

2045: The Year Man Becomes Immortal

From TIME magazine. By Lev Grossman Thursday, Feb. 10, 2011

On Feb. 15, 1965, a diffident but self-possessed high school student named Raymond Kurzweil appeared as a guest on a game show called I've Got a Secret. He was introduced by the host, Steve Allen, then he played a short musical composition on a piano. The idea was that Kurzweil was hiding an unusual fact and the panelists — they included a comedian and a former Miss America — had to guess what it was.

On the show, the beauty queen did a good job of grilling Kurzweil, but the comedian got the win: the music was composed by a computer. Kurzweil got \$200.

Kurzweil then demonstrated the computer, which he built himself — a desk-size affair with loudly clacking relays, hooked up to a typewriter. The panelists were pretty blasé about it; they were more impressed by Kurzweil's age than by anything he'd actually done. They were ready to move on to Mrs. Chester Loney of Rough and Ready, Calif., whose secret was that she'd been President Lyndon Johnson's first-grade teacher.

But Kurzweil would spend much of the rest of his career working out what his demonstration meant. Creating a work of art is one of those activities we reserve for humans and humans only. It's an act of self-expression; you're not supposed to be able to do it if you don't have a self. To see creativity, the exclusive domain of humans, usurped by a computer built by a 17-year-old is to watch a line blur that cannot be unblurred, the line between organic intelligence and artificial intelligence.

That was Kurzweil's real secret, and back in 1965 nobody guessed it. Maybe not even him, not yet. But now, 46 years later, Kurzweil believes that we're approaching a moment when computers will become intelligent, and not just intelligent but more intelligent than humans. When that happens, humanity — our bodies, our minds, our civilization — will be completely and irreversibly transformed. He believes that this moment is not only inevitable but imminent. According to his calculations, the end of human civilization as we know it is about 35 years away.

Computers are getting faster. Everybody knows that. Also, computers are getting faster faster — that is, the rate at which they're getting faster is increasing.

True? True.

So if computers are getting so much faster, so incredibly fast, there might conceivably

come a moment when they are capable of something comparable to human intelligence. Artificial intelligence. All that horsepower could be put in the service of emulating whatever it is our brains are doing when they create consciousness — not just doing arithmetic very quickly or composing piano music but also driving cars, writing books, making ethical decisions, appreciating fancy paintings, making witty observations at cocktail parties.

If you can swallow that idea, and Kurzweil and a lot of other very smart people can, then all bets are off. From that point on, there's no reason to think computers would stop getting more powerful. They would keep on developing until they were far more intelligent than we are. Their rate of development would also continue to increase, because they would take over their own development from their slower-thinking human creators. Imagine a computer scientist that was itself a super-intelligent computer. It would work incredibly quickly. It could draw on huge amounts of data effortlessly. It wouldn't even take breaks to play Farmville.

Probably. It's impossible to predict the behavior of these smarter-than-human intelligences with which (with whom?) we might one day share the planet, because if you could, you'd be as smart as they would be. But there are a lot of theories about it. Maybe we'll merge with them to become super-intelligent cyborgs, using computers to extend our intellectual abilities the same way that cars and planes extend our physical abilities. Maybe the artificial intelligences will help us treat the effects of old age and prolong our life spans indefinitely. Maybe we'll scan our consciousnesses into computers and live inside them as software, forever, virtually. Maybe the computers will turn on humanity and annihilate us. The one thing all these theories have in common is the transformation of our species into something that is no longer recognizable as such to humanity circa 2011. This transformation has a name: the Singularity.

The difficult thing to keep sight of when you're talking about the Singularity is that even though it sounds like science fiction, it isn't, no more than a weather forecast is science fiction. It's not a fringe idea; it's a serious hypothesis about the future of life on Earth. There's an intellectual gag reflex that kicks in anytime you try to swallow an idea that involves super-intelligent immortal cyborgs, but suppress it if you can, because while the Singularity appears to be, on the face of it, preposterous, it's an idea that rewards sober, careful evaluation.

People are spending a lot of money trying to understand it. The three-year-old Singularity University, which offers inter-disciplinary courses of study for graduate students and executives, is hosted by NASA. Google was a founding sponsor; its CEO and co-founder Larry Page spoke there last year. People are attracted to the Singularity for

the shock value, like an intellectual freak show, but they stay because there's more to it than they expected. And of course, in the event that it turns out to be real, it will be the most important thing to happen to human beings since the invention of language.

The Singularity isn't a wholly new idea, just newish. In 1965 the British mathematician I.J. Good described something he called an "intelligence explosion":

Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an "intelligence explosion," and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the last invention that man need ever make.

