female humanoid was developed by the Japanese firm Kokoro and Osaka University; Chinese roboticist Zou Ren Ti, of the Xi'an Chaoren Sculpture Research Institute, sits next to his android twin [right]; Anybots, in Mountain View, Calif., is developing tele-operated mechanical servants; Toyota's Partner robots include little droids that play the trumpet and the violin; Osaka University researchers built the CB2 robot to mimic the appearance and behavior of a toddler.

PHOTOS, CLOCKWISE FROM BOTTOM LEFT: HONDA; YVES GELLIE/CORBIS; INDIANA UNIVERSITY; ANYBOTS; TOYOTA; ANDRONIKI CHRISTODOULOU/WPN

build an AGI—and its predecessors—we'll need to give them a physical constitution.

At this point, I can guess what you're wondering. What will AGIs look like and when will they be here? What will it be like to interact with them? Will they be sociable, fun to be around?

I believe robots will have myriad sizes and shapes. Many will continue to be simply boxes on wheels. But I don't see why, by the middle of this century, we shouldn't have humanoid robots with agile legs and dexterous arms and hands. You won't have to read a manual or enter commands in C++ to operate them. You will just speak to them, tell them what to do. They will wander around our homes, offices, and factories, performing certain tasks as if they were people. Our environments were designed and built for our bodies, so it will be natural to have these human-shaped robots around to perform chores like taking out the garbage, cleaning the bathtub, and carrying groceries.

Will they have complex emotions, personalities, desires, and dreams? Some will, some won't. Emotions wouldn't be much of an asset for a bathtub-cleaning robot. But if the robot is reminding me to take my meds or helping me put the groceries away, I will want a little more personal interaction, with the sort of feedback that lets me know not just whether it's understanding me but how it's understanding me. So I believe the AGIs of the future will not only be able to act intelligently but also convey emotions, intentions, and free will.

So now the big question is: Will those emotions be real or just a very sophisticated simulation? Will they be the same kind of stuff as our own emotions? All I can give you is my hypothesis: the robot's emotional behavior can be seen as the real thing. We are made of biomolecules; the robots will be made of something else. Ultimately, the emotions created in each medium will be indistinguishable. In fact, one of my dreams is to develop a robot that people feel bad about switching off, as if they were extinguishing a life. As I wrote in my book *Flesh and Machines* (Pantheon, 2002), "We had better be careful just what we build, because we might end up liking them, and then we will be morally responsible for their well-being. Sort of like children."

Many of the advocates of the singularity appear to the more sober observers of technology to have a messianic fervor about their predictions, an unshakable faith in the certainty of their predicted future.

To an outsider, a lot of their convictions seem to have many commonalities with religious beliefs. Many singularitarians believe people will conquer death by downloading their consciousnesses into machines before their bodies give out. They expect this option will become available, conveniently enough, within their own lifetimes.

But for the sake of argument, let's accept all the wildest hopes of the singularitarians and accept that we will somehow construct an AGI in the next three or four decades. My view is that we will not live in the techno-utopia future that is so fervently hoped for. There are many possible alternative futures that fit within the themes of the singularity but are very different in their outcomes.

One scenario often considered by singularitarians, and Hollywood, too, is that an AGI might emerge spontaneously on a large computer network. But perhaps such an AGI won't have quite the relationship with humans that the singularitarians expect. The AGI may not know about us, and we may not know about it.

In fact, maybe some kind of AGI already exists on the Google servers, probably the single biggest network of computers on our planet, and we aren't aware of it. So at the 2007 Singularity Summit, I asked Peter Norvig, Google's chief scientist, if the company had noticed any unexpected emergent properties in its network—not full-blown intelligence, but any unexpected emergent property. He replied that they had not seen anything like that. I suspect we are a long, long way from consciousness unexpectedly showing up in the Google network. (Unless it is already there and cleverly concealing its tracks!)

Here's another scenario: the AGI might know about us and we know about it, but it might not care about us at all. Think of chipmunks. You see them wandering around your garden as you look out the window at breakfast, but you certainly do not know them as individuals and probably do not give much thought to which ones survive the winter. To an AGI, we may be nothing more than chipmunks.

From there it's only a short step to the question I'm asked over and over again: Will machines become smarter than us and decide to take over?

I don't think so. To begin with, there will be no "us" for them to take over from. We, human beings, are already starting to change ourselves from purely biological entities into mixtures of biology and technology. My prediction is that we are more likely to see a merger of ourselves and our robots before we see a standalone superhuman intelligence.

Our merger with machines is already happening. We replace hips and other parts of our bodies with titanium and steel parts. More than 50 000 people have tiny computers surgically implanted in their heads with direct neural connections to their cochleas to enable them to hear. In the testing stage, there are retina microchips to restore vision and motor implants to give quadriplegics the ability to control computers with thought. Robotic prosthetic legs, arms, and hands are becoming more sophisticated. I don't think I'll live long enough to get a wireless Internet brain implant, but my kids or their kids might.

And then there are other things still further out, such as drugs and genetic and neural therapies to enhance our senses and strength. While we become more robotic, our robots will become more biological, with parts made of artificial and yet organic materials. In the future, we might share some parts with our robots.

We need not fear our machines because we, as human-machines, will always be a step ahead of them, the machine-machines, because we will adopt the new technologies used to build those machines right into our own heads and bodies. We're going to build our robots incrementally, one after the other, and we're going to decide the things we like having in our robots—humility, empathy, and patience—and things we don't, like megalomania, unrestrained ambition, and arrogance. By being careful about what we instill in our machines, we simply won't create the specific conditions necessary for a runaway, self-perpetuating artificial-intelligence explosion that runs beyond our control and leaves us in the dust.

