THE *NEW YORK TIMES* BESTSELLER

THE SINGULARITY IS NEAR

WHEN HUMANS TRANSCEND BIOLOGY

"Startling in scope and bravado."
—Janet Maslin, *The New York Times*

RAY KURZWEIL

AUTHOR OF *THE AGE OF SPIRITUAL MACHINES* AND *HOW TO CREATE A MIND*

A radical and optimistic view of the future course of human development by "the best person I know at predicting the future of artificial intelligence" (Bill Gates).

At the onset of the twenty-first century, humanity stands on the verge of the most transforming and thrilling period in its history. It will be an era in which the very nature of what it means to be human will be both enriched and challenged as our species breaks the shackles of its genetic legacy and achieves inconceivable heights of intelligence, material progress, and longevity. While the social and philosophical ramifications of these changes will be profound, and the threats they pose considerable, celebrated futurist Ray Kurzweil presents a view of the coming age that is both a dramatic culmination of centuries of technological ingenuity and a genuinely inspiring vision of our ultimate destiny.

"Artfully envisions a breathtakingly better world." —*Los Angeles Times*

"Elaborate, smart, and persuasive." —*The Boston Globe*

"A pleasure to read." —*The Wall Street Journal*

A Penguin Book
Science

U.S. $24.00
CAN. $32.00

ISBN 978-0-14-303788-0

52400

EAN 9 780143 037880

WWW.PENGUIN.COM
COVER DESIGN: HERB THORNBY

Contents

CHAPTER FOUR

Achieving the Software of Human Intelligence: How to Reverse Engineer the Human Brain 143

CHAPTER NINE

Response to Critics 427

PROLOGUE

The Power of Ideas

> I do not think there is any thrill that can go through the human heart like that felt by the inventor as he sees some creation of the brain unfolding to success.
>
> —NIKOLA TESLA, 1896, INVENTOR OF ALTERNATING CURRENT

At the age of five, I had the idea that I would become an inventor. I had the notion that inventions could change the world. When other kids were wondering aloud what they wanted to be, I already had the conceit that I knew what I was going to be. The rocket ship to the moon that I was then building (almost a decade before President Kennedy's challenge to the nation) did not work out. But at around the time I turned eight, my inventions became a little more realistic, such as a robotic theater with mechanical linkages that could move scenery and characters in and out of view, and virtual baseball games.

Having fled the Holocaust, my parents, both artists, wanted a more worldly, less provincial, religious upbringing for me.[1] My spiritual education, as a result, took place in a Unitarian church. We would spend six months studying one religion—going to its services, reading its books, having dialogues with its leaders—and then move on to the next. The theme was "many paths to the truth." I noticed, of course, many parallels among the world's religious traditions, but even the inconsistencies were illuminating. It became clear to me that the basic truths were profound enough to transcend apparent contradictions.

At the age of eight, I discovered the Tom Swift Jr. series of books. The plots of all of the thirty-three books (only nine of which had been published when I started to read them in 1956) were always the same: Tom would get himself into a terrible predicament, in which his fate and that of his friends, and often the rest of the human race, hung in the balance. Tom would retreat to his basement lab and think about how to solve the problem. This, then, was the dramatic tension in each book in the series: what ingenious idea would Tom and

CHAPTER ONE

The Six Epochs

Everyone takes the limits of his own vision for the limits of the world.

—ARTHUR SCHOPENHAUER

I am not sure when I first became aware of the Singularity. I'd have to say it was a progressive awakening. In the almost half century that I've immersed myself in computer and related technologies, I've sought to understand the meaning and purpose of the continual upheaval that I have witnessed at many levels. Gradually, I've become aware of a transforming event looming in the first half of the twenty-first century. Just as a black hole in space dramatically alters the patterns of matter and energy accelerating toward its event horizon, this impending Singularity in our future is increasingly transforming every institution and aspect of human life, from sexuality to spirituality.

What, then, is the Singularity? It's 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. Although neither utopian nor dystopian, this epoch will transform the concepts that we rely on to give meaning to our lives, from our business models to the cycle of human life, including death itself. Understanding the Singularity will alter our perspective on the significance of our past and the ramifications for our future. To truly understand it inherently changes one's view of life in general and one's own particular life. I regard someone who understands the Singularity and who has reflected on its implications for his or her own life as a "singularitarian."[1]

I can understand why many observers do not readily embrace the obvious implications of what I have called the law of accelerating returns (the inherent acceleration of the rate of evolution, with technological evolution as a continuation of biological evolution). After all, it took me forty years to be able to see what was right in front of me, and I still cannot say that I am entirely comfortable with all of its consequences.

