

Questions and Answers on the Singularity

The Singularity is Near, When Humans Transcend Biology

By Ray Kurzweil, Viking Press

So what is the Singularity?

Within a quarter century, nonbiological intelligence will match the range and subtlety of human intelligence. It will then soar past it because of the continuing acceleration of information-based technologies, as well as the ability of machines to instantly share their knowledge. Intelligent nanorobots will be deeply integrated in our bodies, our brains, and our environment, overcoming pollution and poverty, providing vastly extended longevity, full-immersion virtual reality incorporating all of the senses (like “The Matrix”), “experience beaming” (like “Being John Malkovich”), and vastly enhanced human intelligence. The result will be an intimate merger between the technology-creating species and the technological evolutionary process it spawned.

And that’s the Singularity?

No, that’s just the precursor. Nonbiological intelligence will have access to its own design and will be able to improve itself in an increasingly rapid redesign cycle. We’ll get to a point where technical progress will be so fast that unenhanced human intelligence will be unable to follow it. That will mark the Singularity.

When will that occur?

I set the date for the Singularity—representing a profound and disruptive transformation in human capability—as 2045. The nonbiological intelligence created in that year will be one billion times more powerful than all human intelligence today.

Why is this called the Singularity?

The term “Singularity” in my book is comparable to the use of this term by the physics community. Just as we find it hard to see beyond the event horizon of a black hole, we also find it difficult to see beyond the event horizon of the historical Singularity. How can we, with our limited biological brains, imagine what our future civilization, with its intelligence multiplied trillions-fold, be capable of thinking and doing? Nevertheless, just as we can draw conclusions about the nature of black holes through our conceptual thinking, despite never having actually been inside one, our thinking today is powerful enough to have meaningful insights into the implications of the Singularity. That’s what I’ve tried to do in this book.

Okay, let’s break this down. It seems a key part of your thesis is that we will be able to capture the intelligence of our brains in a machine.

Indeed.

So how are we going to achieve that?

We can break this down further into hardware and software requirements. In the book, I show how we need about 10 quadrillion (10^{16}) calculations per second (cps) to provide a functional equivalent to all the regions of the brain. Some estimates are lower than this by a factor of 100. Supercomputers are already at 100 trillion (10^{14}) cps, and will hit 10^{16} cps around the end of this decade. Several supercomputers with 1 quadrillion cps are already on the drawing board, with two Japanese efforts targeting 10 quadrillion cps around the end of the decade. By 2020, 10 quadrillion cps will be available for around \$1,000. Achieving the hardware requirement was controversial when my last book on this topic, *The Age of Spiritual Machines*, came out in 1999, but is now pretty much of a mainstream view among informed observers. Now the controversy is focused on the algorithms.

And how will we recreate the algorithms of human intelligence?

To understand the principles of human intelligence we need to reverse-engineer the human brain. Here, progress is far greater than most people realize. The spatial and temporal (time) resolution of brain scanning is also progressing at an exponential rate, roughly doubling each year, like most everything else having to do with information. Just recently, scanning tools can see individual interneuronal connections, and watch them fire in real time. Already, we have mathematical models and simulations of a couple dozen regions of the brain, including the cerebellum, which comprises more than half the neurons in the brain. IBM is now creating a simulation of about 10,000 cortical neurons, including tens of millions of connections. The first version will simulate the electrical activity, and a future version will also simulate the relevant chemical activity. By the mid 2020s, it's conservative to conclude that we will have effective models for all of the brain.

So at that point we'll just copy a human brain into a supercomputer?

I would rather put it this way: At that point, we'll have a full understanding of the methods of the human brain. One benefit will be a deep understanding of ourselves, but the key implication is that it will expand the toolkit of techniques we can apply to create artificial intelligence. We will then be able to create nonbiological systems that match human intelligence in the ways that humans are now superior, for example, our pattern-recognition abilities. These superintelligent computers will be able to do things we are not able to do, such as share knowledge and skills at electronic speeds.

