



# THE SINGULARITY

WALKING UPRIGHT. THE DISCOVERY OF FIRE. SPEECH. ARE WE ON THE VERGE OF ANOTHER STEP IN THE ASCENT OF MAN? WILL WE ENHANCE OUR OWN MINDS AND GAIN A PEEK AT IMMORTALITY? OR WILL WE JUST FIND OURSELVES SURROUNDED BY A SUPERINTELLIGENT ARMY OF MACHINES?

Let's say you transfer your mind into a computer—not all at once but gradually, having electrodes inserted into your brain and then wirelessly outsourcing your faculties. Your vision is rerouted through cameras, your memories are stored in a net of microprocessors and so on, until at last the transfer is complete. As neuroengineers get to work boosting the performance of your uploaded brain so you can now think as a god, your fleshy brain is heaved into a bag of medical waste. As you—for now let's just call it “you”—start a new chapter of existence exclusively within a machine, an existence that will last as long as there are server farms and hard-disk space and the solar power to run them, are “you” still actually you?

This question was being considered carefully and thoroughly by a 43-year-old man standing on a giant stage backed by high black curtains. He had the bedraggled hair and beard of a Reagan-era metalhead. He wore a black leather coat and an orange-and-red T-shirt covered in

stretched-out figures from a Stone Age cave painting.

He was not, in fact, insane.

The man was David Chalmers, one of the world's leading philosophers of the mind. He has written some of the most influential papers on the nature of consciousness. He is director of the Centre for Consciousness at Australian National University and is also a visiting professor at New York University. In other words, he has his wits about him.

Chalmers was speaking midway through a conference in New York called Singularity Summit 2009, where computer scientists, neuroscientists and other researchers were offering their visions of the future of intelligence. Some ideas were tentative, while others careened into what seemed like science fiction. At their most extreme the speakers foresaw a time when we would understand the human brain in its fine details, be able to build machines not just with artificial intelligence but with superintelligence and be able to merge our own minds with those machines.

BY  
CARL  
ZIMMER

“This raises all kinds of questions for a philosopher,” Chalmers said. “Question one: Will an uploaded system be conscious? Uploading is going to suck if, once you upload yourself, you’re a zombie.”

Chalmers didn’t see why an uploaded brain couldn’t be conscious. “There’s no difference in principle between neurons and silicon,” he said. But that led him to question number two: “Will an uploaded system be me? It’s not a whole lot better to be conscious as someone else entirely. Good for them, not so good for me.”

To try to answer that question Chalmers asked what it takes to be me. It doesn’t take a certain set of atoms, since our neurons break down their molecules and rebuild them every day. Chalmers pondered the best way to guarantee the survival of your identity: “Gradual uploading is the way to go, neuron by neuron, staying conscious throughout.”


But perhaps that won’t be an option. Perhaps you will have died by the time you are uploaded. Chalmers didn’t find this as alarming as others do. “Let’s call it the Buddhist view,” he said. Every day, he pointed out, we lose consciousness as we

developed a strong immune defense against hype. The Singularity, with all its promises of a techno-rapture, seems tailor-made to bring out the worst in people like me. The writer John Horgan, who has even less patience for the promises science cannot keep, wrote “Science Cult,” a devastating essay about the Singularity, for *Newsweek.com* in May 2009.

He acknowledges part of him enjoys pondering the Singularity’s visions, such as boosting your IQ to 1,000. “But another part of me—the grown-up, responsible part—wonders that so many people, smart people, are taking Kurzweil’s sci-fi fantasies seriously,” he writes. “The last thing humanity needs right now is an apocalyptic cult masquerading as science.”

I decided to check out the Singularity for myself. The summit turned out to be one of the most bizarre experiences I’ve had. Chalmers wasn’t the only speaker to induce hallucinations. Between the talks, as I mingled among people wearing *S* lapel pins and eagerly discussing their personal theories of consciousness, I found myself tempted to reject the whole smorgasbord as half-baked science fiction. But in the end I didn’t.

After the meeting I visited researchers working on the type



HELLO, DAVE. 2001: A SPACE ODYSSEY WAS RIGHT ON ONE COUNT: THERE WILL BE MACHINES THAT ARE MORE CLEVER THAN WE ARE.

fall asleep and then regain it the next morning. “Each waking is a new dawn that’s a bit like the commencement of a new person. But it turns out that’s good enough. That’s what ordinary survival is. We’ve lived there a long time. And if that’s so, then reconstructive uploading will also be good enough.”