The key idea underlying the impending Singularity is that the pace of change of our human-created technology is accelerating and its powers are

The Intuitive Linear View Versus the Historical Exponential View

> When the first transhuman intelligence is created and launches itself into recursive self-improvement, a fundamental discontinuity is likely to occur, the likes of which I can't even begin to predict.
>
> —MICHAEL ANISSIMOV

In the 1950s John von Neumann, the legendary information theorist, was quoted as saying that "the ever-accelerating progress of technology . . . gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue."[3] Von Neumann makes two important observations here: *acceleration* and *singularity*. The first idea is that human progress is exponential (that is, it expands by repeatedly *multiplying* by a constant) rather than linear (that is, expanding by repeatedly *adding* a constant).

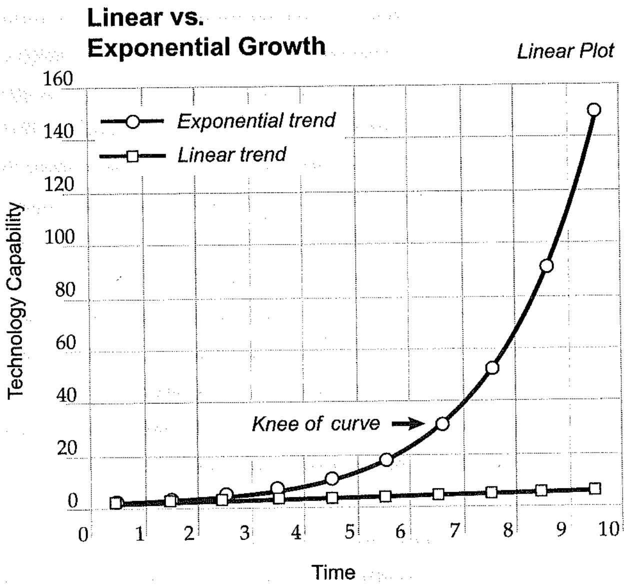

Linear versus exponential: **Linear growth is steady; exponential growth becomes explosive.**

The second is that exponential growth is seductive, starting out slowly and virtually unnoticeably, but beyond the knee of the curve it turns explosive and profoundly transformative. The future is widely misunderstood. Our forebears expected it to be pretty much like their present, which had been pretty much most other elements), giving rise to complicated, information-rich, three-dimensional structures.

The rules of our universe and the balance of the physical constants that govern the interaction of basic forces are so exquisitely, delicately, and exactly appropriate for the codification and evolution of information (resulting in increasing complexity) that one wonders how such an extraordinarily unlikely situation came about. Where some see a divine hand, others see our own hands—namely, the anthropic principle, which holds that only in a universe that allowed our own evolution would we be here to ask such questions.[7] Recent theories of physics concerning multiple universes speculate that new universes are created on a regular basis, each with its own unique rules, but that most of these either die out quickly or else continue without the evolution of any interesting patterns (such as Earth-based biology has created) because their rules do not support the evolution of increasingly complex forms.[8] It's hard to imagine how we could test these theories of evolution applied to early cosmology, but it's clear that the physical laws of our universe are precisely what they need to be to allow for the evolution of increasing levels of order and complexity.[9]