By 2030, a thousand dollars of computation will be about a thousand times more powerful than a human brain. Keep in mind also that computers will not be organized as discrete objects as they are today. There will be a web of computing deeply integrated into the environment, our bodies and brains.

You mentioned the AI tool kit. Hasn't AI failed to live up to its expectations?

There was a boom and bust cycle in AI during the 1980s, similar to what we saw recently in e-commerce and telecommunications. Such boom-bust cycles are often harbingers of true revolutions; recall the railroad boom and bust in the 19th century. But just as the Internet "bust" was not the end of the Internet, the so-called "AI Winter" was not the end of the story for AI either. There are hundreds of applications of "narrow AI" (machine intelligence that equals or exceeds human intelligence for specific tasks) now permeating our modern infrastructure. Every time you send an email or make a cell phone call,

intelligent algorithms route the information. AI programs diagnose electrocardiograms with an accuracy rivaling doctors, evaluate medical images, fly and land airplanes, guide intelligent autonomous weapons, make automated investment decisions for over a trillion dollars of funds, and guide industrial processes. These were all research projects a couple of decades ago. If all the intelligent software in the world were to suddenly stop functioning, modern civilization would grind to a halt. Of course, our AI programs are not intelligent enough to organize such a conspiracy, at least not yet.

Why don't more people see these profound changes ahead?

Hopefully after they read my new book, they will. But the primary failure is the inability of many observers to think in exponential terms. Most long-range forecasts of what is technically feasible in future time periods dramatically underestimate the power of future developments because they are based on what I call the “intuitive linear” view of history rather than the “historical exponential” view. My models show that we are doubling the paradigm-shift rate every decade. Thus the 20th century was gradually speeding up to the rate of progress at the end of the century; its achievements, therefore, were equivalent to about twenty years of progress at the rate in 2000. We'll make another twenty years of progress in just fourteen years (by 2014), and then do the same again in only seven years. To express this another way, we won't experience one hundred years of technological advance in the 21st century; we will witness on the order of 20,000 years of progress (again, when measured by the rate of progress in 2000), or about 1,000 times greater than what was achieved in the 20th century.

The exponential growth of information technologies is even greater: we're doubling the power of information technologies, as measured by price-performance, bandwidth, capacity and many other types of measures, about every year. That's a factor of a thousand in ten years, a million in twenty years, and a billion in thirty years. This goes far beyond Moore's law (the shrinking of transistors on an integrated circuit, allowing us to double the price-performance of electronics each year). Electronics is just one example of many. As another example, it took us 14 years to sequence HIV; we recently sequenced SARS in only 31 days.

So this acceleration of information technologies applies to biology as well?

Absolutely. It's not just computer devices like cell phones and digital cameras that are accelerating in capability. Ultimately, everything of importance will be comprised essentially of information technology. With the advent of nanotechnology-based manufacturing in the 2020s, we'll be able to use inexpensive table-top devices to manufacture on-demand just about anything from very inexpensive raw materials using information processes that will rearrange matter and energy at the molecular level.

We'll meet our energy needs using nanotechnology-based solar panels that will capture the energy in .03 percent of the sunlight that falls on the Earth, which is all we need to meet our projected energy needs in 2030. We'll store the energy in highly distributed fuel cells.

I want to come back to both biology and nanotechnology, but how can you be so sure of these developments? Isn't technical progress on specific projects essentially unpredictable?

Predicting *specific* projects is indeed not feasible. But the result of the overall complex, chaotic evolutionary process of technological progress *is* predictable.

People intuitively assume that the current rate of progress will continue for future periods. Even for those who have been around long enough to experience how the pace of change increases over time, unexamined intuition leaves one with the impression that change occurs at the same rate that we have experienced most recently. From the mathematician's perspective, the reason for this is that an exponential curve looks like a straight line when examined for only a brief duration. As a result, even sophisticated commentators, when considering the future, typically use the current pace of change to determine their expectations in extrapolating progress over the next ten years or one hundred years. This is why I describe this way of looking at the future as the "intuitive linear" view. But a serious assessment of the history of technology reveals that technological change is exponential. Exponential growth is a feature of any evolutionary process, of which technology is a primary example.

