# The Singularity and Human Destiny

## by Patrick Tucker, Assistant Editor, THE FUTURIST

In 1949, several decades before the birth of the PC, computer scientist John von Neumann announced, "It would appear we have reached the limits of what is possible to achieve with computer technology." He quickly and prudently added, "One should be careful with such statements; they tend to sound pretty silly in five years."

More than half a century later, it is Neumann's caveat rather than his prediction that is borne out daily. Our computer intelligence is growing at an ever-quickening pace, surpassed only by our dependence on computer technology. For many of us, computers and computer-based devices have become not so much tools as appendages, third arms that are integral to our lives, cumbersome on occasion, and increasingly willful. This begs the obvious question: If computer power advances beyond our control, how will that change us?

In his most recent book, *The Singularity Is Near,* author and inventor **Ray Kurzweil** provides us with a clue. Imagine that with the help of a small device you could create a perfect replica of any object—Gianlorenzo Bernini's David, the hub cap from a '78 Dodge Dart—seemingly from vapor. Imagine that through virtual-reality software you won't be shackled to a particular position in time, and could exist in several locations at once—work, home, a seaside bungalow in Bora Bora—and each setting in which you chose to locate yourself looked, sounded, and felt perfectly real. Imagine that you could live indefinitely in a world in which all poverty, pollution, and scarcity has been vanquished. Imagine that there existed no limitation to what you could do or be, except for those limitations that you imposed yourself.

This scenario may sound like a cross between Thomas Moore's *Utopia* and Christopher Marlow's *Doctor Faustus,* but according to Ray Kurzweil it is our real and fast-approaching future. It will come about as the result of an explosion in our technological abilities. We will incorporate more computer-based processes into our biological functioning until we transcend our crude, earthly bodies entirely and become machine-based, virtually immortal. This coming period of rapid technological progress and its miraculous effects will occur within the next 50 years and is what Kurzweil refers to as the *Singularity.*

The concept is both startling and optimistic, but it immediately provokes certain philosophical concerns. If nanotechnology allows us to create any object, will any object ever again be valuable? What role will responsibility, temperance, and discipline play in a world where any urge can be gratified at almost the same moment it is felt? What will pass for morality when there is no *mortal* consequence to any action?

These questions cannot and should not be answered all at once—either by Ray Kurzweil, his devotees, or his critics. Rather, what is important is that they be asked, repeatedly and earnestly, and by as many people as possible.

To further the debate of these key issues, THE FUTURIST presents Kurzweil's insights and ideas along with invited commentaries from nanotechnology expert **J. Storrs Hall,** acceleration studies scholar **John Smart,** and sociologists **Damien Broderick** and **Richard Eckersely.** Together, they examine this issue of the Singularity to determine how near it is exactly, and explore what it might mean for humanity.

---

Order the March-April 2006 issue of THE FUTURIST from https://www.wfs.org/futuristorder.htm.

©2006 World Future Society • 7910 Woodmont Avenue, Suite 450, Bethesda, MD 20814, U.S.A. • All rights reserved.