When we look back at what we are calling the singularity, we will see it not as a singular event but as an extended transformation. The singularity will be a period in which a collection of technologies will emerge, mature, and enter our environments and bodies. There will be a brave new world of augmented people, which will help us prepare for a brave new world of AGIs. We will still have our emotions, intelligence, and consciousness.

And the machines will have them too. ❑

RICK SMOLAN

EXPERT VIEW:
Esther Dyson

WHO SHE IS:
Commentator and evangelist for emerging technologies, investor and board member for start-ups; currently focused on health care, genetics, private aviation, and commercial space. Ran PC Forum conference until 2007; currently hosts the annual Flight School conference.

THOUGHTS
"The singularity I'm interested in will come from biology rather than machines. We won't be building things; we'll be growing and cultivating them, and then they will grow on their own."

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

HINTS OF THE SINGULARITY'S APPROACH CAN BE FOUND IN THE ARGUMENTS OF ITS CRITICS

BY VERNOR VINGE

SIGNS OF THE SINGULARITY

I think it's likely that with technology we can in the fairly near future create or become creatures of more than human intelligence. Such a technological singularity would revolutionize our world, ushering in a posthuman epoch. If it were to happen a million years from now, no big deal. So what do I mean by "fairly near" future? In my 1993 essay, "The Coming Technological Singularity," I said I'd be surprised if the singularity had not happened by 2030. I'll stand by that claim, assuming we avoid the showstopping catastrophes—things like nuclear war, superplagues, climate crash—that we properly spend our anxiety upon.

BRYAN CHRISTIE DESIGN

SPECTRUM Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page qMags

In that event, I expect the singularity will come as some combination of the following:

- **The AI Scenario:** We create superhuman artificial intelligence (AI) in computers.

- **The IA Scenario:** We enhance human intelligence through human-to-computer interfaces—that is, we achieve intelligence amplification (IA).

- **The Biomedical Scenario:** We directly increase our intelligence by improving the neurological operation of our brains.

- **The Internet Scenario:** Humanity, its networks, computers, and databases become sufficiently effective to be considered a superhuman being.

- **The Digital Gaia Scenario:** The network of embedded microprocessors becomes sufficiently effective to be considered a superhuman being.

The essays in this issue of *IEEE Spectrum* use similar definitions for the technological singularity but variously rate the notion from likely to totally bogus. I'm going to respond to arguments made in these essays and also mine them for signs of the oncoming singularity that we might track in the future.

Philosopher Alfred Nordmann criticizes the extrapolations used to argue for the singularity. Using trends for outright forecasting is asking for embarrassment. And yet there are a couple of trends that at least raise the possibility of the technological singularity. The first is a very long-term trend, namely Life's tendency, across aeons, toward greater complexity. Some people see this as unstoppable progress toward betterment. Alas, one of the great insights of 20th-century natural science is that Nature can be the harshest of masters. What we call progress can fail. Still, in the absence of a truly terminal event (say, a nearby gamma-ray burst or another collision such as made the moon), the trend has muddled along in the direction we call forward. From the beginning, Life has had the ability to adapt for survival via natural selection of heritable traits. That computational scheme brought Life a long way, resulting in creatures that could reason about survival problems. With the advent of humankind, Life had a means of solving many problems much faster than natural selection.

In the last few thousand years, humans have begun the next step, creating tools to support cognitive function. For example, writing is an off-loading of memory function. We're building tools—computers, networks, database systems—that can speed up the processes of problem solving and adaptation. It's not surprising that some technology enthusiasts have started talking about possible consequences. Depending on our inventiveness—and our artifacts' inventiveness—there is the possibility of a transformation comparable to the rise of human intelligence in the biological world. Even if the singularity does not happen, we are going to have to put up with singularity enthusiasms for a long time.

Get used to it.

In recent decades, the enthusiasts have been encouraged by an enabling trend: the exponential improvement in computer hardware as described by Moore's Law, according to which the number of transistors per integrated circuit doubles about every two years. At its heart, Moore's Law is about inventions that exploit one extremely durable trick: optical lithography to precisely and rapidly emplace enormous numbers of small components. If the economic demand for improved hardware continues, it looks like Moore's Law can continue for some time—though eventually we'll need novel component technology (perhaps carbon nanotubes) and some new method of high-speed emplacement (perhaps self-assembly). But what about that economic demand? Here is the remarkable thing about Moore's Law: it enables improvement in communications, embedded logic, information storage, planning, and design—that is, in areas that are directly or indirectly important to almost all enterprise. As long as the software people can successfully exploit Moore's Law, the demand for this progress should continue.

Roboticist Hans Moravec may have been the first to draw a numerical connection between computer hardware trends and artificial intelligence. Writing in 1988, Moravec took his estimate of the raw computational power of the brain together with the rate of improvement in computer power and projected that by 2010 computer hardware would be available to support roughly human levels of performance. There are a number of reasonable objections to this line of argument. One objection is that Moravec may have radically underestimated the computational power of neurons. But even if his estimate is a few orders of magnitude too low, that will only delay the transition by a decade or two—assuming that Moore's Law holds.