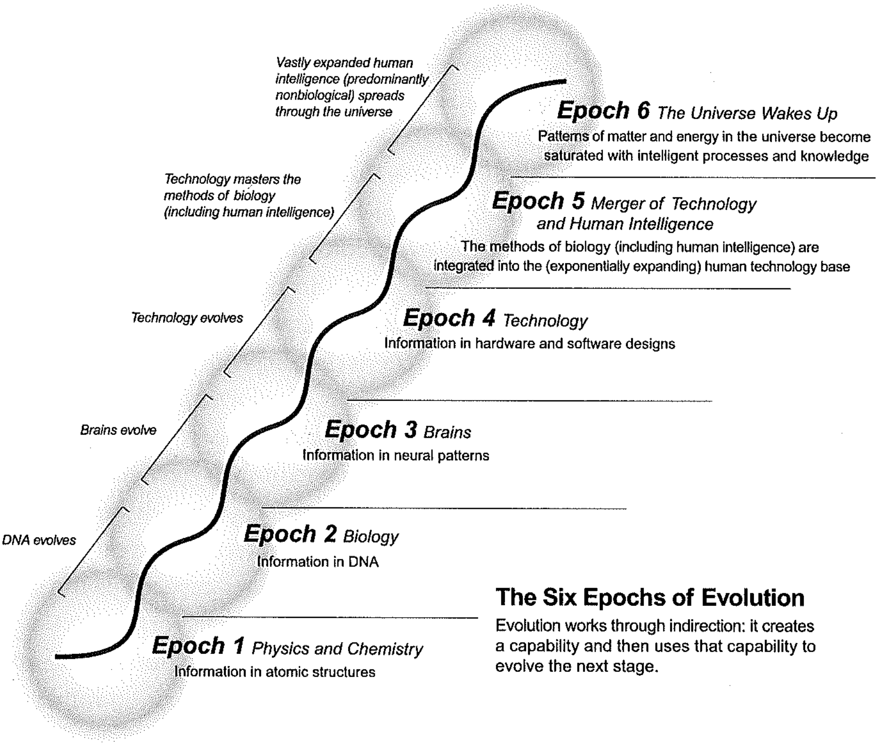

The Six Epochs of Evolution
Evolution works through indirection: it creates a capability and then uses that capability to evolve the next stage.

years ago) and the y-axis (the paradigm-shift time) on logarithmic scales, we find a reasonably straight line (continual acceleration), with biological evolution leading directly to human-directed development.[11]

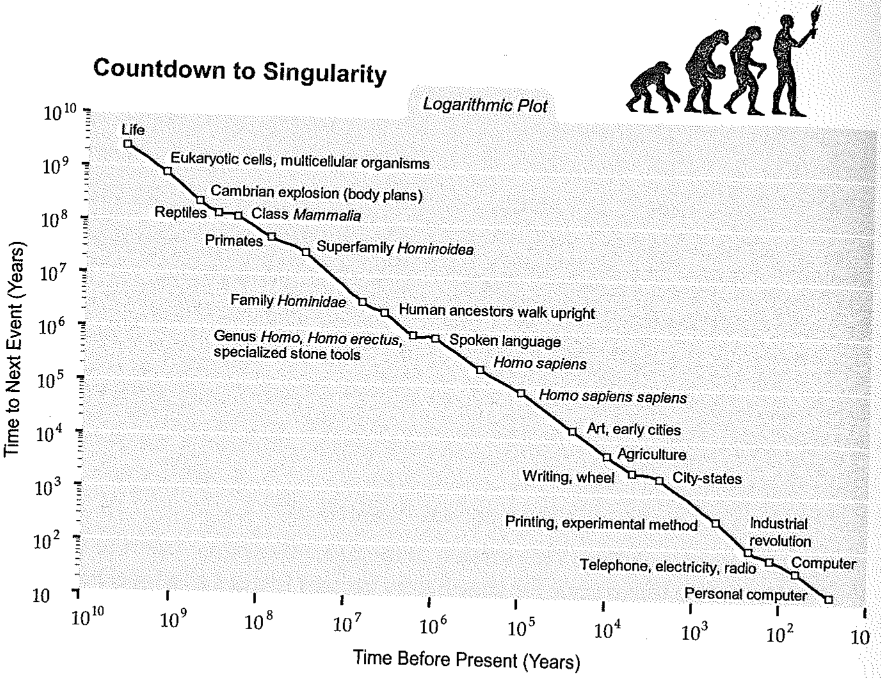

Countdown to Singularity: **Biological evolution and human technology both show continual acceleration, indicated by the shorter time to the next event (two billion years from the origin of life to cells; fourteen years from the PC to the World Wide Web).**

