

Published on *h+ Magazine* (<http://hplusmagazine.com>)

[Home](#) > Ray Kurzweil: The h+ Interview

Ray Kurzweil: The h+ Interview

By *Editor*

Created 12/30/2009 - 11:45



Photo credit: Guru Khalsa

A 3-way conversation with the brilliant and controversial inventor and futurist Ray Kurzweil needs little or no introduction to most *h+* readers. Principal developer of the first omni-font optical character recognition, the first print-to-speech reading machine for the blind, the first CCD flat-bed scanner, the first text-to-speech synthesizer, the first music synthesizer capable of recreating the grand piano and other orchestral instruments, and the first commercially marketed large-vocabulary speech recognition, Ray has been described as “the restless genius” by the *Wall Street Journal*, and “the ultimate thinking machine” by *Forbes, Inc.* The magazine ranked him #8 among entrepreneurs in the United States and called him the “rightful heir to Thomas Edison.” His Kurzweil Technologies, Inc. is an umbrella company for at least eight separate enterprises.

Ray’s writing career rivals his inventions and entrepreneurship. His seminal book, *The Singularity is Near*, presents the Singularity as an overall exponential (doubling) growth trend in technological development, “a future period during which the pace of technological change will be so rapid, its impact so deep, that human life will be irreversibly transformed.”

With his upcoming films, *Transcendent Man* and *The Singularity is Near: A True Story about the Future*, he is becoming an actor, screenplay writer, and director as well.

Sponsored by the Singularity Institute, the first Singularity Summit was held at Stanford University in 2006 to further understanding and discussion about the Singularity concept and the future of technological progress. Founded by Ray, Tyler Emerson, and Peter Thiel, it is a venue for leading thinkers to explore the idea of the Singularity — whether scientist, enthusiast, or skeptic. Ray also founded Singularity University in 2009 with funding from Google and NASA Ames Research Center. Singularity University offers intensive 10-week, 10-day, or 3-day programs examining sets of technologies and disciplines including future studies and forecasting; biotechnology and bioinformatics; nanotechnology; AI, robotics, and cognitive computing; and finance and entrepreneurship.

See Also

- [h+ Magazine Current Issue](#)
- [Transcendent Man: Film About Kurzweil](#)
- [Ray Kurzweil & the Conquest of Mount Improbable](#)
- [Sex & the Singularity](#)

Ray headlined the recent Singularity Summit 2009 in New York City with talks on “The Ubiquity and Predictability of the Exponential Growth of Information Technology” and “Critics of the Singularity.” He was able to take a little time out after the Summit for two separate interview sessions with *h+* Editor-in-Chief R.U. Sirius and Surfdaddy Orca on a variety of topics including consciousness and quantum computing, purposeful complexity, reverse engineering the brain, AI and AGI, GNR technologies and global warming, utopianism and happiness, his upcoming movies, and science fiction.

CONSCIOUSNESS, QUANTUM COMPUTING & COMPLEXITY

RAY KURZWEIL: One area I commented on was the question of a possible link between quantum computing and the brain. Do we need quantum computing to create human level AI? My conclusion is no, mainly because we don't see any quantum computing in the brain. Roger Penrose's conjecture that there was quantum computing in tubules does not seem to have been verified by any experimental evidence.

Quantum computing is a specialized form of computing where you examine in parallel every possible combination of qubits. So it's very good at certain kinds of problems, the classical one being cracking encryption codes by factoring large numbers. But the types of problems that would be vastly accelerated by quantum computing are not things that the human brain is very good at. When it comes to the kinds of problems I just mentioned, the human brain isn't even as good as classical computing. So in terms of what we can do with our brains there's no indication that it involves quantum computing. Do we need quantum computing for consciousness? The only justification for that conjecture from Roger Penrose and Stuart Hameroff is that consciousness is mysterious and quantum mechanics is mysterious, so there must be a link between the two.

I get very excited about discussions about the true nature of consciousness, because I've been thinking about this issue for literally 50 years, going back to junior high school. And it's a very difficult subject. When some article purports to present the neurological basis of consciousness... I read it. And the articles usually start out, “Well, we think that consciousness is caused by...” You know, fill in the blank. And then it goes on with a big extensive examination of that phenomenon. And at the end of the article, I inevitably find myself thinking... where is the link to consciousness? Where is any justification for believing that this phenomenon should cause consciousness? Why would it cause consciousness?