Another roboticist, Rodney Brooks, suggests in this issue that computation may not even be the right metaphor for what the brain does. If we are profoundly off the mark about the nature of thought, then this objection could be a showstopper. But research that might lead to the singularity covers a much broader range than formal computation. There is great variety even in the pursuit of pure AI. In the next decade, those who credit Moravec's timeline begin to expect results. Interestingly powerful computers will become cheap enough for a thousand research groups to bloom. Some of these researchers will pursue the classic computational tradition that Brooks is doubting—and they may still carry the day. Others will be working on their own abstractions of natural mind functions—for instance, the theory that Christof Koch and Giulio Tononi discuss in their article. Some (very likely Moravec and Brooks himself) will be experimenting with robots that cope with many of the same issues that, for animals, eventually resulted in minds that plan and feel. Finally, there will be pure neurological researchers, modeling increasingly larger parts of biological brains *in silico*. Much of this research will benefit from improvements in our tools for imaging

> The best answer to the question, "Will computers ever be as smart as humans?" is probably "Yes, but only briefly"

brain function and manipulating small regions of the brain.

But despite Moravec's estimate and all the ongoing research, we are far short of putting the hardware together successfully. In his essay, Brooks sets several intermediate challenges. Such goals can help us measure the progress that is being made. More generally, it would be good to have indicators and counterindicators to watch for. No single one would prove the case for or against the singularity, but together they would be an ongoing guide for our assessment of the matter. Among the counterindicators (events arguing against the likelihood of the singularity) would be debacles of overweening software ambition: events ranging from the bankruptcy of a major retailer upon the failure of its new inventory management system to the defeat of network-centric war fighters by a transistor-free light infantry. A tradition of such debacles could establish limits on application complexity—independent of any claims about the power of the underlying hardware.

There are many possible positive indicators. The Turing Test—whether a human judge communicating by text alone can distinguish a computer posing as human from a real human—is a subtle but broad indicator. Koch and Tononi propose a version of the Turing Test for machine consciousness in which the computer is presented a scene and asked to "extract the gist of it" for evaluation by a human judge. One could imagine restricted versions of the Turing Test for other aspects of Mind, such as introspection and common sense.

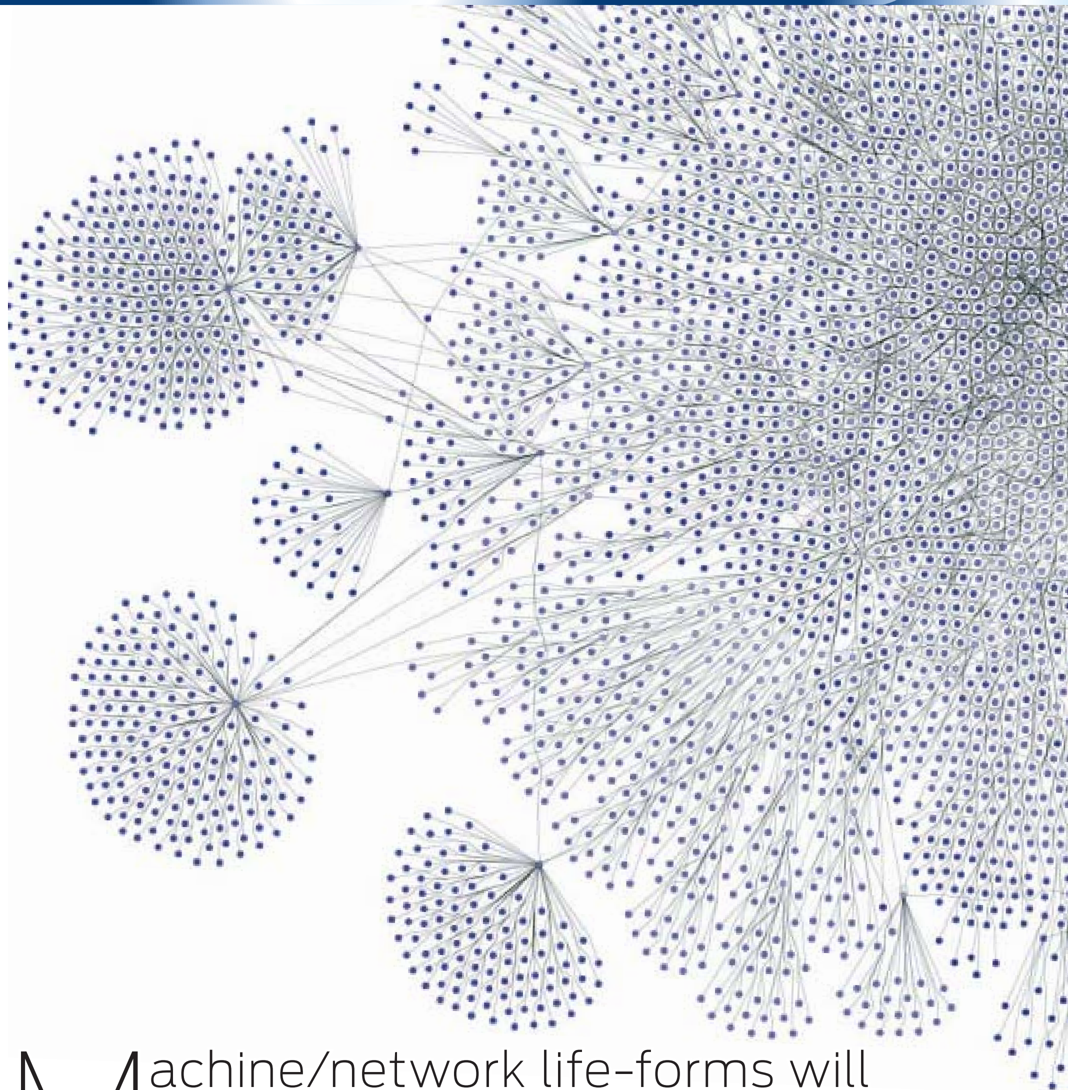

Machine/network life-forms will be faster, more labile, and more varied than what we see in biology. Digital Gaia is a hint of how alien the possibilities are

As with past computer progress, the achievement of some goals will lead to interesting disputes and insights. Consider two of Brooks's challenges: manual dexterity at the level of a 6-year-old child and object-recognition capability at the level of a 2-year-old. Both tasks would be much easier if objects in the environment possessed sensors and effectors and could communicate. For example, the target of a robot's hand could provide location and orientation data, even URLs for specialized manipulation libraries. Where the target has effectors as well as sensors, it could cooperate in the solution of kinematics issues. By the standards of today, such a distributed solution would clearly be cheating. But embedded microprocessors are increasingly widespread. Their coordinated presence may become the assumed environment. In fact, such coordination is much like relationships that have evolved between living things.