If the term *singularity* rings a bell, maybe you’ve read the 2005 bestseller *The Singularity Is Near*. Its author, computer scientist and inventor Ray Kurzweil, confidently predicts intelligence will soon cross a profound threshold. The human brain will be dramatically enhanced with engineering. Artificial intelligence will take on a life of its own. If all goes well, Kurzweil predicts, we will ultimately fuse our minds with this machine superintelligence and find a cybernetic immortality. What’s more, the Singularity is coming soon. Many of us alive today will be part of that transformation.

The Singularity is not only a future milestone but also a peculiar movement today. Along with spaceflight tycoon Peter Diamandis, Kurzweil has launched Singularity University, which brought in its first batch of students in the summer of 2009. Kurzweil is also director of the Singularity Institute for Artificial Intelligence, which held its first annual summit in 2006. The summits are a mix of talks by Kurzweil and other Singularity advocates, along with scientists working on everything from robot cars to gene therapy. For its first three years the Singularity Summit took place around the Bay Area, but in 2009 the institute decided to decamp from its utopian environs and head for the more cynical streets of New York.

I was one of the curious skeptics who heeded the call and came to the 92nd Street Y. I’ve been writing about new advances in science for 20 years, and along the way I’ve

of technology that Kurzweil and others consider the stepping-stones to the Singularity. Not one of the researchers takes Kurzweil’s extreme vision of the future seriously. We will not have some sort of cybernetic immortality in the next few decades. The human brain is far too mysterious and computers far too crude for such a union anytime soon, if ever. In fact some scientists regard all this talk of the Singularity as nothing more than recklessly offering false hope to people struggling with blindness, paralysis and other disorders.

But when I asked these skeptics about the future, even their most conservative visions were unsettling: a future in which people boost their brains with enhancing drugs, for example, or have sophisticated computers implanted in their skulls for life. While we may never be able to upload our minds into a computer, we may still be able to build computers based on the layout of the human brain. I can report I have not drunk the Singularity Kool-Aid, but I have taken a sip.



The future is not new. By the dawn of the 20th century science was moving so fast many people were sure we were on the verge of tremendous change. The blogger Matt Novak collects entertainingly bad predictions at his website *Paleo-Future*. My favorite is a 1900 article by John Watkins that appeared in *Ladies’ Home Journal*, offering readers a long list of predictions from leading thinkers about what life would be like within the next 100 years.

“A man or woman unable to walk 10 miles at a stretch will be regarded as a weakling,” Watkins wrote. “There will be no C, X or Q in our everyday alphabet.”

As science advanced through the 20th century, the future morphed accordingly. When (continued on page 185)

# SINGULARITY

(continued from page 50)

scientists figured out how to culture animal cells in the early 1900s, some claimed such cells would let us live forever. In the 1940s the availability of antibiotics led some doctors to declare the age of infectious diseases over. The engineers who founded NASA were sure we would build cities on the moon, perhaps even Mars. And as scientists began to develop computers and the programs to run them, they began to predict that someday—someday soon—computers would gain a human intelligence.

Their confidence grew not just from their own research. Neuroscientists were discovering that our brains act a lot like computers, so it seems logical we would someday be able to translate the neural code, build computers that process information similarly and even join brains and machines together.

In 1993 the science-fiction writer Vernor Vinge wrote an essay on this particular kind of future. He entitled it “The Coming Technological Singularity,” borrowing a term astrophysicists use to describe a place where the ordinary rules of gravity and other forces break down. “Within 30 years we will have the technological means to create superhuman intelligence,” he claimed. “Shortly after, the human era will be ended.”

By the late 1990s Kurzweil emerged as the leading champion of the coming end of life as we know it. He started as a tremendously successful computer scientist, having invented print-to-speech machines for the blind, music synthesizers and a host of other devices, and in 1990 he published his first forward-looking book, *The Age of Intelligent Machines*. He argues that within a few decades computers will be as intelligent as humans, if not more so.

As the years passed, his predictions grew more extreme. In his 1999 book, *The Age of Spiritual Machines*, he imagines life in 2099. “The number of software-based humans vastly exceeds those still using native neuron-cell-based computation,” he writes. In 2005 he brought Vinge’s term to wide attention in his book *The Singularity Is Near*, in which he bemoans how hobbled we are by feeble neurons, bones and muscles. “The Singularity,” he writes, “will allow us to transcend these limitations of our biological bodies and brains.”