As I show in the book, this has also been true of biological evolution. Indeed, technological evolution emerges from biological evolution. You can examine the data in different ways, on different timescales, and for a wide variety of technologies, ranging from electronic to biological, as well as for their implications, ranging from the amount of human knowledge to the size of the economy, and you get the same exponential—not linear—progression. I have over forty graphs in the book from a broad variety of fields that show the exponential nature of progress in information-based measures. For the price-performance of computing, this goes back over a century, well before Gordon Moore was even born.

Aren't there are a lot of predictions of the future from the past that look a little ridiculous now?

Yes, any number of bad predictions from other futurists in earlier eras can be cited to support the notion that we cannot make reliable predictions. In general, these prognosticators were not using a methodology based on a sound theory of technology evolution. I say this not just looking backwards now. I've been making accurate forward-looking predictions for over twenty years based on these models.

But how can it be the case that we can reliably predict the overall progression of these technologies if we cannot even predict the outcome of a single project?

Predicting which company or product will succeed is indeed very difficult, if not impossible. The same difficulty occurs in predicting which technical design or standard will prevail. For example, how will the wireless-communication protocols Wimax, CDMA, and 3G fare over the next several years? However, as I argue extensively in the book, we find remarkably precise and predictable exponential trends when assessing the overall effectiveness (as measured in a variety of ways) of information technologies. And as I mentioned above, information technology will ultimately underlie everything of value.

But how can that be?

We see examples in other areas of science of very smooth and reliable outcomes resulting from the interaction of a great many unpredictable events. Consider that predicting the

path of a single molecule in a gas is essentially impossible, but predicting the properties of the *entire* gas—comprised of a great many chaotically interacting molecules—can be done very reliably through the laws of thermodynamics. Analogously, it is not possible to reliably predict the results of a specific project or company, but the overall capabilities of information technology, comprised of many chaotic activities, can nonetheless be dependably anticipated through what I call "the law of accelerating returns."

What will the impact of these developments be?

Radical life extension, for one.

Sounds interesting, how does that work?

In the book, I talk about three great overlapping revolutions that go by the letters "GNR," which stands for genetics, nanotechnology, and robotics. Each will provide a dramatic increase to human longevity, among other profound impacts. We're in the early stages of the genetics—also called biotechnology—revolution right now. Biotechnology is providing the means to actually change your genes: not just designer babies but designer baby boomers. We'll also be able to rejuvenate all of your body's tissues and organs by transforming your skin cells into youthful versions of every other cell type. Already, new drug development is precisely targeting key steps in the process of atherosclerosis (the cause of heart disease), cancerous tumor formation, and the metabolic processes underlying each major disease and aging process. The biotechnology revolution is already in its early stages and will reach its peak in the second decade of this century, at which point we'll be able to overcome most major diseases and dramatically slow down the aging process.

That will bring us to the nanotechnology revolution, which will achieve maturity in the 2020s. With nanotechnology, we will be able to go beyond the limits of biology, and replace your current "human body version 1.0" with a dramatically upgraded version 2.0, providing radical life extension.

And how does that work?

The "killer app" of nanotechnology is "nanobots," which are blood-cell sized robots that can travel in the bloodstream destroying pathogens, removing debris, correcting DNA errors, and reversing aging processes.

Human body version 2.0?

We're already in the early stages of augmenting and replacing each of our organs, even portions of our brains with neural implants, the most recent versions of which allow patients to download new software to their neural implants from outside their bodies. In the book, I describe how each of our organs will ultimately be replaced. For example, nanobots could deliver to our bloodstream an optimal set of all the nutrients, hormones, and other substances we need, as well as remove toxins and waste products. The gastrointestinal tract could be reserved for culinary pleasures rather than the tedious biological function of providing nutrients. After all, we've already in some ways separated the communication and pleasurable aspects of sex from its biological function.