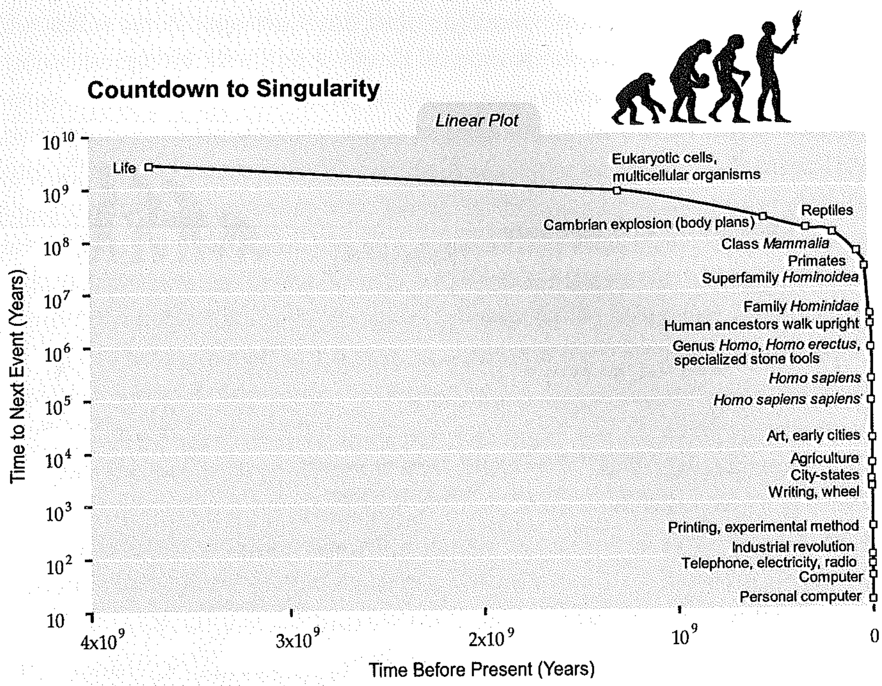

Linear view of evolution: **This version of the preceding figure uses the same data but with a linear scale for time before present instead of a logarithmic one. This shows the acceleration more dramatically, but details are not visible. From a linear perspective, most key events have just happened "recently."**

The above figures reflect my view of key developments in biological and technological history. Note, however, that the straight line, demonstrating the continual acceleration of evolution, does not depend on my particular selection of events. Many observers and reference books have compiled lists of important events in biological and technological evolution, each of which has its own idiosyncrasies. Despite the diversity of approaches, however, if we combine lists from a variety of sources (for example, the *Encyclopaedia Britannica*, the American Museum of Natural History, Carl Sagan's "cosmic calendar," and others), we observe the same obvious smooth acceleration. The following plot combines fifteen different lists of key events.[12] Since different thinkers assign different dates to the same event, and different lists include similar or overlapping events selected according to different criteria, we see an expected "thickening" of the trend line due to the "noisiness" (statistical variance) of this data. The overall trend, however, is very clear.

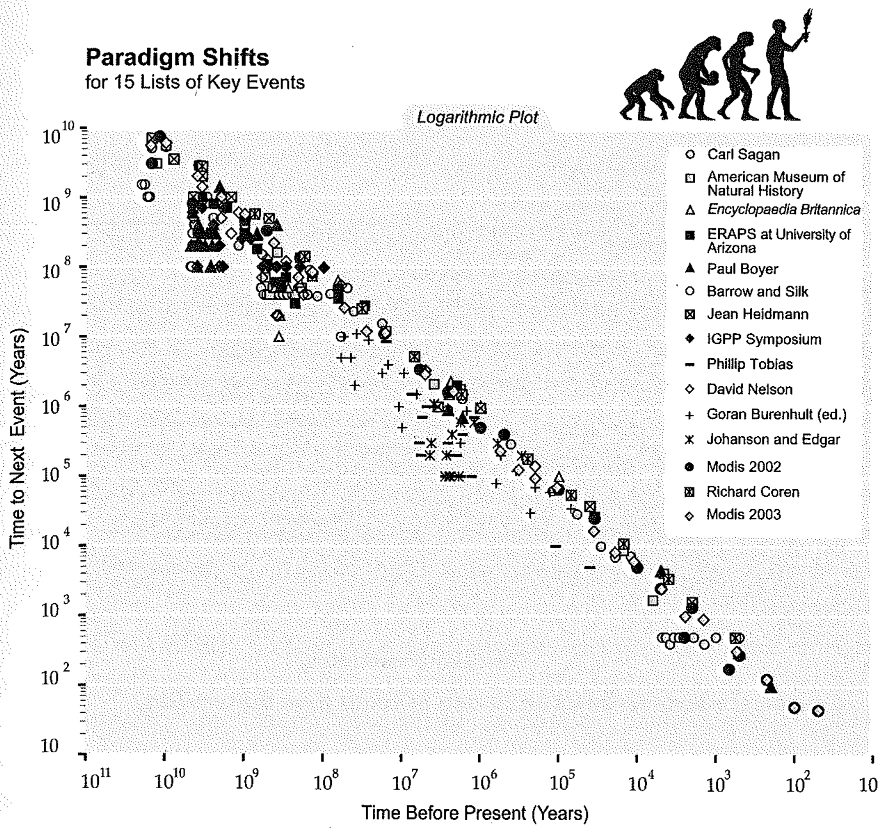

Fifteen views of evolution: **Major paradigm shifts in the history of the world, as seen by fifteen different lists of key events. There is a clear trend of smooth acceleration through biological and then technological evolution.**