# Reinventing Humanity:

## The Future of Machine–Human Intelligence

By Ray Kurzweil

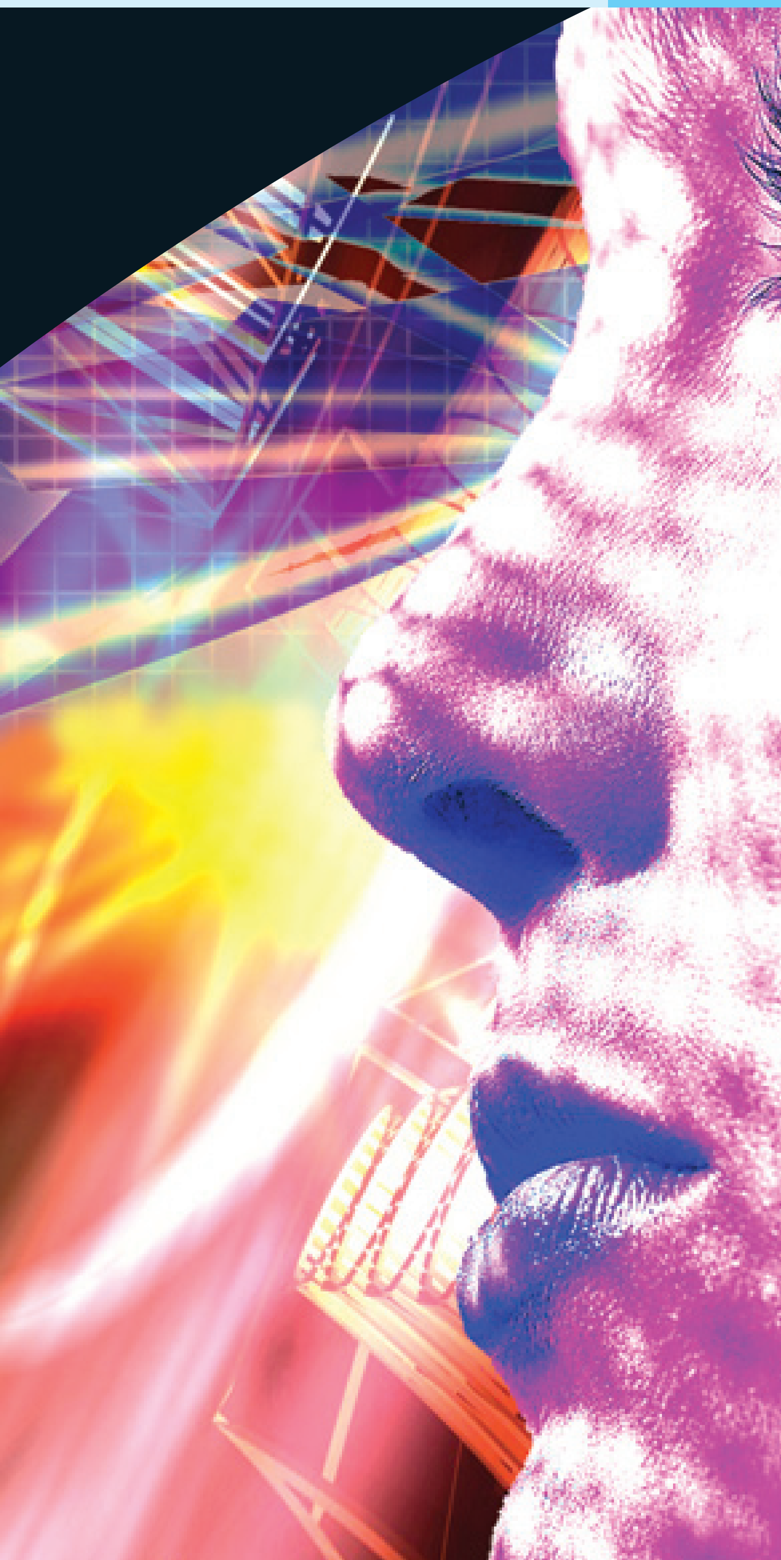

**Author and inventor Ray Kurzweil sees a radical evolution of the human species in the next 40 years.**

We stand on the threshold of the most profound and transformative event in the history of humanity, the "Singularity."

What is the Singularity? From my perspective, the Singularity is a future period during which the pace of technological change will be so fast and far-reaching that human existence on this planet will be irreversibly altered. We will combine our brain power—the knowledge, skills, and personality quirks that make us human—with our computer power in order to think, reason, communicate, and create in ways we can scarcely even contemplate today.

This merger of man and machine, coupled with the sudden explosion in machine intelligence and rapid innovation in gene research and nanotechnology, will result in a world where there is no distinction between the biological and the mechanical, or between physical and virtual reality. These technological revolutions will allow us to transcend our frail bodies with all their limitations. Illness, as we know it, will be eradicated. Through the use of nanotechnology, we will be able to manufacture almost any physical product upon demand, world hunger and poverty will be solved, and pollution will vanish. Human existence will undergo a quantum leap in evolution. We will be able to live as long as we choose. The coming into being of such a world is, in essence, the Singularity.

How is it possible that we could be so close to this enormous change and not see it? The answer is the quickening nature of technological innovation. In thinking about the future, few people take into consideration the fact that human scientific progress is exponential: It expands by repeatedly multiplying by a constant (10 times 10 times 10, and so on) rather than linear (10 plus 10 plus 10, and so on). I emphasize the exponential-versus-linear perspective because it's the most important failure that prognosticators make in considering future trends.

Our forebears expected what lay ahead of them to resemble what they had already experienced, with few exceptions. Because they lived during a time when the rate of technological innovation was so slow as to be unnoticeable, their expectations of an unchanged future were continually fulfilled. Today, we have witnessed the acceleration of the curve. Therefore, we anticipate continuous technological progress and the social repercussions that follow. We see the future as being different from the present. But the future will be far more surprising than most people realize, because few observers have truly internalized the implications of the fact that the rate of change is itself accelerating.

Exponential growth starts out slowly and virtually unnoticeably, but beyond the knee of the curve it turns explosive and profoundly transformative. My models show that we are doubling the paradigm-shift rate for technology innovation every decade. In other words, the twentieth century was gradually speeding up to today's rate of progress; its achievements, therefore, were equivalent to about 20 years of progress at the rate of 2000. We'll make another "20 years" of progress in just 14 years (by 2014), and then do the same again in only seven years. To express this another way, we won't experience 100 years of technological advance in the twenty-first century; we will witness on the order of 20,000 years of progress (again, when measured by today's progress rate), or progress on a level of about 1,000 times greater than what was achieved in the twentieth century.

### How Will We Know the Singularity Is Upon Us?

The first half of the twenty-first century will be characterized by three overlapping revolutions—in genetics, nanotechnology, and robotics. These will usher in the beginning of this period of tremendous change I refer to as the Singularity. We are in the early stages of the genetics revolution today. By understanding the information processes underlying life, we are learning to reprogram our biology to achieve the virtual elimination of disease, dramatic expansion of human potential, and radical life extension. However, Hans Moravec of Carnegie Mellon University's Robotics Institute points out that, no matter how successfully we fine-tune our DNA-based biology, biology will never be able to match what we will be able to engineer once we fully understand life's principles of operation. In other words, we will always be "second-class robots."

The nanotechnology revolution will enable us to redesign and rebuild—molecule by molecule—our bodies and brains and the world with which we interact, going far beyond the limitations of biology.

But the most powerful impending revolution is the robotic revolution. By robotic, I am not referring exclusively—or even primarily—to humanoid-looking droids that take up physical space, but rather to artificial intelligence in all its variations.

Following, I have laid out the principal components underlying each of these coming technological revolutions. While each new wave of progress will solve the problems from earlier transformations, each will also introduce new perils. Each, operating both separately and in concert, underpins the Singularity.

### The Genetic Revolution

Genetic and molecular science will extend biology and correct its obvious flaws (such as our vulnerability to disease). By the year 2020, the full effects of the genetic revolution will be felt across society. We are rapidly gaining the knowledge and the tools to drastically extend the usability of the "house" each of us calls his body and brain.

Nanomedicine researcher Robert Freitas estimates that eliminating 50% of medically preventable conditions would extend human life expectancy to 150 years. If we were able to prevent 99% of naturally occurring medical problems, we'd live to be more than 1,000 years old.

We can see the beginnings of this awesome medical revolution today. The field of genetic biotechnology is fueled by a growing arsenal of tools. Drug discovery was once a matter of finding substrates (chemicals) that produced some beneficial result without excessive side effects, a research method similar to early humans' seeking out rocks and other natural implements that could be used for helpful purposes. Today, we are discovering the precise biochemical pathways that underlie both disease and aging processes. We are

*continued on page 42*

PHOTOS.COM

**In the future, more drugs will be** designed to carry out precise missions at the molecular level.

# Ray Kurzweil's Plan for Cheating Death

**A cure for aging may be found in the next fifty years. The trick now is to live long enough to be there when it happens. Here's how Kurzweil is trying to do it.**

*By Terry Grossman*

I first met Ray Kurzweil in 1999 at a Foresight Institute meeting in Palo Alto. I was there to get some background information on nanotechnology for a new book I was writing. As I stood in the lunch line, a healthy-appearing man in front of me was engaged in animated conversation with a not nearly so healthy-looking second man. Their topic of conversation was vitamins and nutritional supplementation, a topic of great interest to me, a nutritionally oriented M.D.

