SCIENTIFIC AMERICAN

Permanent Address: **http://www.scientificamerican.com/article.cfm?id=e-zimmer-can-you-live-forever**

Can You Live Forever? Maybe Not--But You Can Have Fun Trying

In this chapter from his new e-book, journalist Carl Zimmer tries to reconcile the visions of techno-immortalists with the exigencies imposed by real-world biology

By Carl Zimmer | Wednesday, December 22, 2010 | 14 comments

Editor's Note: Carl Zimmer, author of this month's article, "100 Trillion Connections," has just brought out a much-acclaimed e-book, Brain Cuttings: 15 Journeys Through the Mind *(Scott & Nix), that compiles a series of his writings on neuroscience. In this chapter, adapted from an article that was first published in* Playboy, *Zimmer takes the reader on a tour of the 2009 Singularity Summit in New York City. His ability to contrast the fantastical predictions of speakers at the conference with the sometimes more skeptical assessments from other scientists makes his account a fascinating read.*

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. Someone reroutes your vision through cameras. Someone stores your memories on a net of microprocessors. Step by step your metamorphosis continues until at last the transfer is complete. As engineers get to work boosting the performance of your electronic mind so you can now think as a god, a nurse heaves your fleshy brain 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, 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.

"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 particular 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 think this possibility was anything to be scared of. "Let's call it the Buddhist view," he said. Every day, we lose consciousness as we fall asleep and then regain it the next morning.

"Each waking is really like a new dawn that's a bit like the commencement of a new person," Chalmers said. "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, that may be because 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 a part of it.

The Singularity is more than just hypothetic milestone in history. It's 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. Writing about the brain and other scientific subjects had given me a strong immune defense against hype. The Singularity, with all its promises of a technorapture, seems tailor-made to bring out the worst in people like me. The writer John Horgan wrote a devastating essay about the Singularity in 2009 called "Science Cult."

Horgan acknowledged 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—worries that so many people, smart people, are taking Kurzweil's sci-fi fantasies seriously," he wrote. "The last thing humanity needs right now is an apocalyptic cult masquerading as science."

I decided to check out the Singularity for myself. 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 decided to visit to researchers working on the type of technology that people such as Kurzweil consider the steppingstones to the Singularity. Not one of them takes Kurzweil's own 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 a reckless promise of false hope to the afflicted.

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 humans 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 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 success of antibiotics led some doctors to declare an end to the age of infectious diseases. The engineers who founded NASA were sure that within a few generations 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 only from their own research, but from research on the brain. Neuroscientists were discovering that they could apply the concepts of computer science to the brain. It seemed 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 wrote. "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 argued that within a few decades computers will be as intelligent as humans, if not more so.

As the years passed, Kurzweil's predictions grew more extreme. In his 1999 book, *The Age of Spiritual Machines*, he imagined life in 2099. "The number of software-based humans vastly exceeds those still using native neuron-cell-based computation," he wrote. In 2005 Kurzweil brought Vinge's term to wide attention in *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 brains and bodies."

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

Kurzweil 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 wirelessly communicating with one another and the World Wide Web 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." A decade later, medicine had not lived up to his prophecy. In 2006 —the most recent year for which statistics are available—829,072 people died in the United States of heart disease. Its true that the death rate from heart disease is somewhat lower now than in 1950, but that drop is due mainly to low-tech measures such as getting people to stop smoking. Deaths from cancer, meanwhile, only dropped ten percent between 1975 and 2010. Cancer claimed 562,875 lives in the United States in 2007.

These failed predictions reveal a weakness at the heart of Kurzweil's forecasts: scientific understanding doesn't advance in lockstep with increases in 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 U.S. Influenza Genome Sequencing Project has collected the complete genomes of 4,087 viruses from around the world, and that number is rising close to exponentially.

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. Fortunately, the flu turned out to be no worse than the typical seasonal flu. But that's little comfort for the families of the 12,470 people in the United States who died of the new flu strain in 2009 and 2010.

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 many 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. Sequencing the human genome was promised by some to 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.