Physicist and complexity theorist Theodore Modis analyzed these lists and determined twenty-eight clusters of events (which he called canonical milestones) by combining identical, similar, and/or related events from the different lists.[13] This process essentially removes the "noise" (for example, the variability of dates between lists) from the lists, revealing again the same progression:

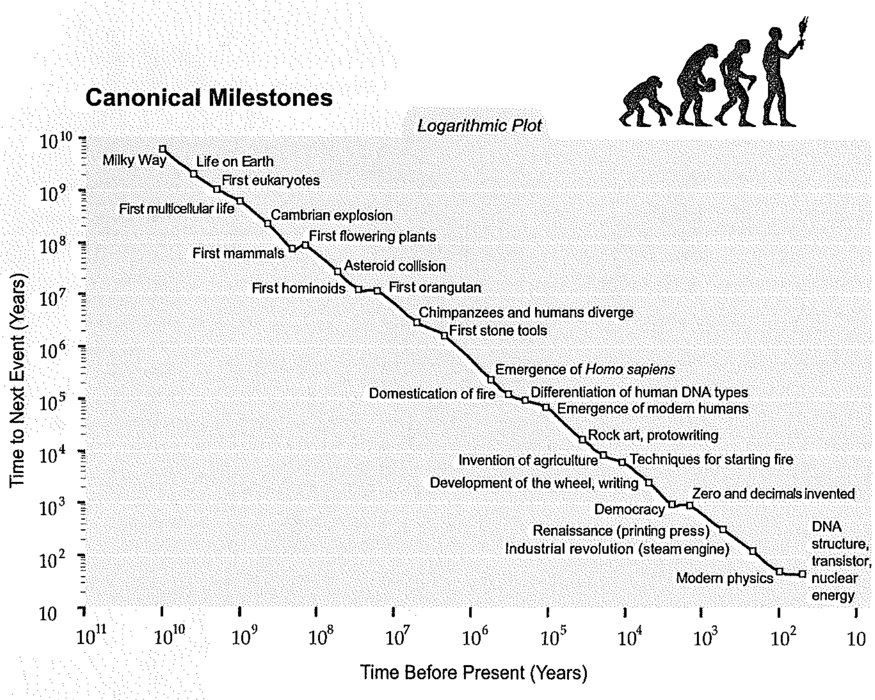

Canonical milestones based on clusters of events from thirteen lists.

The attributes that are growing exponentially in these charts are order and complexity, concepts we will explore in the next chapter. This acceleration matches our commonsense observations. A billion years ago, not much happened over the course of even one million years. But a quarter-million years ago epochal events such as the evolution of our species occurred in time frames of just one hundred thousand years. In technology, if we go back fifty thousand years, not much happened over a one-thousand-year period. But in the recent past, we see new paradigms, such as the World Wide Web, progress from inception to mass adoption (meaning that they are used by a quarter of the population in advanced countries) within only a decade.

Epoch Five: The Merger of Human Technology with Human Intelligence. Looking ahead several decades, the Singularity will begin with the fifth epoch. It will result from the merger of the vast knowledge embedded in our own brains with the vastly greater capacity, speed, and knowledge-sharing ability of our technology. The fifth epoch will enable our human-machine civilization to transcend the human brain's limitations of a mere hundred trillion extremely slow connections.[14]

CHAPTER TWO

A Theory of Technology Evolution

The Law of Accelerating Returns

The further backward you look, the further forward you can see.

—WINSTON CHURCHILL

Two billion years ago, our ancestors were microbes; a half-billion years ago, fish; a hundred million years ago, something like mice; ten million years ago, arboreal apes; and a million years ago, proto-humans puzzling out the taming of fire. Our evolutionary lineage is marked by mastery of change. In our time, the pace is quickening.

—CARL SAGAN

Our sole responsibility is to produce something smarter than we are; any problems beyond that are not *ours* to solve. . . . [T]here are no hard problems, only problems that are hard to a certain level of intelligence. Move the smallest bit upwards [in level of intelligence], and some problems will suddenly move from "impossible" to "obvious." Move a substantial degree upwards, and all of them will become obvious.