I joined the conversation, and the healthy-looking man introduced himself as Ray Kurzweil. Ray and I continued our dialogue via e-mail after the conference ended, and a few months later, he flew from his home in Boston to Frontier Medical Institute, my longevity clinic in Denver, for a comprehensive longevity medical evaluation. We performed a comprehensive battery of tests designed to uncover any health risks he might still have so that together we could optimize his already very sophisticated program for health and longevity.

From the beginning, it was obvious that Ray would be a unique patient. I have many engineers as patients in my practice (and Ray is an engineer by training), so I am not surprised when patients come to see me with a notebook of spreadsheets detailing various data extracted from their daily lives: blood pressure, weight, cholesterol, blood sugar levels, amount of exercise, etc., carefully tabulated for several years. But all previous data collections I had seen, even those organized into Excel and meticulously graphed, paled in comparison to Ray's. His data collection was so thorough and meticulous that he could tell me what he ate for lunch on June 23, 1989 (as well as what he ate for lunch every other day for several years before that time). And not only what he ate, but the number of grams of each serving and calories consumed, as well as the number of calories he burned that day through exercise—every day for decades!

As a result, it came as less of a surprise for me to learn that Ray was taking over 200 supplement pills a day. Ray's approach had been to accurately assess his personal health risks and then quite simply to "reprogram his biochemistry." Ongoing testing indicates that he is doing a remarkable job, as measurement of his biological age in my clinic indicates that he is now almost two decades younger than his chronological age, and all of his health risks appear under optimal control.

Ray was already working on his new book, *The Singularity Is Near*, at that time, and I had just completed my first book, *The Baby Boomers' Guide to Living Forever*. It was natural that our e-mail dialogue moved into discussion of the prospects for truly radical life extension for people of all ages, including older boomers like ourselves. As our e-mails multiplied into the many thousands, we decided to organize the information and see if we had the makings of a new book that we would co-author. I created a preliminary table of contents, Ray organized the information from our e-mails, and, another 10,000 e-mails or so later, our joint book, *Fantastic Voyage: Live Long Enough to Live Forever*, was written in the midst of Ray's writing of *The Singularity Is Near*.

Ray felt that he was writing these books together as a unit and that there was synergy between them. *The Singularity Is Near* details Ray's vision of the astounding possibilities of the world of the near future as the Singularity unfolds sometime within the next few decades. In *Fantastic Voyage*, we provide readers with the information they need to live long enough and remain healthy enough to fully experience the wonders of life in the post-Singularity world. In writing these two books, Ray has painted a clear picture of the future and provided a blueprint for how to get there.

**About the Author**
Terry Grossman, M.D., is the founder and medical director of Frontier Medical Institute in Denver, Colorado. He is the co-author of *Fantastic Voyage: The Science Behind Radical Life Extension* with Ray Kurzweil. Telephone 303-233-4247; Web site www.fmiclinic.com.

*continued from page 40*
able to design drugs to carry out precise missions at the molecular level. With recently developed gene technologies, we're on the verge of being able to control how genes express themselves. Gene expression is the process by which cellular components (specifically RNA and the ribosomes) produce proteins according to a precise genetic blueprint. While every human cell contains a complete DNA sample, and thus the full complement of the body's genes, a specific cell, such as a skin cell or a pancreatic islet cell, gets its characteristics from only the fraction of genetic information relevant to that particular cell type.

Gene expression is controlled by peptides (molecules made up of sequences of up to 100 amino acids) and short RNA strands. We are now beginning to learn how these processes work. Many new therapies currently in development and testing are based on manipulating peptides either to turn off the expression of disease-causing genes or to turn on desirable genes that may otherwise not be expressed in a particular type of cell. A new technique called *RNA interference* is able to destroy the messenger RNA expressing a gene and thereby effectively turn that gene off.

Accelerating progress in biotechnology will enable us to reprogram our genes and metabolic processes to propel the fields of genomics (influencing genes), proteomics (understanding and influencing the role of proteins), gene therapy (suppressing gene expression as well as adding new genetic information), rational drug design (formulating drugs that target precise changes in disease and aging processes), as well as the therapeutic cloning of rejuvenated cells, tissues, and organs.

## The Nanotechnology Revolution

Nanotechnology promises the tools to rebuild the physical world, our bodies, and our brains, molecular fragment by molecular fragment and potentially atom by atom. We are shrinking the key features (working parts), in accordance with the law of accelerating returns, at an exponential rate (over four per linear dimension per decade, or about 100 per 3-D volume.) At this rate, the key feature sizes for most electronic and many mechanical technologies will be in the nanotechnology range—generally considered to be less than 100 nanometers (one-billionth of one meter)—by the 2020s. Electronics has already dipped below this threshold, although not yet in three-dimensional structures and not yet in structures that are capable of assembling other similar structures, an essential step before nanotechnology can reach its promised potential. Meanwhile, rapid progress has been made recently in preparing the conceptual framework and design ideas for the coming age of nanotechnology.

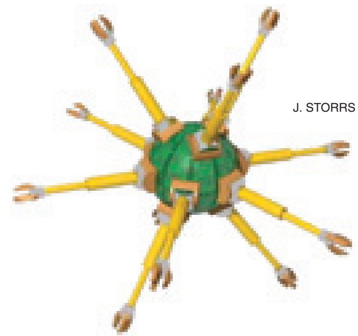

J. STORRS HALL

**A nanoscale, self-replicating robot or utility foglet** could join together with other foglets to form a solid wall that would change in shape and appearance as desired by the user. Foglet technology will allow for full-immersion virtual-reality environments by the 2030s.

Nanotechnology has expanded to include any technology in which a machine's key features are measured by fewer than 100 nanometers. Just as contemporary electronics has already quietly slipped into this nano realm, the area of biological and medical applications has already entered the era of nanoparticles, in which nanoscale objects are being developed to create more-effective tests and treatments.