The word singularity is borrowed from astrophysics: it refers to a point in space-time — for example, inside a black hole — at which the rules of ordinary physics do not apply. In the 1980s the science-fiction novelist Vernor Vinge attached it to Good's intelligence-explosion scenario. At a NASA symposium in 1993, Vinge announced that "within 30 years, we will have the technological means to create super-human intelligence. Shortly after, the human era will be ended."

By that time Kurzweil was thinking about the Singularity too. He'd been busy since his appearance on *I've Got a Secret*. He'd made several fortunes as an engineer and inventor; he founded and then sold his first software company while he was still at MIT. He went on to build the first print-to-speech reading machine for the blind — Stevie Wonder was customer No. 1 — and made innovations in a range of technical fields, including music synthesizers and speech recognition. He holds 39 patents and 19 honorary doctorates. In 1999 President Bill Clinton awarded him the National Medal of Technology.

But Kurzweil was also pursuing a parallel career as a futurist: he has been publishing his thoughts about the future of human and machine-kind for 20 years, most recently in *The Singularity Is Near*, which was a best seller when it came out in 2005. A documentary by the same name, starring Kurzweil, Tony Robbins and Alan Dershowitz, among others, was released in January. (Kurzweil is actually the subject of two current documentaries. The other one, less authorized but more informative, is called *The Transcendent Man*.) Bill Gates has called him "the best person I know at predicting the future of artificial intelligence

In real life, the transcendent man is an unimposing figure who could pass for Woody Allen's even nerdier younger brother. Kurzweil grew up in Queens, N.Y., and you can still hear a trace of it in his voice. Now 62, he speaks with the soft, almost hypnotic calm of someone who gives 60 public lectures a year. As the Singularity's most visible champion, he has heard all the questions and faced down the incredulity many, many times before. He's good-natured about it. His manner is almost apologetic: I wish I could bring you less exciting news of the future, but I've looked at the numbers, and this is what they say, so what else can I tell you?

Kurzweil's interest in humanity's cyborganic destiny began about 1980 largely as a practical matter. He needed ways to measure and track the pace of technological progress. Even great inventions can fail if they arrive before their time, and he wanted to make sure that when he released his, the timing was right. "Even at that time, technology was moving quickly enough that the world was going to be different by the time you finished a project," he says. "So it's like skeet shooting — you can't shoot at the target." He knew about Moore's law, of course, which states that the number of transistors you can put on a microchip doubles about every two years. It's a surprisingly reliable rule of thumb. Kurzweil tried plotting a slightly different curve: the change over time in the amount of computing power, measured in MIPS (millions of instructions per second), that you can buy for \$1,000.

As it turned out, Kurzweil's numbers looked a lot like Moore's. They doubled every couple of years. Drawn as graphs, they both made exponential curves, with their value increasing by multiples of two instead of by regular increments in a straight line. The curves held eerily steady, even when Kurzweil extended his backward through the decades of pretransistor computing technologies like relays and vacuum tubes, all the way back to 1900.

Kurzweil then ran the numbers on a whole bunch of other key technological indexes — the falling cost of manufacturing transistors, the rising clock speed of microprocessors, the plummeting price of dynamic RAM. He looked even further afield at trends in biotech and beyond — the falling cost of sequencing DNA and of wireless data service and the rising numbers of Internet hosts and nanotechnology patents. He kept finding the same thing: exponentially accelerating progress. "It's really amazing how smooth these trajectories are," he says. "Through thick and thin, war and peace, boom times and recessions." Kurzweil calls it the law of accelerating returns: technological progress happens exponentially, not linearly.

Then he extended the curves into the future, and the growth they predicted was so phenomenal, it created cognitive resistance in his mind. Exponential curves start slowly,

then rocket skyward toward infinity. According to Kurzweil, we're not evolved to think in terms of exponential growth. "It's not intuitive. Our built-in predictors are linear. When we're trying to avoid an animal, we pick the linear prediction of where it's going to be in 20 seconds and what to do about it. That is actually hardwired in our brains."

Here's what the exponential curves told him. We will successfully reverse-engineer the human brain by the mid-2020s. By the end of that decade, computers will be capable of human-level intelligence. Kurzweil puts the date of the Singularity — never say he's not conservative — at 2045. In that year, he estimates, given the vast increases in computing power and the vast reductions in the cost of same, the quantity of artificial intelligence created will be about a billion times the sum of all the human intelligence that exists today.

The Singularity isn't just an idea. it attracts people, and those people feel a bond with one another. Together they form a movement, a subculture; Kurzweil calls it a community. Once you decide to take the Singularity seriously, you will find that you have become part of a small but intense and globally distributed hive of like-minded thinkers known as Singularitarians.