There are more general indicators. Does the distinction between neurological and AI researchers continue to blur? Does cognitive biomimetics become a common source of performance improvement in computer applications? From an entirely different direction, consider economist Robin Hanson's "shoreline" metaphor for the boundary between those tasks that can be done by machines and those that can be done only by human beings. Once upon a time, there was a continent of human-only tasks. By the end of the 1900s, that continent had become an archipelago. We might recast much of our discussion in terms of the question, "Is any place on the archipelago safe from further inundation?" Perhaps we could track this process with an objective economic index—say, wages divided by world product. However much human wealth and welfare may increase, a sustained decline in the ratio

YWORKS

of wages to world product would argue a decline in the human contribution to the economy.

Some indicators relate different areas of technological speculation. In his essay, physicist Richard A.L. Jones critiques molecular nanotechnology (MNT). Even moderate success with MNT could support Moore's Law long enough to absorb a number of order-of-magnitude errors in our estimates of the computing power of the brain. At the same time, some of the advanced applications that K. Eric Drexler describes—things like cell-repair machines—depend on awesome progress with software. Thus, while success with MNT probably does not need the technological singularity (or vice versa), each would be a powerful indicator for the other.

Several of the essays discuss the plausibility of mind uploads and consequent immortality for "our digitized psyches," ideas that have recently appeared in serious nonfiction, most notably Ray Kurzweil's *The Singularity Is Near*. As with nanotechnology, such developments aren't prerequisites for the singularity. On the other hand, the goal of enhancing human intelligence through human-computer interfaces (the IA Scenario) is both relevant and in view. Today a well-trained person with a suitably provisioned computer can look very smart indeed. Consider just a slightly more advanced setup, in which an Internet search capability plus math and modeling systems are integrated with a head-up display. The resulting overlays could give the user a kind of synthetic intuition about his or her surroundings. At a more intimate but still noninvasive level, DARPA's Cognitive Technology Threat Warning System is based on the idea of monitoring the user's mental activities and feeding the resulting analysis back to the user as a supplement to his or her own attention. And of course there are the researchers working with direct neural connections to machines. Larger numbers of implanted connections may allow selection for effective subsets of connections. The human and the machine sides can train to accommodate each other.

To date, research on neural prostheses has mainly involved hearing, vision, and communication. Prostheses that could restore any cognitive function would be a very provocative indicator. In his essay, John Horgan discusses neural research, including that of T.W. Berger, into prostheses for memory function. In general, Horgan and I reach very different conclusions, but I don't think we have much disagreement about the facts; Horgan cites them to show how distant today's technology is from anything like the singularity—and I am saying, "Look here, these are the sorts of things we should track going forward, as signs of progress toward the singularity (or not)."

The Biomedical Scenario—directly improving the functioning of our own brains—has a lot of similarities to the IA Scenario, though computers would be only indirectly involved, in support of bioinformatics. In the near future, drugs for athletic ability may be only a small problem compared with drugs for intellect. If these mind drugs are not another miserable fad of uppers and downers, if they enable real improvements to memory and creativity, that would be a strong indicator for this scenario. Much further out—for both logistical and ethical reasons—is the possibility of embryo optimization and germ-line engineering. Biomedical enhancement, even the extreme varieties, probably does not scale very well; however, it might help biological minds maintain some influence over other progress.

If the singularity happens, the world passes beyond human ken

Brooks suggests that the singularity might happen—and yet we might not notice. Of the scenarios I mentioned at the beginning of this essay, I think a pure Internet Scenario—where humanity plus its networks and databases become a superhuman being—is the most likely to leave room to argue about whether the singularity has happened or not. In this future, there might be all-but-magical scientific breakthroughs. The will of the people might manifest itself as a seamless transformation of demand and imagination into products and policy, with environmental and geopolitical disasters routinely finessed. And yet there might be no explicit evidence of a superhuman player.

A singularity arising from networks of embedded microprocessors—the Digital Gaia Scenario—would probably be less deniable, if only because of the palpable strangeness of the everyday world: reality itself would wake up. Though physical objects need not be individually sapient, most would know what they are, where they are, and be able to communicate with their neighbors (and so potentially with the world). Depending on the mood of the network, the average person might notice a level of convenience that simply looks like marvelously good luck. The Digital Gaia would be something beyond human intelligence, but nothing like human. In general, I suspect that machine/network life-forms will be faster, more labile, and more varied than what we see in biology. Digital Gaia is a hint of how alien the possibilities are.

In his essay, Hanson focuses on the economics of the singularity. As a result, he produces spectacular insights while avoiding much of the distracting weirdness. And yet weirdness necessarily leaks into the latter part of his discussion (even leaving Digital Gaia possibilities aside). AI at the human level would be a revolution in our worldview, but we can already create human-level intelligences; it takes between nine months and 21 years, depending on whom you're talking to. The consequences of creating human-level artificial intelligence would be profound, but it would still be explainable to present-day humans like you and me.