At the Singularity Summit Kurzweil came onstage to offer the latest iteration of his case for the Singularity. The audience broke into fierce applause when he appeared, a few people standing and pounding their hands in slow motion. Kurzweil was, as ever, sharply dressed, wearing a tailored blue suit, an open striped shirt and narrow glasses. (In 1984 a writer for *Business Week* noted he “wears expensive Italian suits and a gold Mickey Mouse watch.”)

He launched into his talk, leaning back on one foot as he spoke, his small frame angled diagonally to the audience, his eyebrows softly raised, his entire body seemingly caught in a perpetual shrug. Rather than slamming the audience with an infomercial pitch, his languid body language seemed to be saying, “Look, I don’t care if you believe me or not, but these are the facts.”

He talked about everything from quantum physics to comas to speech recognition, but at the heart of his talk was a series of graphs. They showed an exponential growth in the power of technology, from the speed of DNA sequencing efforts to the power of computers to the growth of the Internet. This exponential growth has been so relentless that Kurzweil has dubbed it the law of accelerating returns.

“It really belies the common wisdom that you can’t predict the future,” he said as he gazed at the graphs.

Thanks to the law of accelerating returns, he said, technology will continue to leap forward and astonish us. “Thirty linear steps take you to 30,” he said. “Thirty exponential steps take you to a billion.”

One way to judge whether Kurzweil is right about the future is to see how well he has done in the past. In the realm of computer science he has done pretty well. In 1990 he predicted the world chess champion would be a computer by 1998. IBM’s Deep Blue computer beat Garry Kasparov in 1997. In 1999 Kurzweil predicted that in 10 years computers communicating with one another and the world wide web wirelessly would be commonplace. It’s already becoming hard to recall when computers were lashed to modems.

But many of Kurzweil’s predictions have failed. In 1999 he predicted that by 2009 “bioengineered treatments for cancer and heart disease [will] have greatly reduced the mortality from these diseases.” In 2006, the most recent year for which statistics are available, 829,072 people died in the United States of heart disease. Fortunately the death rate from heart disease is lower now than in 1950, but that drop is due mainly to low-tech measures such as getting people to stop smoking. Meanwhile the death rate from cancer has dropped only five percent since 1950.

These failed predictions reveal a weakness at the heart of Kurzweil’s forecasts. Scientific understanding doesn’t advance in lockstep with technological horsepower. It was funny, in a morbid way, to watch Kurzweil make his case inside the 92nd Street Y just as a surge of swine flu viruses was sweeping the city. There was a time when sequencing the entire genome of a single flu virus was a colossal, budget-busting project; now it costs a few hundred dollars. As of October 2009 the Influenza Genome Sequencing Project has collected the complete genomes of 4,115 viruses from around the world, and that number is rising rapidly. All that raw genetic information has certainly allowed scientists to learn important things about the flu, but it has not given New York any fancy new way to stop it in its tracks. New Yorkers could only wash their hands as they waited for the delivery of new vaccines and hoped their hospitals didn’t get overwhelmed by people in need of respirators. All thanks to a virus with just 10 genes.

Even as flu viruses multiplied through the city, Kurzweil happily continued his talk, mocking the skeptics who scoffed when scientists were trying to sequence all

3.5 billion “letters” of DNA in the human genome. For a long time it seemed as if they’d never finish, and then in the early 2000s they were done.

Years later we still have lots of open questions about how the genome actually works. Scientists used to think it contained 100,000 protein-coding genes, but it turns out to have just 20,000, and researchers don’t actually know the function of many of them. What’s more, the genome also has tens of thousands of genes our cells use to make single-strand versions of DNA called RNA. Many play vital roles in the cell, but a lot probably don’t do anything, and scientists are just starting to figure out what does what. Some promised that sequencing the human genome would yield cures to most diseases. Scientists are now searching the genome for genes that increase your risk of high blood pressure, diabetes and other diseases. In most cases they have found lots of genes that raise your risk by only a barely measurable amount. Sequencing the human genome has revealed to scientists that they know less than they thought they did.

When I started contacting experts about the Singularity, something surprising happened—they didn’t laugh and hang up the phone.