In his presentation, Hameroff said consciousness comes from gamma coherence, basically a certain synchrony between neurons that create gamma waves that are in a certain frequency, something like around 10 cycles per second. And the evidence is, indeed, that gamma coherence goes away with anesthesia.



Anesthesia is an interesting laboratory for consciousness because it extinguishes consciousness. However, there's a lot of other things that anesthesia also does away with. Most of the activity of the neocortex stops with anesthesia, but there's a little bit going on still in the neocortex. It

Photo credit: Guru Khalsa brings up an interesting issue. How do we even know that we're not conscious under anesthesia? We don't remember anything, but memory is not the same thing as consciousness. Consciousness seems to be an emergent property of what goes on in the neocortex, which is where we do our thinking. And you could build a neocortex. In fact, they are being built in the Blue Brain project, and Numenta also has some neocortex models. In terms of hierarchies and number of units in the human brain, these projects are much smaller. But they certainly do interesting things. There are no tubules in there, there's no quantum computing, and there doesn't seem to be a need for it.

Another theory is the idea of purposeful complexity. If you achieve a certain level of complexity, then that is conscious. I actually like that theory the most. I wrote about that extensively in *The Singularity is Near*. There have been attempts to measure complexity. You have Claude Shannon's information theory, which basically involves the smallest algorithm that can generate a string of information. But that doesn't deal with random information. Random information is not compressible, and would represent a lot of Shannon information, but it's not really purposeful complexity. So you have to factor out randomness. Then you get the concept of arbitrary lists of information. Like, say, the New York telephone book is not random. It's only compressible to a limited extent, but it's not a high level of complexity. It's largely an arbitrary list.

I describe a more meaningful concept of Purposeful Complexity in the book. I propose that there are ways of measuring purposeful complexity. In this theory, humans are more conscious than cats, but cats are conscious, but not quite as much because they're not quite as complex. A worm is conscious, but much less so. The sun is conscious. It actually has a fair amount of structure and complexity, but probably less than a cat, so...

SO: How do you go about proving something like that?

RK: Well, that's the problem. My thesis is that there's really no way to measure consciousness. There's no "Consciousness Detector" that you could imagine where the green light comes on and you can say, "OK, this one's conscious!" Or, "No, this one isn't conscious."

Even among humans, there's no clear consensus as to who's conscious and who is not. We're discovering now that people who are considered minimally conscious, or even in a vegetative state, actually have quite a bit going on in their neocortex and we've been able to communicate with some of them using either real-time brain scanning or other methods.

Today, nobody worries too much about causing pain and suffering to their computer software. But we will get to a point where the emotional reactions of virtual beings will be convincing, unlike the characters in the computer games today. And that will become a real issue. That's the whole thesis of my movie, *The Singularity is Near*. But it really comes down to the fact that it's not a scientific issue, which is to say there's still a role for philosophy.

How do we even know that we're not conscious under anesthesia? We don't remember anything, but memory is not the same thing as consciousness.

Some scientists say, "Well, it's not a scientific issue, therefore it's not a real issue. Therefore consciousness is just an illusion and we should not waste time on it." But we shouldn't be too quick to throw it overboard because our whole moral system and ethical system is based on consciousness. If I cause suffering to some other conscious human, that's considered immoral and probably a crime. On the other hand, if I destroy some property, it's probably okay if it's my property. If it's your property, it's probably not okay. But that's not because I'm causing pain and suffering to the property. I'm causing pain and suffering to the owner of the property.

And there's recognition that animals are probably conscious and that animal cruelty is not okay. But it's okay to cause pain and suffering to the avatar in your computer, at least today. That may not be the case 20 years from now.

ARTIFICIAL INTELLIGENCE & REVERSE ENGINEERING THE BRAIN

RK: I'm working on a book called *How the Mind Works* and *How to Build One*. It's mostly about the brain, but the reason I call it the mind rather than the brain is to bring in these issues of consciousness. A brain becomes a mind as it melds with its own body, and in fact, our sort of circle of identity goes beyond our body. We certainly interact with our environment. It's not a clear distinction between who we are and what we are not.

My concept of the value of reverse engineering the human brain is not that you just simulate a brain in a sort of mechanistic way, without trying to understand what is going on. David Chalmers says he doesn't think this is a fruitful direction. And I would agree that just simulating a brain... *mindlessly*, so to speak... that's not going to get you far enough. The purpose of reverse engineering the human brain is to understand the basic principles of intelligence.