But what happens a year or two after that? The best answer to the question, "Will computers ever be as smart as humans?" is probably "Yes, but only briefly."

For most of us, the hard part is believing that machines could ever reach parity. If that does happen, then the development of superhuman performance seems very likely—and that is the singularity. In its simplest form, this might be achieved by "running the processor clock faster" on machines that were already at human parity. I call such creatures "weakly superhuman," since they should be understandable if we had enough time to analyze their behavior. Assuming Moore's Law muddles onward, minds will become steadily smarter. Would economics still be an important driver? Economics arises from limitations on resources. Personally, I think there will always be such limits, if only because Mind's reach will always exceed its grasp.

However, what is scarce for the new minds and how they deal with that scarcity will be mostly opaque to us.

The period when economics could help us understand the new minds might last decades, perhaps corresponding to what Brooks describes as "a period, not an event." I'd characterize such a period as a soft takeoff into the singularity. Toward the end, the world would be seriously strange from the point of view of unenhanced humans.

A soft takeoff might be as gentle as changes that humanity has encountered in the past. But I think a hard takeoff is possible instead: perhaps the transition would be fast. One moment the world is like 2008, perhaps more heavily networked. People are still debating the possibility of the singularity. And then something...happens. I don't mean the accidental construction that Brooks describes. What I'm thinking of would probably be the result of intentional research, perhaps a group exploring the parameter space of their general theory. One of their experiments finally gets things right. The result transforms the world—in just a matter of hours. A hard takeoff into the singularity could resemble a physical explosion more than it does technological progress.

I base the possibility of hard takeoff partly on the known potential of rapid malcode (remember the Slammer worm?) but also on an analogy: the most recent event of the magnitude of the technological singularity was the rise of humans within the animal kingdom. Early humans could effect change orders of magnitude faster than other animals could. If we succeed in building systems that are similarly advanced beyond us, we might experience a similar incredible runaway.

Whether the takeoff is hard or soft, the world beyond the singularity contains critters who surpass natural humans in just the ability that has so empowered us: intelligence. In human history, there have been a number of radical technological changes: the invention of fire, the development of agriculture, the Industrial Revolution. One might reasonably apply the term *singularity* to these changes. Each has profoundly transformed our world, with consequences that were largely unimagined beforehand. And yet those consequences could have been explained to earlier humans. But if the transforma-

To Probe Further

Vernor Vinge, "The Coming Technological Singularity," Vision-21 Symposium, NASA Lewis Research Center and the Ohio Aerospace Institute, 30 to 31 March, 1993 (http://rohan.sdsu.edu/faculty/vinge/misc/singularity.html).

Vernor Vinge, "What If the Singularity Does NOT Happen," Seminars About Long-term Thinking, The Long Now Foundation, February 15, 2007 (http://rohan.sdsu.edu/faculty/vinge/longnow/index.html; see also http://www.longnow.org/projects/seminars).

Hans Moravec, *Mind Children*, Harvard University Press, 1988.

Theodore W. Berger et al., "Brain-Implantable Biomimetic Electronics as a Neural Prosthesis for Hippocampal Memory Function," in *Toward Replacement Parts for the Brain*, Berger and Glanzman, eds., MIT Press, 2005.

J. Storrs Hall, *Beyond AI: Creating the Conscience of the Machine*, Prometheus Books, 2007. Hall discusses strategies for reaching a favorable outcome in a soft takeoff.

For the complete references to this article go to: http://spectrum.ieee.org/jun08/singularityprobe.

tion discussed in this issue of *Spectrum* occurs, the world will become intrinsically unintelligible to the likes of us. (And that is why "singularity," as in "black hole singularity of physics," is *the* cool metaphor here.) If the singularity happens, we are no longer the apex of intellect. There will be superhumanly intelligent players, and much of the world will be to their design. Explaining that to one of us would be like trying to explain our world to a monkey.

Both Horgan and Nordmann express indignation that singularity speculation distracts from the many serious, real problems facing society. This is a reasonable position for anyone who considers the singularity to be bogus, but some form of the point should also be considered by less skeptical persons: if the singularity happens, the world passes beyond human ken. So isn't all our singularity chatter a waste of breath? There are reasons, some minor, some perhaps very important, for interest in the singularity. The topic has the same appeal as other great events in natural history (though I am more comfortable with such changes when they are at a paleontological remove). More practically, the notion of the singularity is simply a view of progress that we can use—along with other, competing, views—to interpret ongoing events and revise our local planning. And finally: if we are in a soft takeoff, then powerful components of superintelligence will be available well before any complete entity. Human planning and guidance could help avoid ghastliness, or even help create a world that is too good for us naturals to comprehend.

Horgan concludes that "the singularity is a religious rather than scientific vision." Brooks is more mellow, seeing "commonalities with religious beliefs" in many enthusiasts' ideas. I argue against Horgan's conclusion, but Brooks's observation is more difficult to dispute. If there were no other points to discuss, then those commonalities would be a powerful part of the skeptics' position. But there are other, more substantive arguments on both sides of the issue.