None of these facts tempered Kurzweil's promises of immortality. In his 2010 book, _Transcend: Nine Steps to Living Well Forever_ (co-authored with a doctor named Terry Grossman), he offered tips for keeping yourself alive for twenty years, so that you can take advantage of all the life-extending technology that will be invented by then, which will keep you alive for another twenty years, and so on and so on, bootstrapping yourself through the centuries.

For all the talk of transcendence, a lot of the tips in _Transcend_ are profoundly banal. Maintain a regular exercise routine! Lose excess body weight! _Transcend_ also advises readers to take vitamins and supplements, often relying on shaky evidence. Kurzweil recommends 500 milligrams of vitamin C a day, for example, writing that "vitamin C is a potent free-radical scavenger (anti-oxidant) that has been shown to be of value in preventing and treating cancer." As evidence, he cites a small study in 2002 on the effect of vitamin C on precancerous cells living in culture dishes.

Kurzweil doesn't mention any of the large-scale investigations of the benefits that vitamin C brings to the health of actual people. In December 2009, scientists at Harvard Medical School published one such study in _The Journal of the National Cancer Institute_. For a decade, they tracked 7627 women. Half took 500 milligrams of vitamin C a day, and half took placebo. Vitamin C, the Harvard scientists concluded, "offers no overall benefits in the primary prevention of total cancer incidence or cancer mortality."

You won't read about findings like that in _Transcend_. But you can page through 66 pages of Singularity-friendly recipes, like Fiesta

Omelets and Cajun Salmon Fillets. Whether you'll have to eat them for eternity, Kurzweil doesn't say.

***

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 reach the market. Some are intended to slow the disappearance of memories in people with Alzheimer's disease; others may boost cognition in people with impairments. Farah expects there will be people —maybe a lot of them—who will take these drugs in the hopes of improving an already healthy brain, not to fix a deficit.

In December 2008 Farah and a group of fellow neuroscientists and bioethicists wrote a commentary in *Nature* on cognitive enhancement. Their verdict was that it was okay. "Mentally competent adults should be able to engage in cognitive enhancement using drugs," they wrote. It's impractical, they argued, to try to draw a line between treating a disease and enhancing a healthy brain.

What really matters, Farah said, is for scientists to figure out how much enhancement these drugs can actually 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." Farah is now measuring the performance of students on Adderall and placebos to see if that's actually the case.

Someday, Farah suspects, we may look back at the cognitive enhancers of the early twentieth century as the first step towards a Singularity-like era. "I think this growing practice may be softening us up to accept more drastic brain modifications down the line," she said.

***

"Here are some critters," said Ed Boyden. The serene young leader of the Synthetic Neurobiology Group at mit stood in a laboratory strewn with miniature pieces of electronics. Dishes of neurons were positioned under microscopes. And on the lab benches were also bright green flasks of algae, one of which Boyden was sloshing as he spoke with me about the future of the brain.

I came to Boyden's lab after seeing him give a remarkable talk at the Singularity Summit. Boyden belongs to the young field of neuroengineering, which seeks to restore damaged brains by implanting tiny on-board computers and electrodes. Boyden's research may take change the rules of the neuroengineering game by allowing scientists to manipulate the brain not with electricity, but with light. But in order to do that, Boyden was borrowing genes from the algae he was holding in the green flask. If he succeeds, people will become part machine, and part algae too.

The logic behind brain implants is simple. Many disorders, such 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 emerged in the 1970s. They 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. As of 2010, 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.

Scientists are working not just on input devices, but output ones as well. 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 with 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. They are power-hungry, shortening the life of their batteries. Surgeons have also found that the brain attacks these electrodes, covering them with a protective coat of cells that can render them useless. These problems mean that no one in 2010 can expect to carry a brain implant for life.

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 succeeded in implanting a wireless electrode into a zebra finch. 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.

In 2003, Boyden wondered if there might be a way to use light instead of electricity to communicate with neurons. He knew that algae have light-sensitive channels on the surface of their cells. When a photon hit the channels, they opened, up, allowing charged particles to flow in or out of the cell. Boyden imagined putting those channels on the surface of neurons. Hit a neuron with light and its new channels would open, triggering a signal.