Once you have a simulation working, you can start modifying things. Certain things may not matter. Some things may be very critical. So you learn what's important. You learn the basic principles by which the human brain handles hierarchies and variance, properties of patterns, high-level features and so on. And it appears that the neocortex has this very uniform structure. If we learn those principles, we can then engineer them and amplify them and focus on them. That's engineering.

Now, a big evolutionary innovation with *homo sapiens* is that we have a bigger forehead so that we could fit a larger cortex. But it's still quite limited. And it's running on the substrate that transmits information from one part of the brain to another at a few hundred feet per second, which is a million times slower than electronics. The intra-neural connections compute at about 100 or 200 calculations per second, which is somewhere between 1,000,000 to 10,000,000 times slower than electronics. So if we can get past the substrates, we don't have to settle for a billion of these recognizers. We could have a trillion of them, or a thousand trillion. We could have many more hierarchal levels. We can design it to solve more difficult problems.



Photo credit: Guru Khalsa

So that's really the purpose of reverse engineering the human brain. But there are other benefits. We'll get more insight into ourselves. We'll have better means for fixing the brain. Right now, we have vague psychiatric models as to what's going on in the brain of someone with bipolar disease or schizophrenia. We'll not only understand human function, we'll understand human dysfunction. We'll have better means of fixing those problems. And moreover we'll "fix" the limitation that we all have in thinking in a very small, slow, fairly-messily organized substrate. Of course, we have to be careful. What might seem like just a messy arbitrary complexity that evolution put in may in fact be part of some real principle that we don't understand at first.

I'm not saying this is simple. But on the other hand, I had this debate with John Horgan, who wrote a critical article about my views in IEEE Spectrum. Horgan says that we would need trillions of lines of code to emulate the brain and that's far beyond the complexity that we've shown we can handle in our software. The most sophisticated software programs are only tens of millions of lines of code. But that's complete nonsense. Because, first of all, there's no way the brain could be that complicated. The design of the brain is in the genome. The genome — well... it's 800 million bytes. Well, back up and take a look at that. It's 3 billion base pairs, 6 billion bits, 800 million bytes before compression — but it's replete with redundancies. Lengthy sequences like ALU are repeated hundreds of thousands of times. In *The Singularity is Near*, I show that if you apply lossless compression, you get down to about 50 million bytes. About half of that is the brain, so that's about 25 million bytes. That's about a million lines of code. That's one derivation. You could also look at the amount of complexity that appears to be necessary to perform functional simulations of different brain regions. You actually get about the same answer, about a million lines of code. So with two very different methods, you come up with a similar order of magnitude. There just isn't trillions of lines of code — of complexity — in the design of the brain. There is trillions, or even thousands of trillions of bytes of information, but that's not complexity because there's massive redundancy.

Just simulating a brain... mindlessly, so to speak... that's not going to get you far enough.

For instance, the cerebellum, which comprises half the neurons in the brain and does some of our skill formation, has one module repeated 10 billion times with some random variation with each repetition within certain constraints. And there are only a few genes that describe the wiring of the cerebellum that comprise a few tens of thousands of bytes of design information. As we learn skills like catching a fly ball — then it gets filled up with trillions of bytes of information. But just like we don't need trillions of bytes of design information to design a trillion-byte memory system, there are massive redundancies and repetition and a certain amount of randomness in the implementation of the brain. It's a probabilistic fractal. If you look at the Mandelbrot set, it is an exquisitely complex design.

SO: So you're saying the initial intelligence that passes the Turing test is likely to be a reverse-engineered brain, as opposed to a software architecture that's based on weighted probabilistic analysis, genetic algorithms, and so forth?

RK: I would put it differently. We have a toolkit of AI techniques now. I actually don't draw that sharp a distinction between narrow AI techniques and AGI techniques. I mean, you can list them — Markov models, different forms of neural nets and genetic algorithms, logic systems, search algorithms, learning algorithms. These are techniques. Now, they go by the label AGI. We're going to add to that tool kit as we learn how the human brain does it. And then, with more and more powerful hardware, we'll be able to put together very powerful systems.

My vision is that all the different avenues — studying individual neurons, studying brain wiring, studying brain performance, simulating the brain either by doing neuron-by-neuron simulations or functional simulations — and then, all the AI work that has nothing to do with direct emulation of the brain — it's all helping. And we get from here to there through thousands of little steps like that, not through one grand leap.

GLOBAL WARMING & GNR TECHNOLOGIES

SO: James Lovelock, the ecologist behind the Gaia hypothesis, came out a couple of years ago with a prediction that more than 6 billion people are going to perish by the end of this century, mostly because of climate change. Do you see the GNR technologies coming on line to mitigate that kind of a catastrophe?