Not all of them are Kurzweilians, not by a long chalk. There's room inside Singularitarianism for considerable diversity of opinion about what the Singularity means and when and how it will or won't happen. But Singularitarians share a worldview. They think in terms of deep time, they believe in the power of technology to shape history, they have little interest in the conventional wisdom about anything, and they cannot believe you're walking around living your life and watching TV as if the artificial-intelligence revolution were not about to erupt and change absolutely everything. They have no fear of sounding ridiculous; your ordinary citizen's distaste for apparently absurd ideas is just an example of irrational bias, and Singularitarians have no truck with irrationality. When you enter their mind-space you pass through an extreme gradient in worldview, a hard ontological shear that separates Singularitarians from the common run of humanity. Expect turbulence.

In addition to the Singularity University, which Kurzweil co-founded, there's also a

Singularity Institute for Artificial Intelligence, based in San Francisco. It counts among its advisers Peter Thiel, a former CEO of PayPal and an early investor in Facebook. The institute holds an annual conference called the Singularity Summit. (Kurzweil co-founded that too.) Because of the highly interdisciplinary nature of Singularity theory, it attracts a diverse crowd. Artificial intelligence is the main event, but the sessions also cover the galloping progress of, among other fields, genetics and nanotechnology.

At the 2010 summit, which took place in August in San Francisco, there were not just computer scientists but also psychologists, neuroscientists, nanotechnologists, molecular biologists, a specialist in wearable computers, a professor of emergency medicine, an expert on cognition in gray parrots and the professional magician and debunker James "the Amazing" Randi. The atmosphere was a curious blend of Davos and UFO convention. Proponents of seasteading — the practice, so far mostly theoretical, of establishing politically autonomous floating communities in international waters — handed out pamphlets. An android chatted with visitors in one corner.

After artificial intelligence, the most talked-about topic at the 2010 summit was life extension. Biological boundaries that most people think of as permanent and inevitable Singularitarians see as merely intractable but solvable problems. Death is one of them. Old age is an illness like any other, and what do you do with illnesses? You cure them. Like a lot of Singularitarian ideas, it sounds funny at first, but the closer you get to it, the less funny it seems. It's not just wishful thinking; there's actual science going on here.

For example, it's well known that one cause of the physical degeneration associated with aging involves telomeres, which are segments of DNA found at the ends of chromosomes. Every time a cell divides, its telomeres get shorter, and once a cell runs out of telomeres, it can't reproduce anymore and dies. But there's an enzyme called telomerase that reverses this process; it's one of the reasons cancer cells live so long. So why not treat regular non-cancerous cells with telomerase? In November, researchers at Harvard Medical School announced in *Nature* that they had done just that. They administered telomerase to a group of mice suffering from age-related degeneration. The damage went away. The mice didn't just get better; they got younger.

Aubrey de Grey is one of the world's best-known life-extension researchers and a Singularity Summit veteran. A British biologist with a doctorate from Cambridge and a famously formidable beard, de Grey runs a foundation called SENS, or Strategies for Engineered Negligible Senescence. He views aging as a process of accumulating damage, which he has divided into seven categories, each of which he hopes to one day address using regenerative medicine. "People have begun to realize that the view of aging being something immutable — rather like the heat death of the universe — is simply ridiculous," he says. "It's just childish. The human body is a machine that has a bunch of functions, and it accumulates various types of damage as a side effect of the normal function of the machine. Therefore in principle that damage can be repaired periodically. This is why we have vintage cars. It's really just a matter of paying attention. The whole of medicine consists of messing about with what looks pretty inevitable until you figure out how to make it not inevitable."

Kurzweil takes life extension seriously too. His father, with whom he was very close, died of heart disease at 58. Kurzweil inherited his father's genetic predisposition; he also developed Type 2 diabetes when he was 35. Working with Terry Grossman, a doctor who specializes in longevity medicine, Kurzweil has published two books on his own approach to life extension, which involves taking up to 200 pills and supplements a day. He says his diabetes is essentially cured, and although he's 62 years old from a chronological perspective, he estimates that his biological age is about 20 years younger.

But his goal differs slightly from de Grey's. For Kurzweil, it's not so much about staying healthy as long as possible; it's about staying alive until the Singularity. It's an attempted handoff. Once hyper-intelligent artificial intelligences arise, armed with advanced nanotechnology, they'll really be able to wrestle with the vastly complex, systemic problems associated with aging in humans. Alternatively, by then we'll be able to transfer our minds to sturdier vessels such as computers and robots. He and many other Singularitarians take seriously the proposition that many people who are alive today will wind up being functionally immortal.