In the area of testing and diagnosis, nanoparticles are being employed in experimental biological tests as tags and labels to greatly enhance sensitivity in detecting substances such as proteins. Magnetic nanotags can be used to bind with antibodies that can then be read using magnetic probes while still inside the body. Successful experiments have been conducted with gold nanoparticles that are bound to DNA segments and can rapidly test for specific DNA sequences in a sample. Small nanoscale beads called quantum dots can be programmed with specific codes combining multiple colors, similar to a color bar code, that can facilitate tracking of substances through the body.

In the future, nanoscale devices will run hundreds of tests simultaneously on tiny samples of a given substance. These devices will allow extensive tests to be conducted on nearly invisible samples of blood.

In the area of treatment, a particularly exciting application of this technology is the harnessing of nanoparticles to deliver medication to specific sites in the body. Nanoparticles can guide drugs into cell walls and through the blood-brain barrier. Nanoscale packages can be designed to hold drugs, protect them through the gastrointestinal tract, ferry them to specific locations, and then release them in sophisticated ways that can be influenced and controlled, wirelessly, from outside the body.

Nanotherapeutics in Alachua, Florida, has developed a biodegradable polymer only several nanometers thick that uses this approach. Meanwhile, scientists at McGill University in Montreal have demonstrated a nanopill with structures in the 25 to 45 nanometer range. The nanopill is small enough to pass through the cell wall and deliver medications directly to targeted structures inside the cell.

MicroCHIPS of Bedford, Massachusetts, has developed a computerized device that is implanted under the skin and delivers precise mixtures of medicines from hundreds of nanoscale wells inside the device. Future versions of the device are expected to be able to measure blood levels of substances such as glucose. The system could be used as an artificial pancreas, releasing precise amounts of insulin based on the blood glucose response. The system would also be capable of simulating any other hormone-producing organ. If trials go smoothly, the system could be on the market by 2008. Another innovative proposal is to guide nanoparticles (probably composed of gold) to a tumor site and then heat them with infrared beams to destroy the cancer cells.

The revolution in nanotechnology will allow us to do a great deal more than simply treat disease. Ultimately, nanotech will enable us to redesign and rebuild not only our bodies and brains, but also the world with which we interact. The full realization of nanotechnology, however, will lag behind the biotechnology revolution by about one decade. But by the mid to late 2020s, the effects of the nanotech revolution will be widespread and obvious.

### Nanotechnology and The Human Brain

The most important and radical application particularly of circa-2030 nanobots will be to expand our minds through the merger of biological and nonbiological, or "machine," intelligence. In the next 25 years, we will learn how to augment our 100 trillion very slow interneuronal connections with high-speed virtual connections via nanorobotics. This will allow us to greatly boost our pattern-recognition abilities, memories, and overall thinking capacity, as well as to directly interface with powerful forms of computer intelligence. The technology will also provide wireless communication from one brain to another.

PHOTOS.COM

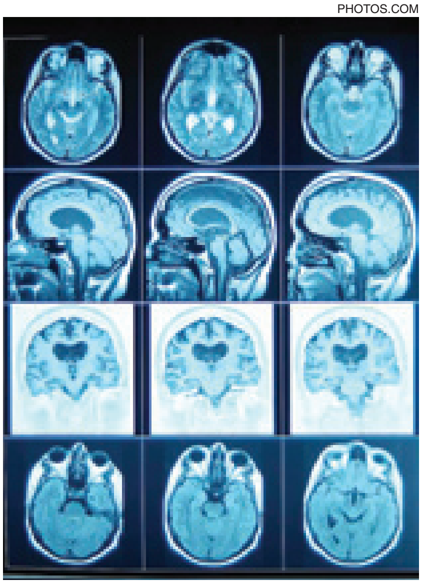

**Our brains today are** relatively fixed in design. As humanity's artificial-intelligence (AI) capabilities begin to upstage our human intelligence at the end of the 2030s, we will be able to move beyond the basic architecture of the brain's neural regions. But artificial Intelligence will be based, at least in part, on a human-made version of a fully functional human brain.

In other words, the age of telepathic communication is almost upon us.

Our brains today are relatively fixed in design. Although we do add patterns of interneuronal connections and neurotransmitter concentrations as a normal part of the learning process, the current overall capacity of the human brain is highly constrained. As humanity's artificial-intelligence (AI) capabilities begin to upstage our human intelligence at the end of the 2030s, we will be able to move beyond the basic architecture of the brain's neural regions.

Brain implants based on massively distributed intelligent nanobots will greatly expand our memories and otherwise vastly improve all of our sensory, pattern-recognition, and cognitive abilities. Since the nanobots will be communicating with one another, they will be able to create any set of new neural connections, break existing connections (by suppressing neural firing), create new hybrid biological and computer networks, and add completely mechanical networks, as well as interface intimately with new computer programs and artificial intelligences.

The implementation of artificial intelligence in our biological systems will mark an evolutionary leap forward for humanity, but it also implies we will indeed become more "machine" than "human." Billions of nanobots will travel through the bloodstream in our bodies and brains. In our bodies, they will destroy pathogens, correct DNA errors, eliminate toxins, and perform many other tasks to enhance our physical well-being. As a result, we will be able to live indefinitely without aging.

In our brains, nanobots will interact with our biological neurons. This will provide full-immersion virtual reality incorporating all of the

ILLUSTRATION: PATRICK TUCKER, WFS / IMAGES: PHOTOS.COM

senses, as well as neurological correlates of our emotions, from within the nervous system. More importantly, this intimate connection between our biological thinking and the machine intelligence we are creating will profoundly expand human intelligence.

Warfare will move toward nanobot-based weapons, as well as cyberweapons. Learning will first move online, but once our brains are fully online we will be able to download new knowledge and skills. The role of work will be to create knowledge of all kinds, from music and art to math and science. The role of play will also be to create knowledge. In the future, there won't be a clear distinction between work and play.

### The Robotic Revolution

Of the three technological revolutions underlying the Singularity (genetic, nano-mechanical, and robotic), the most profound is robotic or, as it is commonly called, the *strong artificial intelligence* revolution. This refers to the creation of computer thinking ability that exceeds the thinking ability of humans. We are very close to the day when fully biological humans (as we now know them) cease to be the dominant intelligence on the planet. By the end of this century, computational or mechanical intelligence will be trillions of trillions of times more powerful than unaided human brain power. I argue that computer, or as I call it *nonbiological intelligence*, should still be considered human since it is fully derived from human–machine civilization and will be based, at least in part, on a human-made version of a fully functional human brain. The merger of these two worlds of intelligence is not merely a merger of biological and mechanical thinking mediums, but also (and more importantly) a merger of method and organizational thinking that will expand our minds in virtually every imaginable way.