Boyden and his colleagues pinpointed a gene for a channel in a species of algae and inserted it into viruses. They then infected neurons with the engineered viruses. Along with their own genes the viruses inserted the light-channel gene from the algae. Because the viruses were harmless, the neurons did not suffer from the infection. Instead they started using the algae gene to build channels of their own.

Boyden exposed the infected neurons to a flickering blue light. The neurons responded by crackling with spikes of voltage. "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."

In 2005 Boyden and his colleagues published that experiment, and since then he has expanded his neuroengineer's toolkit. "We grow organisms to screen for new molecules," he told me. They went on to find new genes for receptors sensitive to different wavelengths of light. Now their engineered neurons can respond to a rainbow of signals. Boyden can get the neurons 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 give different types of neurons different channels by adding specific 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.

I visited Boyden as he was starting to see how his engineered neurons behave in real brains rather than in petri dishes. One virus he selected for his experiments, known as an adeno-associated virus, proved to be promising in human gene-therapy trials in other labs. As of 2010, it had safely delivered genes into the bodies of more than 600 people (though some of the genes have produced unintended side effects). Boyden and his colleagues successfully infected 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, just as they do in a petri dish.

At Boyden's talk at the Singularity to Summit he unveiled 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 the retinal neurons missing their own light receptors. Boyden gave the mice 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. When Boyden and his colleagues put the mice infected with neuron channels in the pool, they headed for the exit almost as often as the healthy mice. As far as Boyden and his colleagues can tell, the mice can see again.

Boyden hopes to treat blind people by putting the channels in their eyes. And ultimately, he wants to install these light-sensitive receptors on neurons deep in the brain as well. To manipulate those deep-brain neurons, he'll need a new kind of electrode. On my visit to his lab, Boyden showed me his newest prototype: 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.

In the far future, Boyden can imagine treating people with Parkinson's, depression, and other disorders with his implants. 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 feed the imagination. Could people with implants see the world in ultraviolet? Could they control a jet fighter with thought alone? Could they become slaves of a computer?

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 says.

And yet, as we talked, the conversation drifted to the subject of 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," says 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 AM: Technical Road Map for Whole Brain Emulation.
10:00 AM: The Time Is Now: We Need Whole Brain Emulation.

This should be interesting, I thought.

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 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. Researchers will then have to 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 clusters of neurons known as cortical columns. 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.

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 made an exit stage right, 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 just behaved like a generic human brain rather than my brain or yours 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.

Sporns is in a good position to judge. He and his colleagues have carried out just about the closest thing to whole-brain emulation that scientists can manage in the early twenty-first century. They use a high-resolution method called diffusion spectrum imaging to map the long fibers that link regions of the brain together. 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.

In 2009, 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, Sporns told me, but he found 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 only work 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. In 2009, for example, iBm won 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. 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 computation 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 ten years—I don't believe it," he said.

I expected some serious curmudgeonliness it came time for Marcus to deliver his own talk the following day. In his 2008 book, *Kluge*, Marcus explored design flaws in the human brain. Our memory is inevitably faulty, for example, because we don't simply store information on a hard disk. Instead, we 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. He came out onto the stage and began to explain the shortfalls of human memory. So far, so good. But then he proceeded to explain why memory's failings meant that it was a good place to start improving the human brain. I had lost another skeptic.

I called Marcus a few days later and told him that he had surprised me. "Human enhancement is a real possibility," he replied. He envisioned a coming age of cognitive prosthetics, implanted devices that would work like a kind of onboard iPhone. The only challenge would be to decipher the human brain's signals well enough to let the iPhone talk to it. Marcus didn't see any reason scientists wouldn't eventually figure that out.

Marcus agreed with Sporns that the true value of whole- brain emulations might turn out to be what they teach us about the nature of intelligence. But that possibility worried Marcus deeply. As we use what we learn about the brain to program new machines, we may succeed too well. We may build a machine with so much intelligence it can start expanding its own intelligence without our help.

"There are going to be machines that are cleverer than we are," Marcus said. It is time, he believed, to start planning for that world. It doesn't matter whether those of us alive today get to see it or not. Even if Fiesta Omelets don't keep us alive for decades, our grandchildren or great-grandchildren may still have to cope with our too-clever creations. We owe it to them to get ready—if not for the Singularity, then at least for a life different from our own.