It's an idea that's radical and ancient at the same time. In "Sailing to Byzantium," W.B. Yeats describes mankind's fleshly predicament as a soul fastened to a dying animal. Why not unfasten it and fasten it to an immortal robot instead? But Kurzweil finds that life extension produces even more resistance in his audiences than his exponential growth

curves. "There are people who can accept computers being more intelligent than people," he says. "But the idea of significant changes to human longevity — that seems to be particularly controversial. People invested a lot of personal effort into certain philosophies dealing with the issue of life and death. I mean, that's the major reason we have religion."

Of course, a lot of people think the Singularity is nonsense — a fantasy, wishful thinking, a Silicon Valley version of the Evangelical story of the Rapture, spun by a man who earns his living making outrageous claims and backing them up with pseudoscience. Most of the serious critics focus on the question of whether a computer can truly become intelligent.

The entire field of artificial intelligence, or AI, is devoted to this question. But AI doesn't currently produce the kind of intelligence we associate with humans or even with talking computers in movies — HAL or C3PO or Data. Actual AIs tend to be able to master only one highly specific domain, like interpreting search queries or playing chess. They operate within an extremely specific frame of reference. They don't make conversation at parties. They're intelligent, but only if you define intelligence in a vanishingly narrow way. The kind of intelligence Kurzweil is talking about, which is called strong AI or artificial general intelligence, doesn't exist yet.

Why not? Obviously we're still waiting on all that exponentially growing computing power to get here. But it's also possible that there are things going on in our brains that can't be duplicated electronically no matter how many MIPS you throw at them. The neurochemical architecture that generates the ephemeral chaos we know as human consciousness may just be too complex and analog to replicate in digital silicon. The biologist Dennis Bray was one of the few voices of dissent at last summer's Singularity Summit. "Although biological components act in ways that are comparable to those in electronic circuits," he argued, in a talk titled "What Cells Can Do That Robots Can't," "they are set apart by the huge number of different states they can adopt. Multiple biochemical processes create chemical modifications of protein molecules, further diversified by association with distinct structures at defined locations of a cell. The resulting combinatorial explosion of states endows living systems with an almost infinite capacity to store information regarding past and present conditions and a unique capacity to prepare for future events." That makes the ones and zeros that computers trade in look pretty crude.

Underlying the practical challenges are a host of philosophical ones. Suppose we did create a computer that talked and acted in a way that was indistinguishable from a human being — in other words, a computer that could pass the Turing test. (Very loosely speaking, such a computer would be able to pass as human in a blind test.) Would that mean that the computer was sentient, the way a human being is? Or would it just be an extremely sophisticated but essentially mechanical automaton without the mysterious spark of consciousness — a machine with no ghost in it? And how would we know?

Even if you grant that the Singularity is plausible, you're still staring at a thicket of unanswerable questions. If I can scan my consciousness into a computer, am I still me? What are the geopolitics and the socioeconomics of the Singularity? Who decides who gets to be immortal? Who draws the line between sentient and nonsentient? And as we approach immortality, omniscience and omnipotence, will our lives still have meaning? By beating death, will we have lost our essential humanity?

Kurzweil admits that there's a fundamental level of risk associated with the Singularity that's impossible to refine away, simply because we don't know what a highly advanced artificial intelligence, finding itself a newly created inhabitant of the planet Earth, would choose to do. It might not feel like competing with us for resources. One of the goals of the Singularity Institute is to make sure not just that artificial intelligence develops but also that the AI is friendly. You don't have to be a super-intelligent cyborg to understand that introducing a superior life-form into your own biosphere is a basic Darwinian error.

If the Singularity is coming, these questions are going to get answers whether we like it or not, and Kurzweil thinks that trying to put off the Singularity by banning technologies is not only impossible but also unethical and probably dangerous. "It would require a totalitarian system to implement such a ban," he says. "It wouldn't work. It would just drive these technologies underground, where the responsible scientists who we're counting on to create the defenses would not have easy access to the tools."

Kurzweil is an almost inhumanly patient and thorough debater. He relishes it. He's

tireless in hunting down his critics so that he can respond to them, point by point, carefully and in detail.

Take the question of whether computers can replicate the biochemical complexity of an organic brain. Kurzweil yields no ground there whatsoever. He does not see any fundamental difference between flesh and silicon that would prevent the latter from thinking. He defies biologists to come up with a neurological mechanism that could not be modeled or at least matched in power and flexibility by software running on a computer. He refuses to fall on his knees before the mystery of the human brain. "Generally speaking," he says, "the core of a disagreement I'll have with a critic is, they'll say, Oh, Kurzweil is underestimating the complexity of reverse-engineering of the human brain or the complexity of biology. But I don't believe I'm underestimating the challenge. I think they're underestimating the power of exponential growth."