Biological human thinking is limited to $10^{16}$ calculations per second (cps) per human brain (based on neuromorphic modeling of brain regions) and about $10^{26}$ cps for all human brains. These figures will not appreciably change, even with bioengineering adjustments to our genome. The processing capacity of nonbiological intelligence or strong AI, in contrast, is growing at an exponential rate (with the rate itself increasing) and will vastly exceed biological intelligence by the mid-2040s.

Artificial intelligence will necessarily exceed human intelligence for several reasons. First, machines can share knowledge and communicate with one another far more efficiently than can humans. As humans, we do not have the means to exchange the vast patterns of interneuronal connections and neurotransmitter-concentration levels that comprise our learning, knowledge, and skills, other than through slow, language-based communication.

Second, humanity's intellectual skills have developed in ways that have been evolutionarily encouraged in natural environments. Those skills, which are primarily based on our abilities to recognize and extract meaning from patterns, enable us to be highly proficient in certain tasks, such as distinguishing faces, identifying objects, and recognizing language sounds. Unfortunately, our brains are less well-suited for dealing with more-complex patterns, such as those that exist in financial, scientific, or product data. The application of computer-based techniques will allow us to fully master pattern-recognition paradigms. Finally, as human knowledge migrates to the Web, machines will demonstrate increased proficiency in reading, understanding, and synthesizing all human–machine information.

### The Chicken or the Egg

A key question regarding the Singularity is whether the "chicken" (strong AI) or the "egg" (nanotechnology) will come first. In other words, will strong AI lead to full nanotechnology (molecular-manufacturing assemblers that can turn information into physical products), or will full nanotechnology lead to strong AI?

The logic of the first premise is that strong AI would be in a position to solve any remaining design problems required to implement full nanotechnology. The second premise is based on the assumption that hardware requirements for strong AI will be met by nanotechnology-based computation. Likewise, the software requirements for engineering strong AI would be facilitated by nanobots. These microscopic machines will allow us to create highly detailed scans of human brains along with diagrams of how the human brain is able to do all the wonderful things that have long mystified us, such as create meaning, contextualize information, and experience emotion. Once we fully understand how the brain functions, we will be able to recreate the phenomenon of human thinking in machines. We will endow computers, already superior to us in the performance of mechanical tasks, with lifelike intelligence.

Progress in both areas (nano and robotic) will necessarily use our most-advanced tools, so advances in each field will simultaneously facilitate the other. However, I do expect that the most important nanotechnological breakthroughs will emerge prior to strong AI, but only by a few years (around 2025 for nanotechnology and 2029 for strong AI).

As revolutionary as nanotechnology will be, strong AI will have far more profound consequences. Nanotechnology is powerful but not necessarily intelligent. We can devise ways of at least trying to manage the enormous powers of nanotechnology, but superintelligence by its nature cannot be controlled.

The nano/robotic revolution will also force us to reconsider the very definition of human. Not only will we be surrounded by machines that will display distinctly human characteristics, but also we will be less human from a literal standpoint.

Despite the wonderful future potential of medicine, real human longevity will only be attained when we move away from our biological bodies entirely. As we move toward a software-based existence, we will gain the means of "backing ourselves up" (storing the key patterns underlying our knowledge, skills, and personality in a digital setting) thereby enabling a virtual immortality. Thanks to nanotechnology, we will have bodies that we can not only modify, but also change into new forms at will. We will be able to quickly change our bodies in full-immersion virtual-reality environments incorporating all of the senses during the 2020s and in real reality in the 2040s.

### Implications of the Singularity

What will be the nature of human experience once computer intelligence predominates? What are the implications for the human–machine civilization when strong AI and nanotechnology can create any product, any situation, any environment that we can imagine at will? I stress the role of imagination here because we will still be constrained in our creations to what we can imagine. But our tools for bringing imagination to life are growing exponentially more powerful.

People often go through three stages in considering the impact of future technology: awe and wonderment at its potential to overcome age-old problems, then a sense of dread at the grave new dangers that accompany these novel technologies, followed finally by the realization that the only viable and responsible path is to set a careful course that can realize the benefits while managing the dangers.

My own expectation is that the creative and constructive applications of these technologies will dominate, as I believe they do today.

However, we need to vastly increase our investment in developing specific defensive technologies. We are at the critical stage where we need to directly implement defensive technologies for nanotechnology during the late teen years of this century.

I believe that a narrow relinquishment of the development of certain capabilities needs to be part of our ethical response to the dangers of twenty-first-century technological challenges. For example, Bill Joy and I wrote a joint op-ed piece in the *New York Times* recently, criticizing the publication of the 1918 flu genome on the Web because it constitutes a dangerous blueprint. Another constructive example of this are the ethical guidelines proposed by the Foresight Institute: namely, that nanotechnologists agree to relinquish the development of physical entities that can self-replicate in a natural environment free of any human control or override mechanism. However, deciding in favor of too many limitations and restrictions would undermine economic progress and is ethically unjustified, given the opportunity to alleviate disease, overcome poverty, and clean up the environment.

We don't have to look past today to see the intertwined promise and peril of technological advancement. Imagine describing the dangers (atomic and hydrogen bombs, for one thing) that exist today to people who lived a couple of hundred years ago. They would think it mad to take such risks. But how many people in 2006 would really want to go back to the short, brutish, disease-filled, poverty-stricken, disaster-prone lives that 99% of the human race struggled through two centuries ago?