And of course, the spirituality card can be played against both skeptics and enthusiasts: Consciousness, intelligence, self-awareness, emotion—even their definitions have been debated since forever, by everyone from sophomores to great philosophers. Now, because of our computers, the applications that we are attempting, and the tools we have for observing the behavior of living brains, there is the possibility of making progress with these mysteries. Some of the hardest questions may be ill-posed, but we should see a continuing stream of partial answers and surprises. I expect that many successes will still be met by reasonable criticism of the form "Oh, but that's not really what intelligence is about" or "That method of solution is just an inflexible cheat." And yet for both skeptics and enthusiasts, this is a remarkable process. For the skeptic, it's a bit like subtractive sculpture, where step-by-step, each partial success is removing more dross, closing in on the ineffable features of Mind—a rather spiritual prospect! Of course, we may remove and remove and find that ultimately we are left with nothing but a pile of sand—and devices that are everything we are, and more. If that is the outcome, then we've got the singularity. ❑

ieee spectrum classifieds

The Norwegian University of Science and Technology (NTNU) in Trondheim represents academic eminence in technology and the natural sciences as well as in other academic disciplines ranging from the social sciences, the arts, medicine, architecture to fine art. Cross-disciplinary cooperation results in innovative breakthroughs and creative solutions with far-reaching social and economic impact.

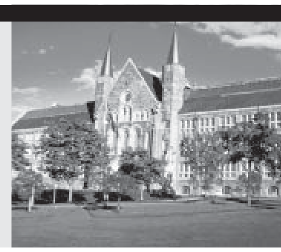

Professor/Associate Professor in Automatic Control in Trondheim, Norway

The position is in the Department of Engineering Cybernetics at the Faculty of Information Technology, Mathematics and Electrical Engineering, the Norwegian University of Science and Technology (NTNU). The Department has for the moment 9 full professors, 5 associate/assistant professors, 8 adjunct professors, 4 postdoc candidates, and about 40 PhD students.

The successful applicant must have a solid background in several areas of automatic control, for example: control of linear and nonlinear systems, adaptive control, optimization and optimal control, control of distributed parameter systems, robust control, stochastic control systems, state estimation, and system identification. The candidate must, moreover, have an outstandig scientific record in at least one of the fields of interest mentioned above. Exceptional candidates in other areas of systems and control theory will be given serious consideration.

The successful candidate will have a particular responsibility, in conjunction with the Department's other faculty members, for research in the field/fields of interest where he/she has documented scientific qualifications. The candidate must, moreover, have a record of relevant applied research and development and will have to take part in the Department's activities in that field. Applicants are expected to submit a research plan for the first five years. A proven track record of attracting external research funding is given weight in the evaluation of applicants.

The full details about the position and how to apply can be found at
http://innsida.ntnu.no/nettopp_lesmer.php?kategori=nyheter&dokid=47f4b9a3567105.56669390

The application deadline is 2008-08-01.

Further information about this professorship can be obtained from
Professor Kristin Y. Pettersen, phone no. + 47 7359 4346, e-mail: kyp@itk.ntnu.no
or Professor Morten Hovd, phone no. + 47 7359 14 26, e-mail: Morten.Hovd@itk.ntnu.no

NTNU would like to increase the percentage of female scientists in academic positions. The following initiatives are relevant for women applying for this position: Startup package for women in male-dominated fields. Qualification scholarships and mentor program for women in academic positions.

Trondheim is an ideal place to combine an academic career with a good personal life. NTNU has long traditions in high-quality engineering research and education, as a successor to the Norwegian Institute of Technology (NTH). The Department of Engineering Cybernetics has a strong international reputation. During the last decade, department faculty has twice won the IEEE Transactions on Control Systems Technology Outstanding Paper Award and twice the Automatica Best Paper Award.

Norway, like other Scandinavian countries, has generous rules for both maternity and paternity leave. Professional daycare for children is easily available. Trondheim has two international schools with English as the language of instruction, including Pre-school, Primary and Middle school.
Having a population of 160000, Trondheim is a small city by international standards, and thus has low crime rates, little pollution, and easy access to a beautiful countryside. Still, the city can rival many larger cities in the availability of cultural facilities, having both a professional theater and a professional symphony orchestra. The Trondheim Municipal School of Music and Culture is outstanding, and caters for all levels of talent and a wide variety of cultural genres. NTNU includes the Trondheim Conservatory of Music, which offers a high quality classical music program, and is one of the World's leading places of learning for jazz musicians.

CICERO ev

NTNU
Norwegian University of Science and Technology

Open Postdoc positions: at EISLAB at Lulea University of Technology in Sweden. Specialization in Analog Electronics and Embedded Systems. More information at http://www.ltu.se/csee/news/postdoc

國立交通大學

National Chiao Tung University at Hsinchu, Taiwan:

The Department of Electrical Engineering invites applications for several tenure track faculty positions at all levels. Areas of interest to the department include microwave/millimeter-wave circuits and antennas, RFIC, and microwave packaging technique. Candidates should have a Ph.D. degree in related fields, have demonstrated record of research accomplishment in these areas and strong commitment to excellence in teaching. Please send (or e-mail) curriculum vitae and names of at least three references with both postal and e-mail addresses to Institute of Communications Engineering, Department of Electrical Engineering, ED819, National Chiao Tung University, 1001 Ta Hsueh Road, Hsinchu, 300 Taiwan (e-mail: **qponing@mail.nctu.edu.tw**). Review of applications will commence in April 2008 and continue until the positions are filled. The department has the country's largest and strongest postgraduate program in communication science and technologies, with 80 full-time faculty members and more than 800 graduate students.

ESIEE

ESIEE-Paris, a member of the University Paris-Est provides electronics, computer science, telecommunications, and control systems and technology management.