“I find some people way too quick to pooh-pooh the idea of an impending Singularity,” said Martha Farah, director of the Center for Cognitive Neuroscience at the University of Pennsylvania.

Farah investigates how people try to enhance their cognition with drugs. Drugs originally designed to treat mental disorders are now being taken by perfectly healthy people. Adderall, a drug for ADHD, is a popular campus drug for boosting concentration. Modafinil, developed for people with narcolepsy, is now a drug of choice for those who want to burn the midnight oil.

In the years to come, Farah anticipates, even more powerful drugs will come to market. Some are intended to slow the disappearance of memories in people with Alzheimer’s disease; others may boost cognition in people with impairments. She expects there will be people—maybe a lot of them—who will take these drugs in the hopes of boosting an already healthy brain, not to fix a deficit.

In December 2008 Farah and a group of fellow neuroscientists and bioethicists published an article in *Nature* claiming this kind of brain boosting is okay. “Mentally competent adults should be able to engage in cognitive enhancement using drugs,” they write. It’s impractical, they argue, to try to draw a line between treating a disease and enhancing a healthy brain.

Farah sees an urgent need now to measure the actual enhancement these drugs can bring. “The effectiveness of Adderall depends crucially on the individual,” she said. “The literature suggests that people who are average or below get the biggest benefit. The high performers may get no benefit or may actually be impaired by it.” She is now measuring the performance of students on Adderall and placebos to see if that’s actually the case.

Farah feels the drug dosing of today will have a profound impact on how we treat our brains: "I think this growing practice may be softening us up to accept more drastic brain modifications down the line."

"Here are some critters," said Ed Boyden. The serene young leader of the Synthetic Neurobiology Group at MIT stood in his laboratory. Tiny pieces of electronics were strewn across the lab benches. Dishes full of neurons were positioned under microscopes. And there were also flasks of algae, one of which Boyden had grabbed and held up for me to see. He sloshed the green fluid around. I came to Boyden's lab after seeing him give a remarkable talk at the Singularity Summit. Boyden is at the forefront of a field known as neuroengineering. Neuroengineers seek to restore damaged brains by implanting tiny on-board computers and electrodes. Boyden's research may take neuroengineering to a new level. Rather than use electricity to manipulate the brain, he wants to use light. But in order to do that, he will have to borrow some genes from the algae he was holding and put them in human neurons. If he succeeds, people will be part machine, but they will also be a little bit algae, too.

The logic behind brain implants is simple. Many disorders, from blindness to paralysis, come down to a break in the flow of signals through our nervous system. Neuroengineers have long dreamed of building tiny machines that could restore that flow. So far, they've had one great success: the cochlear implant, a machine that delivers sound to the brains of the deaf. A cochlear implant picks up sounds with an external microphone and converts them to electronic signals, which travel down wires into a shell in the ear called the cochlea, where they tickle the ends of the auditory nerves.

The first generation of cochlear implants, in the 1970s, were big, awkward devices with wires crossing the skull, raising the risk of infection. They used up power quickly and produced crude perceptions of sound. In the 1990s scientists developed microphones small enough to perch on the ear that transmit sounds wirelessly to an implanted receiver. Today more than 180,000 people use cochlear implants. Scientists continue to make improvements to the implants so they can run on far less energy yet perform even better.

Neuroengineers have also been testing implants that go into the brain itself, but progress has been slower on that front. So far 30,000 people have had electrodes implanted in their brains to help them cope with Parkinson's disease. Pulses of electricity from the implants make it easier for them to issue the commands to move their bodies. Other scientists are experimenting with similar implants to treat other disorders. In October 2009 scientists reported 15 people with Tourette's syndrome had 52 percent fewer tics thanks to deep-brain stimulation. Other scientists are trying to build the visual equivalent of a cochlear implant. They've linked cameras to electrodes implanted in the visual centers of blind people's brains. Stimulating those electrodes allows the subjects to see a few spots of light.

While these electrodes send electricity into the brain, other units that researchers are working on pull information out. At Massachusetts General Hospital doctors have started clinical trials on human volunteers to test brain implants that give paralyzed people the ability to control a computer cursor with thought alone. Other neuroengineers have been able to achieve even more spectacular results on monkeys. At the University of Pittsburgh, for example, monkeys can use their thoughts to feed themselves with a robotic arm.