—ELIEZER S. YUDKOWSKY, *STARING INTO THE SINGULARITY*, 1996

"The future can't be predicted," is a common refrain. . . . But . . . when [this perspective] is wrong, it is profoundly wrong.

—JOHN SMART[1]

The ongoing acceleration of technology is the implication and inevitable result of what I call the law of accelerating returns, which describes the acceleration of the pace of and the exponential growth of the products of an evolutionary process. These products include, in particular, information-bearing technologies such as computation, and their acceleration extends substantially beyond the predictions made by what has become known as Moore's tial is exhausted. When this happens, a paradigm shift occurs, which enables exponential growth to continue.

The Life Cycle of a Paradigm. Each paradigm develops in three stages:

1. Slow growth (the early phase of exponential growth)
2. Rapid growth (the late, explosive phase of exponential growth), as seen in the S-curve figure below
3. A leveling off as the particular paradigm matures

The progression of these three stages looks like the letter S, stretched to the right. The S-curve illustration shows how an ongoing exponential trend can be composed of a cascade of S-curves. Each successive S-curve is faster (takes less time on the time, or *x*, axis) and higher (takes up more room on the performance, or *y*, axis).

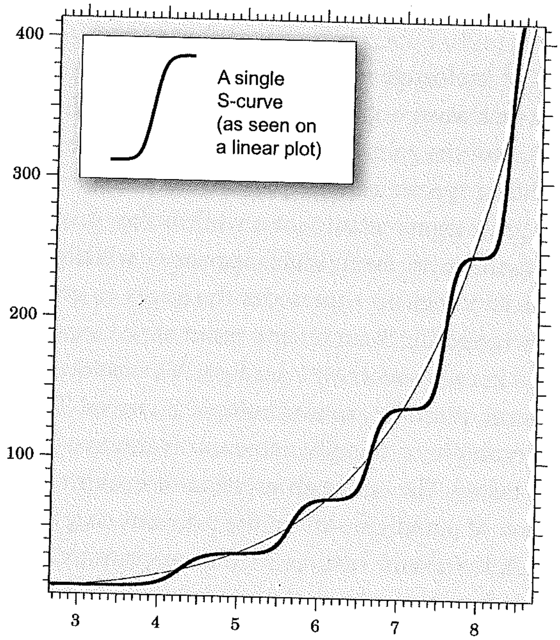

An ongoing exponential sequence made up of a cascade of S-curves (linear plot)

The S-Curve of a Technology as Expressed in Its Life Cycle

A machine is as distinctively and brilliantly and expressively human as a violin sonata or a theorem in Euclid.

—Gregory Vlastos

It is a far cry from the monkish calligrapher, working in his cell in silence, to the brisk "click, click" of the modern writing machine, which in a quarter of a century has revolutionized and reformed business.

—*Scientific American*, 1905

No communication technology has ever disappeared, but instead becomes increasingly less important as the technological horizon widens.

—Arthur C. Clarke

I always keep a stack of books on my desk that I leaf through when I run out of ideas, feel restless, or otherwise need a shot of inspiration. Picking up a fat volume that I recently acquired, I consider the bookmaker's craft: 470 finely printed pages organized into 16-page signatures, all of which are sewn together with white thread and glued onto a gray canvas cord. The hard linen-bound covers, stamped with gold letters, are connected to the signature block by delicately embossed end sheets. This is a technology that was perfected many decades ago. Books constitute such an integral element of our society—both reflecting and shaping its culture—that it is hard to imagine life without them. But the printed book, like any other technology, will not live forever.

The Life Cycle of a Technology

We can identify seven distinct stages in the life cycle of a technology.

1. During the precursor stage, the prerequisites of a technology exist, and dreamers may contemplate these elements coming together. We do not, however, regard dreaming to be the same as inventing, even if the dreams are written down. Leonardo da Vinci drew convincing pictures of airplanes and automobiles, but he is not considered to have invented either.
2. The next stage, one highly celebrated in our culture, is invention, a very brief stage, similar in some respects to the process of birth ▸

after an extended period of labor. Here the inventor blends curiosity, scientific skills, determination, and usually a measure of showmanship to combine methods in a new way and brings a new technology to life.