Applications are open for 7 permanent professor positions
to be involved in teaching and research
(in France from September, 2008)

- **Computer science and networks**
- **Computer Science and imaging**
- **Real-time Distributed Systems**
- **Nanotechnology and Micro Systems**
- **Urban communication systems**
- **Embedded Electronics Systems**
- **Digital Electronics Design**

Online applications : http://www.esiee.fr/groupe/postes.php

The appointee will require a strong interest and track record in scientific research and its industrial applications. Previous R&D experience is a plus. Owing to the international nature of this appointment, knowledge of French is a plus, but teaching can be in English.

Applications must include full curriculum vitae, reprints of a relevant recent publications and a short description of present and future research programs.

Deadline for application : June 15th, 2008
Application materials must be sent to the following address :

Doyen du Corps Professoral
ESIEE - Cité Descartes – BP99
93162 Noisy-le-Grand cedex - France
Tél. : (33) 01 45 92 65 16
Télécopie : (33) 01 45 92 66 99
Email : doyen@esiee.fr

THE SCHOOLS OF INNOVATION AND ADVANCED TECHNOLOGIES

ONR Global is searching for scientists and engineers for international positions.

The mission of the Office of Naval Research (ONR) Global is to serve as the primary international science and technology (S&T) node of the Naval Research Enterprise, provide access to the worldwide S&T community and to ensure technological awareness.

ONR Global is seeking highly qualified individuals to serve on a temporary assignment (2-4 years) as liaison scientists and engineers in our offices in London and Tokyo. Liaison scientists and engineers build relationships with emerging and leading international scientists, provide technical assessments, and facilitate the integration of international research into ONR programs. Currently we are seeking scientists and engineers with expertise in a range of Naval relevant disciplines including:

- **Applied physics**
- **Autonomous systems**
- **Behavioral sciences**
- **Communication networks**
- **Directed energy**
- **Human factors**
- **Materials science**

For information on the Office of Naval Research Global please visit our webpage at www.onrglobal.navy.mil.asp. Please send your resume to cstewart@onrglobal.navy.mil if interested. Review of applications will be ongoing throughout the year. This is a detail opportunity for current federal government employees or an Intergovernmental Personnel Act (IPA) opportunity for employees of universities or non-profit organizations.

U.S. CITIZENSHIP REQUIRED • AN EQUAL OPPORTUNITY EMPLOYER

香 港 大 學
THE UNIVERSITY OF HONG KONG

Founded in 1911, The University of Hong Kong is committed to the highest international standards of excellence in teaching and research, and has been at the international forefront of academic scholarship for many years. Of a number of recent indicators of the University's performance, one is its ranking at 18 among the top 200 universities in the world by the UK's *Times Higher Education Supplement*. The University has a comprehensive range of study programmes and research disciplines, with 20,000 undergraduate and postgraduate students from 50 countries, and a complement of 1,200 academic members of staff, many of whom are internationally renowned.

Assistant Professor in Photonics
in the Department of Electrical and Electronic Engineering
(Ref.: RF-2007/2008-643)

Applications are invited for appointment as Assistant Professor in Photonics in the Department of Electrical and Electronic Engineering, from November 2008 on a three-year fixed-term basis, with consideration for tenure after satisfactory completion of a second contract.

The Department offers B.Eng., M.Sc., M.Phil. and Ph.D. programmes. The B.Eng. programmes comprise Electronic and Communications Engineering, Information Engineering, Electrical Engineering and Computer Engineering (jointly with the Department of Computer Science). The Department consists of 40 full-time teaching staff and has excellent computing resources, well-equipped teaching and research facilities and support. Information about the Department can be obtained at http://www.eee.hku.hk/.

Applicants should possess a Ph.D. degree in Electrical/Electronic/Computer Engineering or a closely related field, and an excellent research record.

Annual salary for Assistant Professorship will be in the range of HK$474,600 - HK$733,440 (approx. US$1 = HK$7.8) (subject to review from time to time at the entire discretion of the University). At current rates, salaries tax does not exceed 17% of gross income. The appointment will attract a contract-end gratuity and University contribution to a retirement benefits scheme, totalling up to 15% of basic salary, as well as leave, and medical/dental benefits. Housing benefits will be provided as applicable.

Further particulars and application forms (272/302 amended) can be obtained at https://www.hku.hk/apptunit/; or from the Appointments Unit (Senior), Human Resource Section, Registry, The University of Hong Kong, Hong Kong (**fax**: (852) 2540 6735 or 2559 2058; **e-mail**: senrappt@hkucc.hku.hk). **Closes June 23, 2008. Candidates who are not contacted within 4 months from the closing date may consider their applications unsuccessful.**

The University is an equal opportunity employer and is committed to a No-Smoking Policy

BOY WONDER: The superpowered, ever benign Astro Boy touched the hearts of millions of fans in post–World War II Japan.

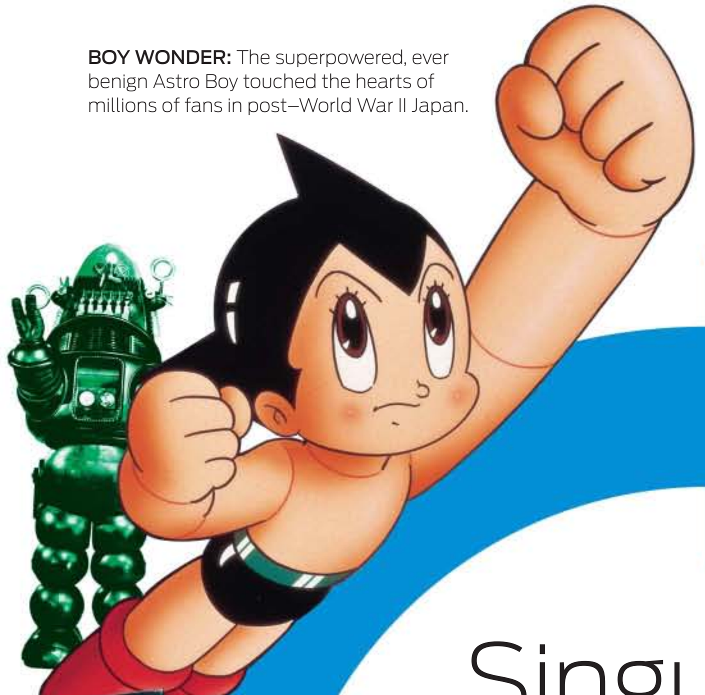

GOOD INTENTIONS: HAL 9000 and Colossus mean well, but their single-mindedness in carrying out their programmed objectives turns out badly for the humans in their midst.