These are all promising results, but brain implants are still fairly crude. When the electrodes release pulses of electricity they can't target particular neurons; they just blast whatever neurons happen to be nearby. The best electrodes for recording brain activity, meanwhile, can pick up only a tiny portion of the chatter in the brain because engineers can implant only a few dozen electrodes in a single person.

Making matters worse, almost all of these implants have to be rigged to wires that snake out of the skull and draw a lot of power, relatively speaking, limiting battery lifetime. Surgeons have also found that the brain attacks these electrodes, covering them with a protective coat of cells that can render

---

*Boyden gave a talk at the summit, unveiling a particularly stunning experiment on congenitally blind mice: They infected the animals with sight.*

---

them useless. All of these problems mean you can't expect to carry a brain implant for life. Dipping into a person's brain from time to time to swap out implants and batteries would not just be expensive but would also pose the risk of infection.

But none of these challenges is necessarily a showstopper. Scientists are working on new designs that can allow brain implants to shrink in size, use less power and deliver better performance. In 2009, for example, a team of scientists at MIT reported how they had implanted a tiny electrode into the brain of zebra finches. The birds could fly freely around an enclosure. But with the press of a button, the scientists could wirelessly transmit a signal to the song-producing region of the bird's brain. The bird instantly stopped singing.

Up to now, many brain implant studies have focused on delivering or receiving electrical signals. Boyden's approach is novel. Isolating a particular gene in algae, it turns out, can make light-emitting devices possible. That's because some algae have channels on special membranes in their cells that respond to light of certain colors by opening up, allowing charged particles to move in or out of the membrane. A few years back Boyden wondered

if he could insert those channels into neurons and use them as an optical switch. Hit a neuron with light and its channels will open, triggering a signal.

Boyden and his colleagues pinpointed the gene for the channel, inserted it into viruses and then put the engineered viruses into a dish of neurons. The viruses infected the neurons and, along with their own genes, inserted the light-channel gene from the algae. The virus was harmless, so the neurons did not suffer from the infection, but they started using the algae gene to build channels of their own. Boyden then exposed the neurons to a flickering blue light. The neurons responded by crackling with spikes of electricity. "We started playing around with it and we got light-driven spikes almost on the first try," he said. "It was an idea whose time had come."

Boyden and his colleagues published that experiment in 2005, and since then he has expanded his neuroengineer's tool kit dramatically. "We grow organisms to screen for new molecules," he said. Their discovery of new light-sensitive channels in algae, bacteria and fungi have allowed them to engineer neurons that respond to a rainbow of light.

Boyden has neurons performing new tricks. He can get them to produce a voltage spike in response to light or to go completely quiet. He can flash a particular pattern of lights to trigger signals. He can also target his channels to different types of neurons by adding different genetic handles to the channel DNA. The genes get inserted into lots of cells, but they get switched on in only one kind of neuron. Flashing different colors of light, he can switch on and shut down different groups of neurons all at once.

Now he is starting to see how his engineered neurons behave in real brains rather than on glass slides. One virus he selected for his experiments, known as an adeno-associated virus, has proven to be promising in human gene-therapy trials in other labs. It has safely delivered genes into the bodies of more than 600 people so far (though some of the genes have produced unintended negative side effects). Also, last April Boyden and his colleagues reported they were able to successfully infect certain neurons in the brains of monkeys without causing harm to the animals. The scientists inserted an optical fiber into the monkey brains and were able to switch the neurons on with flashes of light, just as they do on the slides.

Boyden gave a talk at the Singularity Summit, unveiling a particularly stunning experiment he and his colleague Alan Horsager have run at the University of Southern California on congenitally blind mice: They infected the animals with sight.

The mice were blind thanks to mutations in the light-receptor genes in their retinas. The team wondered if they could make those neurons sensitive to light again. They loaded genes for light-sensitive channels onto viruses and injected them into the mice. The genes were targeted for certain retinal neurons that communicate with the brain. Boyden gave the mice enough time to incorporate the genes into their eyes and, he hoped, make the channels in their neurons. Since mice can't read eye charts out loud, he and his colleagues had to use a behavioral test to see if their eyes

were working. They put the mice into a little pool with barriers arranged into a maze. At one end of the pool the mice could get out of the water by climbing onto an illuminated platform. Regular mice quickly followed the light to the platform, while blind mice swam around randomly. The mice infected with neuron channels headed for the exit far more often than chance and almost as often as the healthy mice. Boyden and his colleagues have founded a company, Eos Neuroscience, to see if they can use this gene therapy to help restore some eyesight to humans.