3. The next stage is development, during which the invention is protected and supported by doting guardians (who may include the original inventor). Often this stage is more crucial than invention and may involve additional creation that can have greater significance than the invention itself. Many tinkerers had constructed finely hand-tuned horseless carriages, but it was Henry Ford's innovation of mass production that enabled the automobile to take root and flourish.
4. The fourth stage is maturity. Although continuing to evolve, the technology now has a life of its own and has become an established part of the community. It may become so interwoven in the fabric of life that it appears to many observers that it will last forever. This creates an interesting drama when the next stage arrives, which I call the stage of the false pretenders.
5. Here an upstart threatens to eclipse the older technology. Its enthusiasts prematurely predict victory. While providing some distinct benefits, the newer technology is found on reflection to be lacking some key element of functionality or quality. When it indeed fails to dislodge the established order, the technology conservatives take this as evidence that the original approach will indeed live forever.
6. This is usually a short-lived victory for the aging technology. Shortly thereafter, another new technology typically does succeed in rendering the original technology to the stage of obsolescence. In this part of the life cycle, the technology lives out its senior years in gradual decline, its original purpose and functionality now subsumed by a more spry competitor.
7. In this stage, which may comprise 5 to 10 percent of a technology's life cycle, it finally yields to antiquity (as did the horse and buggy, the harpsichord, the vinyl record, and the manual typewriter).

In the mid-nineteenth century there were several precursors to the phonograph, including Léon Scott de Martinville's phonautograph, a device that recorded sound vibrations as a printed pattern. It was Thomas Edison, however, who brought all of the elements together and invented the first device that could both record and reproduce sound in 1877. Further refinements were necessary for the phonograph to become commercially ▶

viable. It became a fully mature technology in 1949 when Columbia introduced the 33-rpm long-playing record (LP) and RCA Victor introduced the 45-rpm disc. The false pretender was the cassette tape, introduced in the 1960s and popularized during the 1970s. Early enthusiasts predicted that its small size and ability to be rerecorded would make the relatively bulky and scratchable record obsolete.

Despite these obvious benefits, cassettes lack random access and are prone to their own forms of distortion and lack of fidelity. The compact disc (CD) delivered the mortal blow. With the CD providing both random access and a level of quality close to the limits of the human auditory system, the phonograph record quickly entered the stage of obsolescence. Although still produced, the technology that Edison gave birth to almost 130 years ago has now reached antiquity.

Consider the piano, an area of technology that I have been personally involved with replicating. In the early eighteenth century Bartolommeo Cristofori was seeking a way to provide a touch response to the then-popular harpsichord so that the volume of the notes would vary with the intensity of the touch of the performer. Called *gravicembalo col piano e forte* ("harpsichord with soft and loud"), his invention was not an immediate success. Further refinements, including Stein's Viennese action and Zumpe's English action, helped to establish the "piano" as the preeminent keyboard instrument. It reached maturity with the development of the complete cast-iron frame, patented in 1825 by Alpheus Babcock, and has seen only subtle refinements since then. The false pretender was the electric piano of the early 1980s. It offered substantially greater functionality. Compared to the single (piano) sound of the acoustic piano, the electronic variant offered dozens of instrument sounds, sequencers that allowed the user to play an entire orchestra at once, automated accompaniment, educational programs to teach keyboard skills, and many other features. The only feature it was missing was a good-quality piano sound.

This crucial flaw and the resulting failure of the first generation of electronic pianos led to the widespread conclusion that the piano would never be replaced by electronics. But the "victory" of the acoustic piano will not be permanent. With their far greater range of features and price-performance, digital pianos already exceed the sales of acoustic pianos in homes. Many observers feel that the quality of the "piano" sound on digital pianos now equals or exceeds that of the upright acoustic piano. With the exception of concert and luxury grand pianos (a small part of the market), the sale of acoustic pianos is in decline. ▸