This position doesn't make Kurzweil an outlier, at least among Singularitarians. Plenty of people make more-extreme predictions. Since 2005 the neuroscientist Henry Markram has been running an ambitious initiative at the Brain Mind Institute of the Ecole Polytechnique in Lausanne, Switzerland. It's called the Blue Brain project, and it's an attempt to create a neuron-by-neuron simulation of a mammalian brain, using IBM's Blue Gene super-computer. So far, Markram's team has managed to simulate one neocortical column from a rat's brain, which contains about 10,000 neurons. Markram has said that he hopes to have a complete virtual human brain up and running in 10 years. (Even Kurzweil sniffs at this. If it worked, he points out, you'd then have to educate the brain, and who knows how long that would take?)

By definition, the future beyond the Singularity is not knowable by our linear, chemical, animal brains, but Kurzweil is teeming with theories about it. He positively flogs himself to think bigger and bigger; you can see him kicking against the confines of his aging organic hardware. "When people look at the implications of ongoing exponential growth, it gets harder and harder to accept," he says. "So you get people who really accept, yes, things are progressing exponentially, but they fall off the horse at some point because the implications are too fantastic. I've tried to push myself to really look."

In Kurzweil's future, biotechnology and nanotechnology give us the power to manipulate our bodies and the world around us at will, at the molecular level. Progress hyperaccelerates, and every hour brings a century's worth of scientific breakthroughs. We ditch Darwin and take charge of our own evolution. The human genome becomes just so much code to be bug-tested and optimized and, if necessary, rewritten. Indefinite life extension becomes a reality; people die only if they choose to. Death loses its sting once and for all. Kurzweil hopes to bring his dead father back to life.

We can scan our consciousnesses into computers and enter a virtual existence or swap our bodies for immortal robots and light out for the edges of space as intergalactic godlings. Within a matter of centuries, human intelligence will have re-engineered and saturated all the matter in the universe. This is, Kurzweil believes, our destiny as a species.

Or it isn't. When the big questions get answered, a lot of the action will happen where no one can see it, deep inside the black silicon brains of the computers, which will either bloom bit by bit into conscious minds or just continue in ever more brilliant and powerful iterations of nonsentience.

But as for the minor questions, they're already being decided all around us and in plain sight. The more you read about the Singularity, the more you start to see it peeking out at you, coyly, from unexpected directions. Five years ago we didn't have 600 million humans carrying out their social lives over a single electronic network. Now we have Facebook. Five years ago you didn't see people double-checking what they were saying and where they were going, even as they were saying it and going there, using handheld network-enabled digital prosthetics. Now we have iPhones. Is it an unimaginable step to take the iPhones out of our hands and put them into our skulls?

Already 30,000 patients with Parkinson's disease have neural implants. Google is experimenting with computers that can drive cars. There are more than 2,000 robots fighting in Afghanistan alongside the human troops. This month a game show will once again figure in the history of artificial intelligence, but this time the computer will be the guest: an IBM super-computer nicknamed Watson will compete on Jeopardy! Watson

runs on 90 servers and takes up an entire room, and in a practice match in January it finished ahead of two former champions, Ken Jennings and Brad Rutter. It got every question it answered right, but much more important, it didn't need help understanding the questions (or, strictly speaking, the answers), which were phrased in plain English. Watson isn't strong AI, but if strong AI happens, it will arrive gradually, bit by bit, and this will have been one of the bits.

A hundred years from now, Kurzweil and de Grey and the others could be the 22nd century's answer to the Founding Fathers — except unlike the Founding Fathers, they'll still be alive to get credit — or their ideas could look as hilariously retro and dated as Disney's Tomorrowland. Nothing gets old as fast as the future.

But even if they're dead wrong about the future, they're right about the present. They're taking the long view and looking at the big picture. You may reject every specific article of the Singularitarian charter, but you should admire Kurzweil for taking the future seriously. Singularitarianism is grounded in the idea that change is real and that humanity is in charge of its own fate and that history might not be as simple as one damn thing after another. Kurzweil likes to point out that your average cell phone is about a millionth the size of, a millionth the price of and a thousand times more powerful than the computer he had at MIT 40 years ago. Flip that forward 40 years and what does the world look like? If you really want to figure that out, you have to think very, very far outside the box. Or maybe you have to think further inside it than anyone ever has before.