Ultimately, though, Boyden also wants to install these light-sensitive receptors on neurons deep in the brain. Then, with a flash of light, he can make certain neurons fire or go silent. Boyden and his colleagues have built a peculiar gadget that looks like a miniature glass pipe organ. At the base of each fiber a diode can produce an intense flash of light. Boyden envisions implanting this array into people's brains and then wirelessly programming it to produce a rapid-fire rainbow pattern of light.

His far-off goal is to help treat medical disorders with these implants, and he doesn't give much thought to the possibility of people using implants to enhance their brains, the way they do now with Adderall. Brain implants certainly inspire cool scenarios—what if someone wanted to see in ultraviolet or operate a jet fighter with thought alone?—but Boyden has the luxury of not having to worry about those ethical matters. After all, it's one thing to open a jar and pop a pill; it's quite another to undergo

brain surgery. "I think the invasive techniques won't be used for augmentation for a long time to come," he said.

As the conversation continued we got to talking about Lasik. Once there was a time when having a laser shot into your eye to fix myopia was the ophthalmological equivalent of Russian roulette. "Forty years ago it was daring," said Boyden. "Now there are clinics that do hundreds of these day in and day out."

I opened my Summit schedule to see what was next.

9:35 A.M.: Technical Road Map for Whole Brain Emulation.

10:00 A.M.: The Time Is Now: We Need Whole Brain Emulation.

This should be interesting, I thought.

Beyond drugs and prosthetics is whole brain emulation. If you haven't heard of it, here's a quick definition from a 2008 paper by Nick Bostrom and Anders Sandberg, two scientists at the University of Oxford: "The basic idea is to take a particular brain, scan its structure in detail and construct a software model of it so faithful to the original that, when run on appropriate hardware, it will behave in essentially the same way as the original brain."

At the Singularity Summit, Sandberg strode onto the stage, wearing a science-fair smile and a bright red tie, to explain what it would take to reach that goal. First, scientists would have to decide exactly how much detail they'd need. Would they have to track every single molecule? Would a

write software that can turn all the data into a three-dimensional model and then boot up this virtual brain. Sandberg doesn't think a computer would have to calculate the activity of the neurons atom by atom. Neurons are fairly predictable, so it's already possible to build models that behave a lot like real neurons. Mikael Djurfeldt, a scientist at Sweden's Royal Institute of Technology, and his colleagues have succeeded in modeling one particular kind of neuron cluster known as a cortical column. They created a model of 22 million neurons joined together by 11 billion synapses. When their imitation neurons started to talk to each other, they behaved a lot like real cortical columns. Of course it's important to remember there are about 100 billion neurons in the human brain—several hundred times more than Djurfeldt has simulated. And even if scientists do manage to simulate 100 billion neurons, they'll also need to give the simulated brain a simulated world.


Sandberg has made some rough calculations of how much computing power that would demand and how fast computer power is rising, and he's fairly confident a whole-brain emulation will be possible in a matter of a few decades. Exactly how a whole-brain emulation would behave, he isn't sure.

"I don't know whether a complete simulation of a brain, one to one, would actually produce a mind," he said. "I find it pretty likely, but I don't have evidence for that. I want to test it."

Sandberg exited stage right and was replaced at the podium by Randal Koene, a neuroscientist at the European

technology firm Fatronik-Tecnalia. Koene offered some reasons for why anyone would want to work so hard to make a whole-brain emulation in the first place. Even if it behaved like a generic human brain rather than my or your brain in particular, scientists could still use it to run marvelous new kinds of experiments. They might test drugs for depression, Parkinson's and other disorders. Koene is also a strong advocate of so-called mind uploading—the possibility of not just running a brain-like simulation on a computer but actually transferring a person's mind into a machine. To him, it is the liberation of our species. "We must free the mind," said Koene.