Singularity Index

FRIEND

ENEMY

The singularity—that pivotal moment when machines attain humanlike intelligence—may never arrive, but don't tell Hollywood. Over the years, writers, directors, and set designers have envisioned worlds in which machines rule. Whether benign or evil or somewhere in between, these mechanized souls tend to mirror society's own attitudes toward technology. In the 1927 silent classic *Metropolis*, the robot doppelgänger Maria foments rebellion among the human workers—a reflection of that era's real-world struggles over labor and class. The trust-in-technology 1950s and early 1960s, by contrast, yielded a fleet of friendly helpers, from Robby the Robot in *Forbidden Planet* to Rosey, the automaton maid in "The Jetsons." Recent incarnations of humanlike machines have been more subtly drawn: though often physically superior, they remain conflicted about their existence and uneasy about their human creators. Here we offer a few depictions of the singularity, as seen through the lens of pop culture. ❑

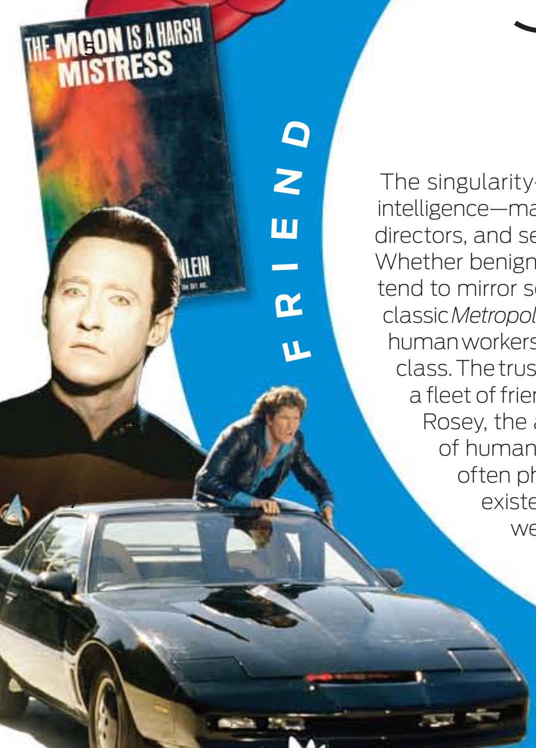

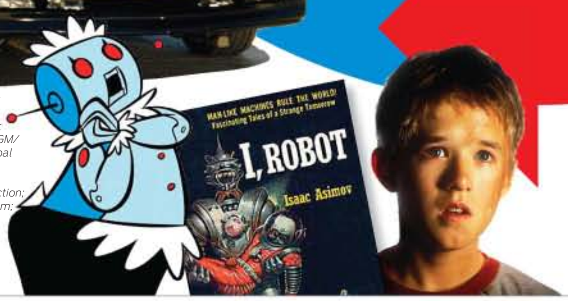

PHOTOS, CLOCKWISE FROM TOP: Carolco/The Kobal Collection; Warner Bros./The Kobal Collection; Universal/The Kobal Collection; MGM/The Kobal Collection; UFA/The Kobal Collection; Ace Hardcover; Sci Fi; Cat's Collection/Corbis; Amblin/Dreamworks/WB/The Kobal Collection; Spectra; Turner Broadcasting System; Universal TV/The Kobal Collection; Paramount Television/The Kobal Collection; G.P. Putnam's Sons; MGM/The Kobal Collection; Courtesy of Tezuka Productions Co.

DEADLY BUT BEAUTIFUL: Number Six, the robotic femme fatale of the sci-fi series "Battlestar Galactica," engineers the deaths of billions of people—yet longs for their approval.

THINKING MACHINES [clockwise from top]: *The Terminator* (film, 1984); Agent Smith, from *The Matrix* (film, 1999); Colossus supercomputer, from *Colossus: The Forbin Project* (film, 1970); HAL 9000, from *2001: A Space Odyssey* (film, 1968); Robot Maria, from *Metropolis* (film, 1927); William Gibson's *Neuromancer* (novel, 1984); Number Six, a Cylon from "Battlestar Galactica" (TV series, 2003–); Roy Batty, a replicant from *Blade Runner* (film, 1982); the boy-robot David, from *Artificial Intelligence: A.I.* (film, 2001); Isaac Asimov's *I, Robot* (short-story collection, 1950); Rosey the Robot, from "The Jetsons" (TV series, 1962–63); K.I.T.T., the talking car, from "Knight Rider" (TV series, 1982–86); android Lt. Commander Data, from "Star Trek: The Next Generation" (TV series, 1987–94); Robert Heinlein's *The Moon Is a Harsh Mistress* (novel, 1966); Robby, from *Forbidden Planet* (film, 1956); Astro Boy (Japanese manga series, original run 1952–68).