And the third revolution?

The robotics revolution, which really refers to “strong” AI, that is, artificial intelligence at the human level, which we talked about earlier. We’ll have both the hardware and software to recreate human intelligence by the end of the 2020s. We’ll be able to improve these methods and harness the speed, memory capabilities, and knowledge-sharing ability of machines.

We’ll ultimately be able to scan all the salient details of our brains from inside, using billions of nanobots in the capillaries. We can then back up the information. Using nanotechnology-based manufacturing, we could recreate your brain, or better yet reconstitute it in a more capable computing substrate.

Which means?

Our biological brains use chemical signaling, which transmit information at only a few hundred feet per second. Electronics is already millions of times faster than this. In the book, I show how one cubic inch of nanotube circuitry would be about one hundred million times more powerful than the human brain. So we’ll have more powerful means of instantiating our intelligence than the extremely slow speeds of our interneuronal connections.

So we’ll just replace our biological brains with circuitry?

I see this starting with nanobots in our bodies and brains. The nanobots will keep us healthy, provide full-immersion virtual reality from within the nervous system, provide direct brain-to-brain communication over the Internet, and otherwise greatly expand human intelligence. But keep in mind that nonbiological intelligence is doubling in capability each year, whereas our biological intelligence is essentially fixed in capacity. As we get to the 2030s, the nonbiological portion of our intelligence will predominate.

The closest life extension technology, however, is biotechnology, isn’t that right?

There’s certainly overlap in the G, N and R revolutions, but that’s essentially correct.

So tell me more about how genetics or biotechnology works.

As we are learning about the information processes underlying biology, we are devising ways of mastering them to overcome disease and aging and extend human potential. One powerful approach is to start with biology’s information backbone: the genome. With gene technologies, we’re now on the verge of being able to control how genes express themselves. We now have a powerful new tool called RNA interference (RNAi), which is capable of turning specific genes off. It blocks the messenger RNA of specific genes, preventing them from creating proteins. Since viral diseases, cancer, and many other diseases use gene expression at some crucial point in their life cycle, this promises to be a breakthrough technology. One gene we’d like to turn off is the fat insulin receptor gene, which tells the fat cells to hold on to every calorie. When that gene was blocked in mice, those mice ate a lot but remained thin and healthy, and actually lived 20 percent longer.

New means of adding new genes, called gene therapy, are also emerging that have overcome earlier problems with achieving precise placement of the new genetic information. One company I’m involved with, United Therapeutics, cured pulmonary

hypertension in animals using a new form of gene therapy and it has now been approved for human trials.

So we're going to essentially reprogram our DNA.

That's a good way to put it, but that's only one broad approach. Another important line of attack is to regrow our own cells, tissues, and even whole organs, and introduce them into our bodies without surgery. One major benefit of this "therapeutic cloning" technique is that we will be able to create these new tissues and organs from versions of our cells that have also been made younger—the emerging field of rejuvenation medicine. For example, we will be able to create new heart cells from your skin cells and introduce them into your system through the bloodstream. Over time, your heart cells get replaced with these new cells, and the result is a rejuvenated "young" heart with your own DNA.

Drug discovery was once a matter of finding substances that produced some beneficial effect without excessive side effects. This process was similar to early humans' tool discovery, which was limited to simply finding rocks and natural implements that could be used for helpful purposes. Today, we are learning the precise biochemical pathways that underlie both disease and aging processes, and are able to design drugs to carry out precise missions at the molecular level. The scope and scale of these efforts is vast.

But perfecting our biology will only get us so far. The reality is that biology will never be able to match what we will be capable of engineering, now that we are gaining a deep understanding of biology's principles of operation.

Isn't nature optimal?