RK: Absolutely. Those projections are based on linear thinking, as if nothing's going to happen over the next 50 or 100 years. It's ridiculous. For example, we're applying nanotechnology to solar panels. The cost per watt of solar energy is coming down dramatically. As a result, the amount of solar energy is growing exponentially. It's doubling every two years, reliably, for the last 20 years. People ask, "Is there really enough solar energy to meet all of our energy needs?" It's actually 10,000 times more than we need. And



Photo credit: transcendentman.com

yes you lose some with cloud cover and so forth, but we only have to capture one part in 10,000. If you put efficient solar collection panels on a small percentage of the deserts in the world, you would meet 100% of our energy needs. And there's also the same kind of progress being made on energy storage to deal with the intermittency of solar. There are only eight doublings to go before solar meets 100% of our energy needs. We're awash in sunlight and

these new technologies will enable us to capture that in a clean and renewable fashion. And then, geothermal — you have the potential for incredible amounts of energy.

Global warming — regardless of what you think of the models and whether or not it's been human-caused — it's only been one degree Fahrenheit in the last 100 years. There just isn't a dramatic global warming so far. I think there are lots of reasons we want to move away from fossil fuels, but I would not put greenhouse gasses at the top of the list.

These new energy technologies are decentralized. They're relatively safe. Solar energy, unlike say nuclear power plants and other centralized facilities, are safe from disaster and sabotage and are non-polluting. So I believe that's the future of energy, and of resource utilization in general.

THE SINGULARITY, UTOPIA & HAPPINESS

R.U. SIRIUS: Have any critics of your ideas offered a social critique that gives you pause?

RK: I still think Bill Joy articulated the concerns best in his *Wired* cover story of some years ago. My vision is not a utopian one. For example, I'm working with the U.S. Army on developing a rapid response system for biological viruses, and that's actually the approach that I advocate — that we need to put resources and attention to the downsides. But I think we do have the scientific tools to create a rapid response system in case of a biological viral attack. It took us five years to sequence HIV; we can sequence a virus now in one day. And we could, in a matter of days, create an RNA interference medication based on sequencing a new biological virus. This is something we created to contend with software viruses. And we have a technological immune system that works quite well.

And we also need ethical standards for responsible practitioners of AI, similar to the Asilomar Guidelines for biotech, or the Forsyth Institute Guidelines for nanotech, which are based on the Asilomar Guidelines. So it's a complicated issue. We can't just come up with a simple solution and then just cross it off our worry list. On the other hand, these technologies can vastly expand our creativity. They've already democratized the tools of creativity. And they are overcoming human suffering, extending our longevity and can provide not only radical life extension but radical life expansion.

There's a lot of talk about existential risks. I worry that painful episodes are even more likely. You know, 60 million people were killed in WWII. That was certainly exacerbated by the powerful destructive tools that we had then. I'm fairly optimistic that we will make it through. I'm less optimistic that we can avoid painful episodes. I do think decentralized communication actually helps reduce violence in the world. It may not seem that way because you just turn on CNN and you've got lots of violence right in your living room. But that kind of visibility actually helps us to solve problems.



Photo credit: Guru Khalsa

RUS: You've probably heard the phrase from critics of the Singularity — they call it the “Rapture of the Nerds.” And a lot of people who are into this idea do seem to envision the Singularity as a sort of magical place where pretty much anything can happen and all your dreams come true. How do you separate your view of the Singularity from a utopian view?

RK: I don't necessarily think they are utopian. I mean, the whole thing is difficult to imagine. We have a certain level of intelligence and it's difficult to imagine what it would mean and what would happen when we vastly expand that. It would be like asking cavemen and women, “Well, gee, what would you like to have?” And they'd say, “Well, we'd like a bigger rock to keep the animals out of our cave and we'd like to keep this fire from burning out?” And you'd say, “Well, don't you want a good web site? What about your Second Life habitat?” They couldn't imagine these things. And those are just technological innovations.

So the future does seem magical. But that gets back to that Arthur C. Clark quote that any sufficiently developed technology is indistinguishable from magic. That's the nature of technology — it transcends limitations that exist without that technology. Television and radio seem magical — you have these waves going through the air, and they're invisible, and they go at the speed of light and they carry pictures and sounds. So think of a substrate that's a million times faster. We'll be overcoming problems at a very rapid rate, and that will seem magical. But that doesn't mean it's not rooted in science and technology.