For a little ground-truthing I called Olaf Sporns, a neuroscientist at Indiana University. "This is not going to happen," he said.



rough approximation of all 100 billion neurons suffice? Sandberg suspected scientists would need a scan of a brain that could provide details down to a few nanometers (less than a millionth of an inch). Today researchers at Texas A&M have figured out how to take images of the brain at a resolution of just 160 nanometers, but they've scanned only a rice-grain-size piece of mouse brain in any one trial. To scan the brain tissue the scientists must stain it with color-producing chemicals, dunk it in plastic for hardening and then shave away one layer at a time. For now brain emulation is a zero-sum game.

But let's assume for the moment scientists can scan an entire human brain at nanometer scale. It turns out the hard work is just beginning. Researchers will then have to

Sporns is in a good position to judge. He and his colleagues have carried out just about the closest thing to whole-brain emulation given today's technology. They map human brains using a high-resolution method called diffusion spectrum imaging. They chart the long fibers that link regions of the brain together like computers on the Internet. Sporns and his colleagues have analyzed the connections between 1,000 regions and have found the brain's network is organized according to some of the same rules of other large networks—including the Internet. For example, several regions act as hubs, while most regions are connected to only a few others. Sporns and his colleagues created a computer model of this brain network and let each region produce signals that could spread down the fibers. They found their simulation of a brain at rest produced distinctive waves that spread back and forth around the entire brain similar to the way waves spread across our own.

Whole-brain emulations will become more sophisticated in the future, said Sporns, but he finds it ridiculous to expect them to be able to capture an individual's mind. In fact, mind uploading is a distraction from the truly revolutionary impact whole-brain emulations will have. By experimenting with them, researchers may discover some of the laws for building thinking networks. It may turn out human brains can work only if their networks have certain arrangements. "It's like learning the laws of physics when you want to build an airplane. It helps," said Sporns.

Discovering those laws may allow computer scientists to finally build machines that have mental processes similar to ours. Sporns thinks scientists are already moving in that direction. IBM, for example, now has a contract with the military to "fabricate a multichip neural system of about 108 neurons and instantiate into a robotic platform performing at 'cat' level."

"That is going to happen," said Sporns.



"Yes, she has a divine right...and her left one ain't bad either!"

This is the best I can manage for skeptics. Uploading your mind is science fiction. But endowing robots with a humanlike cognition? Only a matter of time.

"I hope you didn't believe a word of that last talk."

I stood in a long line for the men's room during a coffee break. The women's room next door was empty. I thought how only at a meeting like the Singularity Summit would I find myself in this situation. When I heard someone talking to me, I turned to see a tousle-haired psychologist named Gary Marcus.

Good, I thought. Maybe Marcus would demolish the Singularity and leave nothing behind but smoking wreckage.

Marcus, who teaches at New York University, has spent years studying computation. Computers are good at certain kinds of computations such as sorting things into categories. But they're not good at the things we do effortlessly, such as generating rules from experience.

Marcus was annoyed by one of the talks at the Summit, in which a computer scientist promised that humanlike artificial intelligence was nigh. "Figuring this stuff out in 10 years—I don't believe it," he said.

I expected some serious curmudgeonliness when Marcus delivered his talk the following day. In his 2008 book, *Kluge*, Marcus explores design flaws in the human brain. We don't simply store information on a hard disk, for example, but embed it in a web of associations. That's why memories may escape us until something—perhaps the taste of a cookie—brings back the right associations. Marcus explains how these quirks are locked into our brains thanks to our evolutionary history. We did not evolve to be computers but animals that could learn to find food and avoid being eaten.

So I imagined Marcus would declare the human brain unimprovable, but he didn't. He stood up, explained the shortfalls of human memory and suggested memory would be a good place to start improving the human brain.

I called Marcus later and told him I was surprised. "Human enhancement is a real possibility," he replied. He thought a powerful way to enhance the brain would be with "cognitive prosthetics," a kind of onboard iPhone. The only challenge would be to decipher our brains' code well enough to let the iPhone talk to our brains. Marcus didn't see any reason scientists wouldn't eventually figure that out.

Like Sporns, Marcus agreed whole-brain emulations might turn out to be most valuable for what they reveal about the nature of intelligence. In the end that's what left Marcus worried. We just might succeed too well and program a machine with so much intelligence it can boost its own intelligence by itself.

"There are going to be machines that are cleverer than we are," he said.

It is time, Marcus said, to start planning for that world. It doesn't matter whether we live to see it or not; our grandchildren or great-grandchildren might. We owe it to them to get ready—if not for the Singularity then at least for a world different from our own.