Not at all. Our interneuronal connections compute at about 200 transactions per second, at least a million times slower than electronics. As another example, a nanotechnology theorist, Rob Freitas, has a conceptual design for nanobots that replace our red blood cells. A conservative analysis shows that if you replaced 10 percent of your red blood cells with Freitas' "respirocytes," you could sit at the bottom of a pool for four hours without taking a breath.

If people stop dying, isn't that going to lead to overpopulation?

A common mistake that people make when considering the future is to envision a major change to today's world, such as radical life extension, as if nothing else were going to change. The GNR revolutions will result in other transformations that address this issue. For example, nanotechnology will enable us to create virtually any physical product from information and very inexpensive raw materials, leading to radical wealth creation. We'll have the means to meet the material needs of any conceivable size population of biological humans. Nanotechnology will also provide the means of cleaning up environmental damage from earlier stages of industrialization.

So we'll overcome disease, pollution, and poverty—sounds like a utopian vision.

It's true that the dramatic scale of the technologies of the next couple of decades will enable human civilization to overcome problems that we have struggled with for eons. But these developments are not without their dangers. Technology is a double edged

sword—we don't have to look past the 20th century to see the intertwined promise and peril of technology.

What sort of perils?

G, N, and R each have their downsides. The existential threat from genetic technologies is already here: the same technology that will soon make major strides against cancer, heart disease, and other diseases could also be employed by a bioterrorist to create a bioengineered biological virus that combines ease of transmission, deadliness, and stealthiness, that is, a long incubation period. The tools and knowledge to do this are far more widespread than the tools and knowledge to create an atomic bomb, and the impact could be far worse.

So maybe we shouldn't go down this road.

It's a little late for that. But the idea of relinquishing new technologies such as biotechnology and nanotechnology is already being advocated. I argue in the book that this would be the wrong strategy. Besides depriving human society of the profound benefits of these technologies, such a strategy would actually make the dangers worse by driving development underground, where responsible scientists would not have easy access to the tools needed to defend us.

So how do we protect ourselves?

I discuss strategies for protecting against dangers from abuse or accidental misuse of these very powerful technologies in chapter 8. The overall message is that we need to give a higher priority to preparing protective strategies and systems. We need to put a few more stones on the defense side of the scale. I've given testimony to Congress on a specific proposal for a "Manhattan" style project to create a rapid response system that could protect society from a new virulent biological virus. One strategy would be to use RNAi, which has been shown to be effective against viral diseases. We would set up a system that could quickly sequence a new virus, prepare a RNA interference medication, and rapidly gear up production. We have the knowledge to create such a system, but we have not done so. We need to have something like this in place before its needed.

Ultimately, however, nanotechnology will provide a completely effective defense against biological viruses.

But doesn't nanotechnology have its own self-replicating danger?

Yes, but that potential won't exist for a couple more decades. The existential threat from engineered biological viruses exists right now.

Okay, but how will we defend against self-replicating nanotechnology?

There are already proposals for ethical standards for nanotechnology that are based on the Asilomar conference standards that have worked well thus far in biotechnology. These standards will be effective against unintentional dangers. For example, we do not need to provide self-replication to accomplish nanotechnology manufacturing.

But what about intentional abuse, as in terrorism?

We'll need to create a nanotechnology immune system—good nanobots that can protect us from the bad ones.

Blue goo to protect us from the gray goo!

Yes, well put. And ultimately we'll need the nanobots comprising the immune system to be self-replicating. I've debated this particular point with a number of other theorists, but I show in the book why the nanobot immune system we put in place will need the ability to self-replicate. That's basically the same "lesson" that biological evolution learned.

Ultimately, however, strong AI will provide a completely effective defense against self-replicating nanotechnology.

Okay, what's going to protect us against a pathological AI?

Yes, well, that would have to be a yet more intelligent AI.

This is starting to sound like that story about the universe being on the back of a turtle, and that turtle standing on the back of another turtle, and so on all the way down. So what if this more intelligent AI is unfriendly? Another even smarter AI?