We may romanticize the past, but up until fairly recently most of humanity lived extremely fragile lives in which one all-too-common misfortune could spell disaster. Two hundred years ago, life expectancy for females in the record-holding country (Sweden) was roughly 35 years—very brief compared with the longest life expectancy today, almost 85 years for Japanese women. Life expectancy for males was roughly 33 years, compared with the current 79 years. Half a day was often required to prepare an evening meal, and hard labor characterized most human activity. There were no social safety nets. Substantial portions of our species still live in this precari-

ILLUSTRATION: PATRICK TUCKER, WFS / IMAGES: PHOTOS.COM

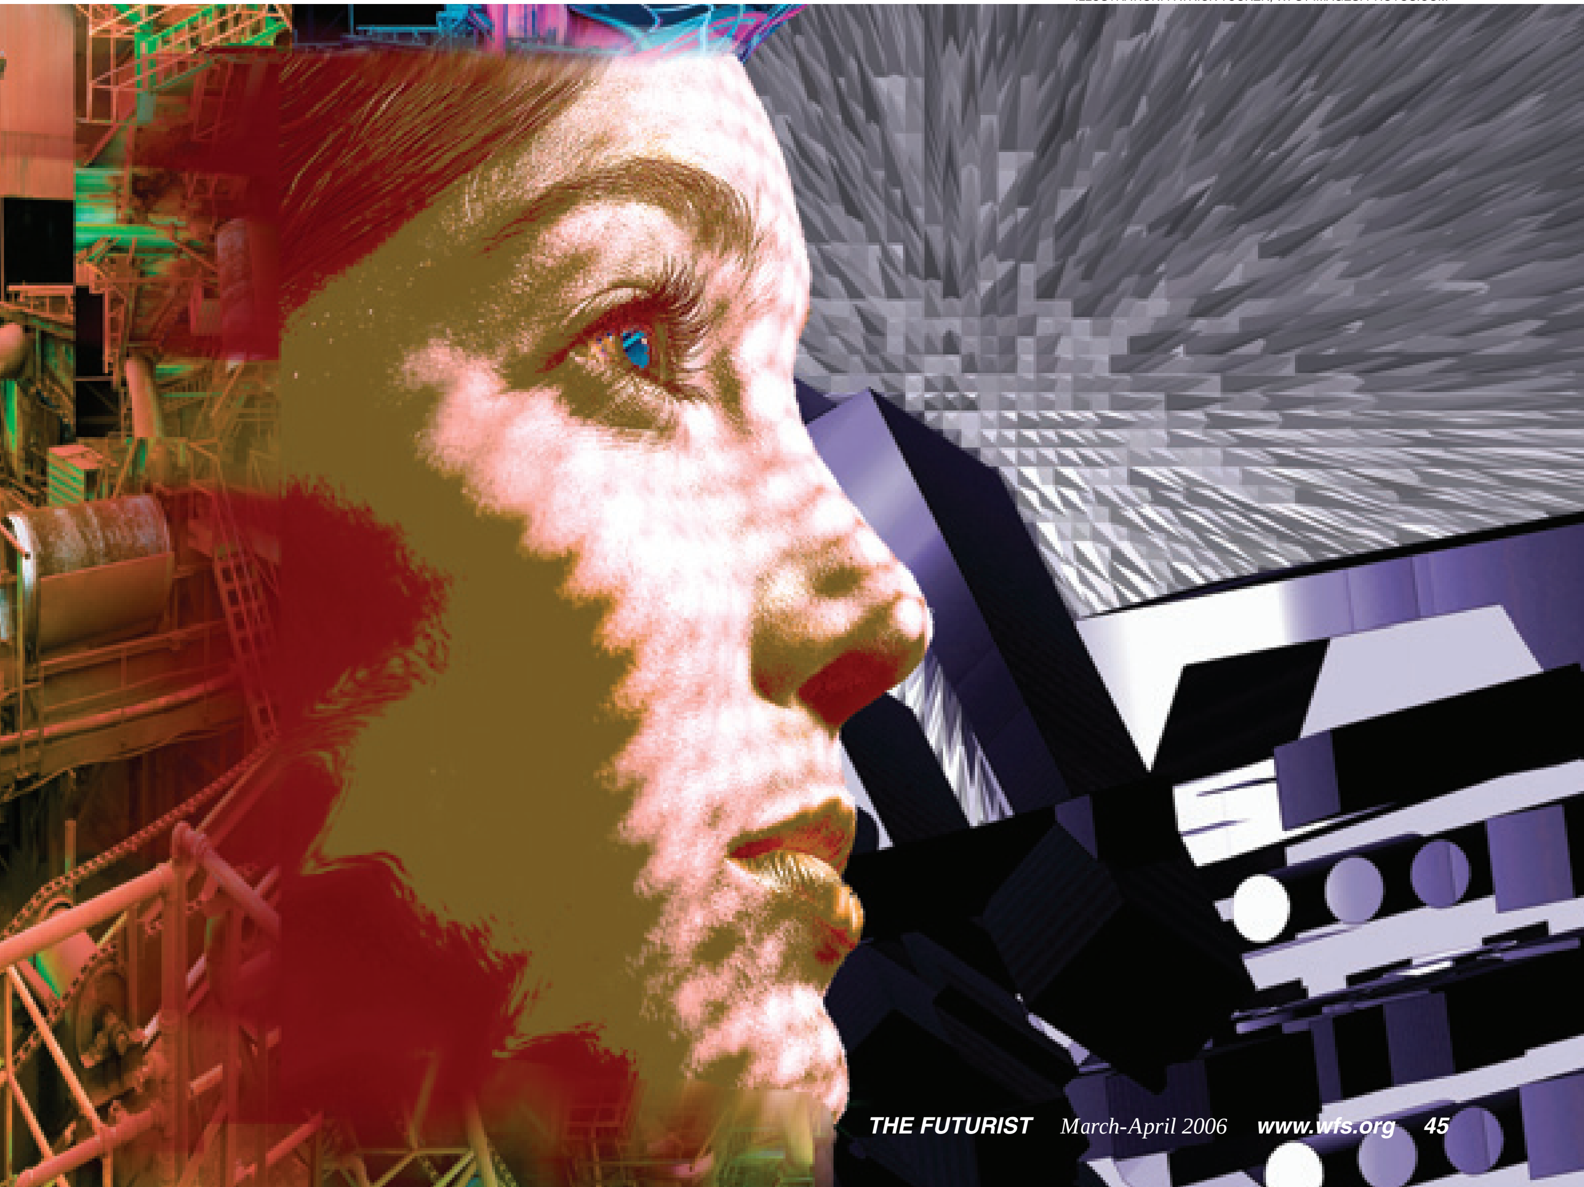

ous way, which is at least one reason to continue technological progress and the economic improvement that accompanies it. Only technology, with its ability to provide orders of magnitude of advances in capability and affordability, has the scale to confront problems such as poverty, disease, pollution, and the other overriding concerns of society today. The benefits of applying ourselves to these challenges cannot be overstated.