I say it's not utopian because it also introduces new problems. Artificial Intelligence is the most difficult to contend with, because whereas we can articulate technical solutions to the dangers of biotech, there's no purely technical solution to a so-called unfriendly AI. We can't just say, “We'll just put this little software code sub-routine in our AIs, and that'll keep them safe.” I mean, it really comes down to what the goals and intentions of that artificial intelligence are. We face daunting challenges.

RUS: I think when most people think of utopia, they probably just think about everybody being happy and feeling good.

RK: I really don't think that's the goal. I think the goal has been demonstrated by the multi-billion-year history of biological evolution and the multi-thousand-year history of technological evolution. The goal is to be creative and create entities of beauty, of insight, that solve problems. I mean, for myself as an inventor, that's what makes me happy. But it's not a state that you would seek to be in at all times, because it's fleeting. It's momentary.

To sit around being happy all the time is not the goal. In fact, that's kind of a downside. Because if we were to just stimulate our pleasure centers and sit around in a morphine high at all times — that's been recognized as a downside and it ultimately leads to a profound unhappiness. We can identify things that make us unhappy. If we have diseases that rob our faculties or cause physical or emotional pain — that makes us unhappy and prevents us from having these moments of connection with another person, or a connection with an idea, then we should solve that. But happiness is not the right goal. I think it represents the cutting edge of the evolutionary condition to seek greater horizons and to always want to transcend whatever our limitations are at the time. And so it's not our nature just to sit back and be happy.

MOVIES & SCIENCE FICTION

RUS: You've got two films coming out, *Transcendent Man* and *The Singularity is Near*. What do you think the impact will be of having those two films out in the world?

I think there are anti-technology movements that continue to spread among the intelligentsia that are actually pretty ignorant.

RK: Well, *Transcendent Man* has already premiered at the Tribeca film festival and it will have an international premier at the Amsterdam documentary film festival next month. There's quite a lot of interest in that, and there are discussions with distributors. So it's expected to have a theatrical release both in this country and internationally early next year. And *The Singularity is Near* will follow.

Movies are a really different venue. They cover less content than a book but they have more emotional impact. It's a big world out there. No matter how many times I speak — and even with all the press coverage of all these ideas, whether it's featuring me or others — I'm impressed by how many otherwise thoughtful people still haven't heard of these ideas.

I think it's important that people not just understand the Singularity, which is some decades away, but the impact right now, and in the fairly near future, of the exponential growth of information technology. It's not an obscure part of the economy and the social scene. Every new period is going to bring new opportunities and new challenges. These are the issues that people should be focusing on. It's not just the engineers who should be worrying about the downsides of biotechnology or nanotechnology, for example. And people should also understand the opportunities. And I think there are anti-technology movements that continue to spread among the intelligentsia that are actually pretty ignorant.

RUS: Do you read science fiction novels and watch science fiction television, or science fiction movies?

RK: I have seen most of the popular science fiction movies.

RUS: Any that you find particularly interesting or enjoyable?

RK: Well, one problem with a lot of science fiction — and this is particularly true of movies — is they take one change, like the human-level cyborgs in the movie *AI*, and they put it in a world that is otherwise unchanged. So in *AI*, the coffee maker is the same and the cars are the same. There's no virtual reality, but you had human-level cyborgs. Part of the reason for that is the limitation of the form. To try to present a world in which everything is quite different would take the whole movie, and people wouldn't be able to follow it very easily. It's certainly a challenge to do that. I am in touch with some movie makers who want to try to do that.

I thought *The Matrix* was pretty good in its presentation of virtual reality. And they also had sort of AI-based people in that movie, so it did present a number of ideas. Some of the concepts were arbitrary as to how things work in the matrix, but it was pretty interesting.

See Also

- [h+ Magazine Current Issue](#)
- [Transcendent Man: Film About Kurzweil](#)
- [Ray Kurzweil & the Conquest of Mount Improbable](#)
- [Sex & the Singularity](#)

Resources:

Kurzweil AI

<http://www.kurzweilai.net/index.html?flash=1>

Transcendental Man

<http://www.transcendentman.com/>

The Singularity is Near (Movie)

<http://singularity.com/themovie/>

[AI](#)

- [about us](#)
- [contact us](#)
- [advertising](#)
- [editor's blog](#)
- [articles](#)
- [multimedia](#)
- [directory](#)
- [community](#)
- [archives](#)
- [privacy policy](#)
- [terms & conditions](#)



Source URL: <http://hplusmagazine.com/articles/ai/ray-kurzweil-h-interview>