History teaches us that the more intelligent civilization—the one with the most advanced technology—prevails. But I do have an overall strategy for dealing with unfriendly AI, which I discuss in chapter 8.

Okay, so I'll have to read the book for that one. But aren't there limits to exponential growth? You know the story about rabbits in Australia—they didn't keep growing exponentially forever.

There are limits to the exponential growth inherent in each paradigm. Moore's law was not the first paradigm to bring exponential growth to computing, but rather the fifth. In the 1950s they were shrinking vacuum tubes to keep the exponential growth going and then that paradigm hit a wall. But the exponential growth of computing didn't stop. It kept going, with the new paradigm of transistors taking over. Each time we can see the end of the road for a paradigm, it creates research pressure to create the next one. That's happening now with Moore's law, even though we are still about fifteen years away from the end of our ability to shrink transistors on a flat integrated circuit. We're making dramatic progress in creating the sixth paradigm, which is three-dimensional molecular computing.

But isn't there an overall limit to our ability to expand the power of computation?

Yes, I discuss these limits in the book. The ultimate 2 pound computer could provide 10^{42} cps, which will be about 10 quadrillion (10^{16}) times more powerful than all human brains put together today. And that's if we restrict the computer to staying at a cold temperature. If we allow it to get hot, we could improve that by a factor of another 100 million. And, of course, we'll be devoting more than two pounds of matter to computing. Ultimately, we'll use a significant portion of the matter and energy in our vicinity. So, yes, there are limits, but they're not very limiting.

And when we saturate the ability of the matter and energy in our solar system to support intelligent processes, what happens then?

Then we'll expand to the rest of the Universe.

Which will take a long time I presume.

Well, that depends on whether we can use wormholes to get to other places in the Universe quickly, or otherwise circumvent the speed of light. If wormholes are feasible, and analyses show they are consistent with general relativity, we could saturate the universe with our intelligence within a couple of centuries. I discuss the prospects for this in the chapter 6. But regardless of speculation on wormholes, we'll get to the limits of computing in our solar system within this century. At that point, we'll have expanded the powers of our intelligence by trillions of trillions.

Getting back to life extension, isn't it natural to age, to die?

Other natural things include malaria, Ebola, appendicitis, and tsunamis. Many natural things are worth changing. Aging may be “natural,” but I don't see anything positive in losing my mental agility, sensory acuity, physical limberness, sexual desire, or any other human ability.

In my view, death is a tragedy. It's a tremendous loss of personality, skills, knowledge, relationships. We've rationalized it as a good thing because that's really been the only alternative we've had. But disease, aging, and death are problems we are now in a position to overcome.

Wait, you said that the golden era of biotechnology was still a decade away. We don't have radical life extension today, do we?

In my last book, *Fantastic Voyage, Live Long Enough to Live Forever*, which I coauthored with Terry Grossman, M.D., we describe a detailed and personalized program you can implement now (which we call “bridge one”) that will enable most people to live long enough to get to the mature phase of the biotechnology evolution (“bridge two”). That in turn will get us to “bridge three,” which is nanotechnology and strong AI, which will result in being able to live indefinitely.

Okay, but won't it get boring to live many hundreds of years?

If humans lived many hundreds of years with no other change in the nature of human life, then, yes, that would lead to a deep ennui. But the same nanobots in the bloodstream that will keep us healthy—by destroying pathogens and reversing aging processes—will also vastly augment our intelligence and experiences. As is its nature, the nonbiological portion of our intelligence will expand its powers exponentially, so it will ultimately predominate. The result will be accelerating change—so we will not be bored.

Won't the Singularity create the ultimate “digital divide” due to unequal access to radical life extension and superintelligent computers?