As the Singularity approaches, we will have to reconsider our ideas about the nature of human life and redesign our human institutions. Intelligence on and around Earth will continue to expand exponentially until we reach the limits of matter and energy to support intelligent computation. As we approach this limit in our corner of the galaxy, the intelligence of our civilization will expand outward into the rest of the universe, quickly reaching the fastest speed possible. We understand that speed to be the speed of light, but there are suggestions that we may be able to circumvent this apparent limit (conceivably by taking shortcuts through "wormholes," or hypothetical shortcuts through space and time).

A common view is that science has consistently been correcting our overly inflated view of our own significance. Stephen Jay Gould said, "The most important scientific revolutions all include, as their only common feature, the dethronement of human arrogance from one pedestal after another of previous convictions about our centrality in the cosmos."

Instead, it turns out we are central. Our ability to create models—virtual realities—in our brains, combined with our modest-looking thumbs, has been sufficient to usher in another form of evolution: technology. That development enabled the persistence of the accelerating pace that started with biological evolution. It may continue until the entire universe is at our fingertips.

---

**About the Author**

Ray Kurzweil is a scientist, inventor, and entrepreneur. He has received 12 honorary doctorates in science, engineering, music, and humane letters from Rensselaer Polytechnic Institute, Hofstra University, and other leading colleges and universities. He has been inducted into the National Inventors Hall of Fame and received the 1999 National Medal of Technology, among numerous other awards. His Web site, www.KurzweilAI.net, has more than a million readers and includes a free daily e-newsletter. This article draws, in part, on his most recent book, *The Singularity Is Near: When Humans Transcend Biology* (Viking, 2005).

## Technology and Human Enhancement

*By John Smart*

I have a few differences of opinion with Kurzweil about the coming Singularity.

I think he is being overly optimistic about biotechnology's ability to create substantially better biological human beings. While we'll certainly learn to push human capacities to their natural limits in coming decades, I see nothing on the horizon that would allow us to exceed those limits. Biology seems far too frail, slow, complex, and well defended (both at the molecular level and with regard to social custom) for that to be plausible within any reasonable time frame. Furthermore, by the time we are able to substantially improve our biology, we probably won't want to, as there will be far more interesting and powerful technological environments available to us instead. This points to the importance of understanding the relative accelerations of various technologies (in this case, biological vs. technological).

In his book, Kurzweil makes a major contribution to the literature on acceleration studies by clearly explaining technological acceleration curves. These acceleration curves show that the longer we use a technology, the more we get out of it: We use less energy, space, and time, and we get more capacity for less cost. Technological acceleration curves are a little-understood area, but thanks to pioneers like Kurzweil, interest and research in the field are advancing.

The common notion that the "future can't be predicted" is demonstrably false with regard to a wide number of accelerating physical-computational trends, even though we do not yet know specifically how those technologies will be implemented. We can no longer ignore the profound technological changes occurring all around us.

It's also time we acknowledged the slowness of human biology compared with our technological progeny. Our machines are increasingly exceeding us in the performance of more and more tasks, from guiding objects like missiles or satellites to assembling other machines. They are merging with us ever more intimately and are learning how to reconfigure our biology in new and significantly faster technological domains.

Something very interesting is happening, and human beings are selective catalysts, not absolute controllers, of this process. Let us face this openly, and investigate it actively, so that we may guide these developments as wisely as possible.

---

**About the Author**

John Smart is a developmental systems theorist and the president of the Acceleration Studies Foundation, 2227 Amirante, San Pedro, California 90732. Telephone 310-831-4191; Web site www.accelerating.org.

## Runaway Artificial Intelligence?

*By J. Storrs Hall*

Some years ago, I reviewed Ray Kurzweil's earlier book, *The Age of Spiritual Machines*, for the Foresight Nanotech Institute's newsletter. Shortly thereafter I met him in person at a Foresight event, and he had read the review. He told me, "Of all the people who reviewed my book, you were the only one who said I was too conservative!"

*The Singularity Is Near* is very well researched, and in general, Kurzweil's predictions are about as good as it's possible to get for things that far in advance. I still think he's too conservative in one specific area: Synthetic computer-based artificial intelligence will become available well before nanotechnology makes neuron-level brain scans possible in the 2020s.

What's happening is that existing technologies like functional MRI are beginning to give us a high-level functional block diagram of the brain's processes. At the same time, the hardware capable of running a strong, artificially intelligent computer, by most estimates, is here now, though it's still pricey.

Existing AI software techniques can build programs that are experts in any well-defined field. The breakthroughs necessary for such programs to learn for themselves could happen easily in the next decade—one or two decades before Kurzweil predicts.

Kurzweil finesses the issue of runaway AI by proposing a pathway where machine intelligence is patterned after human brains, so that they would have our morals and values built in. Indeed, this would clearly be the wise and prudent course. Unfortunately, it seems all too likely that a shortcut exists without that kind of safeguard. Corporations already use huge computer systems for data mining and decision support that employ sophisticated algorithms no human manager understands. It's a very short step to having such a system make better decisions than the managers do, as far as the corporation's bottom line is concerned.

The Singularity may mean different things to different people. To me, it is that point where intelligences significantly greater than our own control so many of the essential processes that figure in our lives that mere humans can't predict what happens next. This future may be even nearer than Ray Kurzweil has predicted.


**About the Author**

J. Storrs Hall is chief scientist of Nanorex Inc. and a fellow of the Molecular Engineering Research Institute. His address is Nanorex Inc., 2738 Turtle Ridge Drive, Bloomfield Hills, Michigan 48302.

He is the author of *Nanofuture: What's Next for Nanotechnology* (Prometheus Books, 2005, $28), which is available from the Futurist Bookshelf, www.wfs.org/bkshelf.htm.


## Nanofactories, Gang Wars, and "Feelies"

*By Damien Broderick*

BETH GWINN

A quarter century ago, we'd have laughed at the prospect of "Dick Tracy" cell phones with cameras; now they're everywhere, and nobody noticed after the first few days. So the jump to the idea of a Singularity is not really extraordinary. But should we really expect ever more substantial changes to follow the same accelerating, headlong pace?

It's reasonable to expect affordable computers to be smaller and more powerful, 1,000 times improved in a decade, one million times in 20 years, one billion in 30. By then, some machines might have capabilities to rival the human mind. A new intelligent species might share the planet with us.

In addition, developing technologies such as molecular manufacture—nanotechnology—will allow the very engines of productivity to be copied cheaply and distributed widely. If that happens the gap between rich and poor might diminish. However, it will only occur if we find ways to prevent portable nanofactories from making lethal weapons available to any child or psychopath. We'll be able to solve most of the problems that currently vex us—global warming (to the extent that it's caused by humans), water and food shortages, provision of clean, cheap power, and so on.

There is a scary downside that I discussed nearly a decade ago in my book *The Spike:* Dirt-cheap molecular manufacturing may end poverty and strife, but there exists a risk that a world of lotus-eaters will degenerate into gang wars among those for whom life retains no discipline or meaning outside of arbitrary local status and violence. People (young men especially) with full bellies gained effortlessly, but lacking meaning in their lives, often find purpose in ganging up on each other in fits of murderous primate chest-pounding. Making Huxleian soma, or "feelies," the opiate of the people might help, but that, too, is a sickening prospect.

On the other hand, those strictly unforeseeable and mysterious changes captured in the word "Singularity" are likely to overwhelm and surpass such predictable downsides of any technological utopia or dystopia. The eeriest aspect of accelerating change is that we ourselves, and our children, will be the ones soaking in it. The sooner we start thinking seriously about the prospect, the better prepared we'll be.


**About the Author**

Damien Broderick is a senior fellow in the Department of English and Cultural Studies at the University of Melbourne, Australia. His futurist books include *The Spike* (1997, rev. 2001), *The Last Mortal Generation* (1999), and *Ferocious Minds* (2005). His novels dealing with the Singularity include *Transcension* (2002), *The Hunger of Time* (with Rory Barnes, 2003), *Godplayers* (2005), and *K-Machines* (forthcoming, 2006). He lives in Melbourne, Australia, and San Antonio, Texas.

# Techno-Utopia and Human Values

By Richard Eckersley

I have sometimes asked audiences if they are inspired or excited by the sort of techno-utopian vision represented by the Singularity; almost no one is. In my surveys over the past decade, I found dwindling minorities of young people (one-fifth to one-quarter) believe in the sort of technical fixes to human problems that Ray Kurzweil champions, while an increased majority (about three-quarters) believe science and technology are alienating people from each other and from nature.

The question I ask is, why? Why pursue this future? I don't pose this question dismissively, or derogatorily, but out of genuine curiosity and a desire for an open, honest conversation. I'm skeptical of arguments that say pre-technological humans led short, nasty, and brutish lives. Yes, life expectancy was lower—mainly because of high rates of infant mortality—but those who survived often lived socially and spiritually rich lives. It doesn't make evolutionary sense to believe humans lived in misery until we discovered technological progress. Animals in the wild don't live that way, and humans have been, for most of their history, animals in the wild.

The future world that Kurzweil describes bears almost no relationship to human well-being that I am aware of. In essence, human health and happiness comes from being connected and engaged, from being suspended in a web of relationships and interests—personal, social, and spiritual—that give meaning to our lives. The intimacy and support provided by close personal relationships seem to matter most; isolation exacts the highest price. The need to belong is more important than the need to be rich. Meaning matters more than money and what it buys.

We are left with the matter of destiny: It is our preordained fate, Kurzweil suggests, to advance technologically "until the entire universe is at our fingertips." The question then becomes, preordained by whom or what? Biological evolution has not set this course for us. Is technology itself the planner? Perhaps it will eventually be, but not yet. Is God the entity doing the ordaining? A lot of religious people would have something to say about that, and are likely to strenuously, and even violently, oppose what the Singularity promises, as I have argued before (THE FUTURIST, November-December 2001).

We are left to conclude that we will do this because it is we who have decided it is our destiny. But we have made no such decision, really, as the observations with which I began this commentary show.


**About the Author**

Richard Eckersley is a fellow at the National Centre for Epidemiology and Population Health at the Australian National University, Canberra, Australia, and author of *Well & Good: Morality, Meaning and Happiness* (Text Publishing, Melbourne, 2005). E-mail Richard.Eckersley@anu.edu.au.


**Feedback:** Send your comments about this article to letters@wfs.org.

***Ray Kurzweil reponds:***

Richard Eckersley's idyllic notion of human life hundreds of years ago belies our scientific knowledge of history. Two hundred years ago, there was no understanding of sanitation, so bacterial infections were rampant. There were no antibiotics and no social safety nets, so an infectious disease was a disaster plunging a family into desperation. Thomas Hobbes's characterization in 1651 of human life as solitary, poor, nasty, brutish, and short was on the mark. Even ignoring infant mortality, life expectancy was in the 30s only a couple of hundred years ago. Schubert's and Mozart's deaths at 31 and 35 respectively were typical.

Eckersley bases his romanticized idea of ancient life on communication and the relationships fostered by communication. But much of modern technology is directed at just this basic human need. The telephone allowed people to be together even if far apart geographically. The Internet is the quintessential communication technology. Social networks and the panoply of new ways to make connection are creating communities based on genuine common interests rather than the accident of geography.

This decentralized electronic communication is also highly democratizing. In a book I wrote in the mid-1980s, I predicted the demise of the Soviet Union from the impact of the then-emerging communication networks, and that is indeed what happened in the early 1990s. The democracy movement we saw in the 1990s and since is similarly fueled by our unprecedented abilities to stay in touch.

If Eckersley really sticks to his own philosophy, he won't be around for very long to influence the debate. But I hope that he will take advantage of the life extension—and enhancement—technologies that will emerge in the decades ahead, so that we can continue this dialogue through this century and beyond. □