We need to consider an important feature of the law of accelerating returns, which is a 50 percent annual deflation factor for information technologies, a factor which itself will

increase. Technologies start out affordable only by the wealthy, but at this stage, they actually don't work very well. At the next stage, they're merely expensive, and work a bit better. Then they work quite well and are inexpensive. Ultimately, they're almost free. Cell phones are now at the inexpensive stage. There are countries in Asia where most people were pushing a plow fifteen years ago, yet now have thriving information economies and most people have a cell phone. This progression from early adoption of unaffordable technologies that don't work well to late adoption of refined technologies that are very inexpensive is currently a decade-long process. But that too will accelerate. Ten years from now, this will be a five year progression, and twenty years from now it will be only a two- to three-year lag.

This model applies not just to electronic gadgets but to anything having to do with information, and ultimately that will be mean everything of value, including all manufactured products. In biology, we went from a cost of ten dollars to sequence a base pair of DNA in 1990 to about a penny today. AIDS drugs started out costing tens of thousands of dollars per patient per year and didn't work very well, whereas today, effective drugs are about a hundred dollars per patient per year in poor countries. That's still more than we'd like, but the technology is moving in the right direction. So the digital divide and the have-have not divide is diminishing, not exacerbating. Ultimately, everyone will have great wealth at their disposal.

Won't problems such as war, intolerance, environmental degradation prevent us from reaching the Singularity?

We had a lot of war in the 20th century. Fifty million people died in World War II, and there were many other wars. We also had a lot of intolerance, relatively little democracy until late in the century, and a lot of environmental pollution. All of these problems of the 20th century had no effect on the law of accelerating returns. The exponential growth of information technologies proceeded smoothly through war and peace, through depression and prosperity.

The emerging 21st century technologies tend to be decentralized and relatively friendly to the environment. With the maturation of nanotechnology, we will also have the opportunity to clean up the mess left from the crude early technologies of industrialization.

But won't there still be objections from religious and political leaders, not to mention the common man and woman, to such a radical transformation of humanity?

There were objections to the plow also, but that didn't stop people from using it. The same can be said for every new step in technology. Technologies do have to prove themselves. For every technology that is adopted, many are discarded. Each technology has to demonstrate that it meets basic human needs. The cell phone, for example, meets our need to communicate with one another. We are not going to reach the Singularity in some single great leap forward, but rather through a great many small steps, each seemingly benign and modest in scope.

But what about controversies such as the stem cell issue? Government opposition is clearly slowing down progress in that field.

I clearly support stem cell research, but it is not the case that the field of cell therapies has been significantly slowed down. If anything, the controversy has accelerated creative ways of achieving the holy grail of this field, which is transdifferentiation, that is, creating new differentiated cells you need from your own cells—for example, converting skin cells into heart cells or pancreatic Islet cells. Transdifferentiation has already been demonstrated in the lab. Objections such as those expressed against stem- cell research end up being stones in the water: the stream of progress just flows around them.

Where does God fit into the Singularity?

Although the different religious traditions have somewhat different conceptions of God, the common thread is that God represents unlimited—infinite—levels of intelligence, knowledge, creativity, beauty, and love. As systems evolve—through biology and technology—we find that they become more complex, more intelligent and more knowledgeable. They become more intricate and more beautiful, more capable of higher emotions such as love. So they grow exponentially in intelligence, knowledge, creativity, beauty, and love, all of the qualities people ascribe to God without limit. Although evolution does not reach a literally infinite level of these attributes, it does accelerate towards ever greater levels, so we can view evolution as a spiritual process, moving ever closer to this ideal. The Singularity will represent an explosion of these higher values of complexity.

So are you trying to play God?

Actually, I'm trying to play a human. I'm trying to do what humans do well, which is solve problems.

But will we still be human after all these changes?

That depends on how you define human. Some observers define human based on our limitations. I prefer to define us as the species that seeks—and succeeds—in going beyond our limitations.

Many observers point out how science has thrown us off our pedestal, showing us that we're not as central as we thought, that the stars don't circle around the Earth, that we're not descended from the Gods but rather from monkeys, and before that earthworms.

All of that is true, but it turns out that we are central after all. Our ability to create models—virtual realities—in our brains, combined with our modest-looking thumbs, are enabling us to expand our horizons without limit.