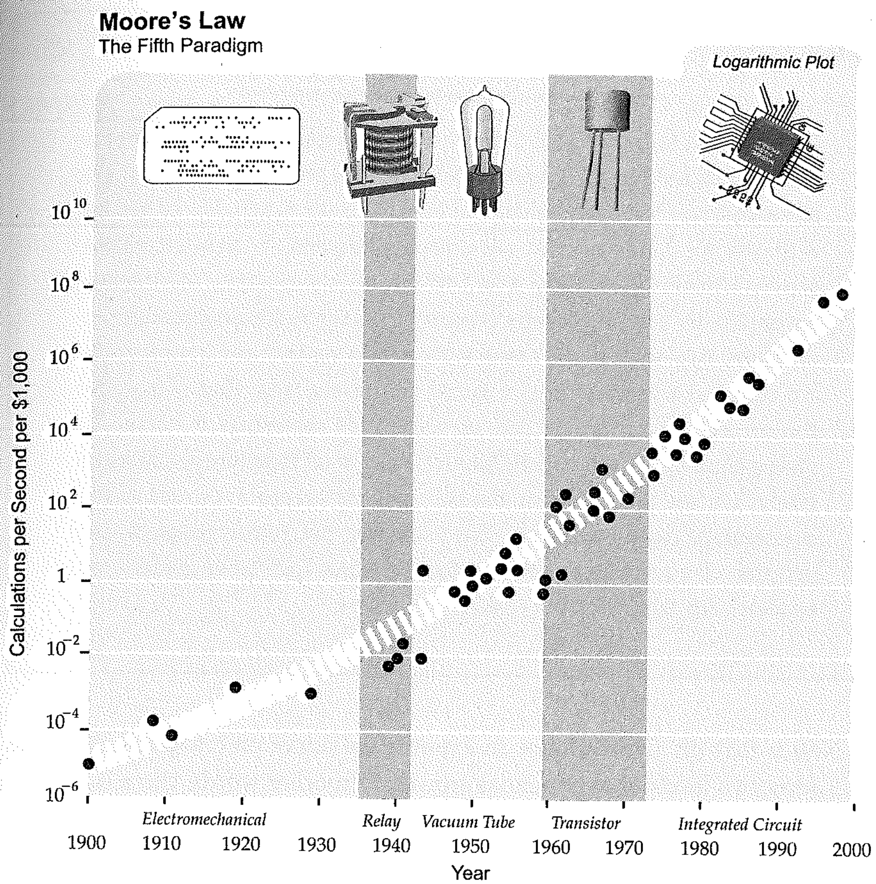

The five paradigms of exponential growth of computing: **Each time one paradigm has run out of steam, another has picked up the pace.**

As the figure demonstrates, there were actually four different paradigms—electromechanical, relays, vacuum tubes, and discrete transistors—that showed exponential growth in the price-performance of computing long before integrated circuits were even invented. And Moore's paradigm won't be the last. When Moore's Law reaches the end of its S-curve, now expected before 2020, the exponential growth will continue with three-dimensional molecular computing, which will constitute the sixth paradigm.

Evolution of Computer Power/Cost

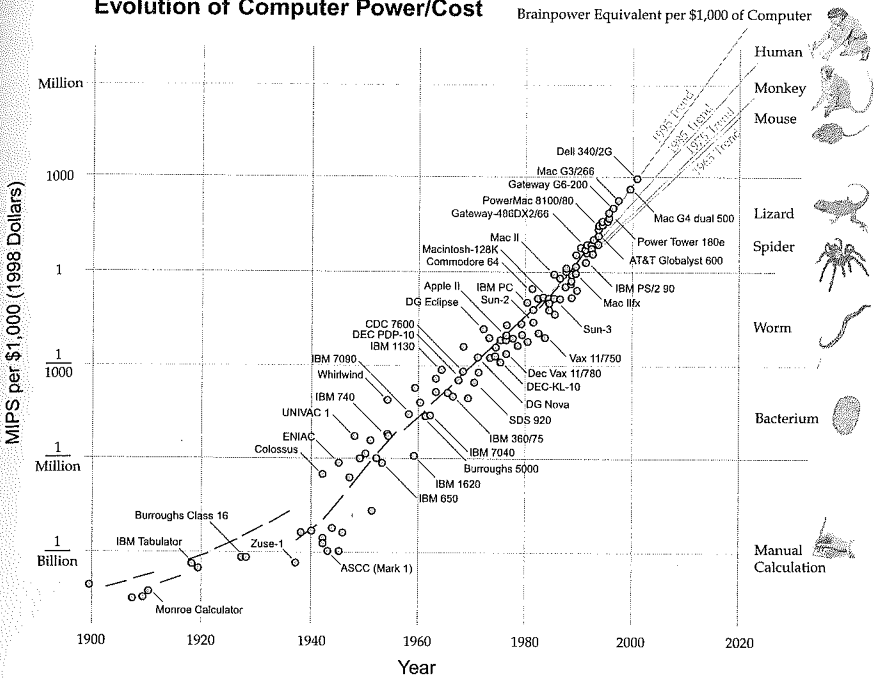

If we project these computational performance trends through this next century, we can see in the figure below that supercomputers will match human brain capability by the end of this decade and personal computing will achieve it by around 2020—or possibly sooner, depending on how conservative an estimate of human brain capacity we use. (We'll discuss estimates of human brain computational speed in the next chapter.)[39]

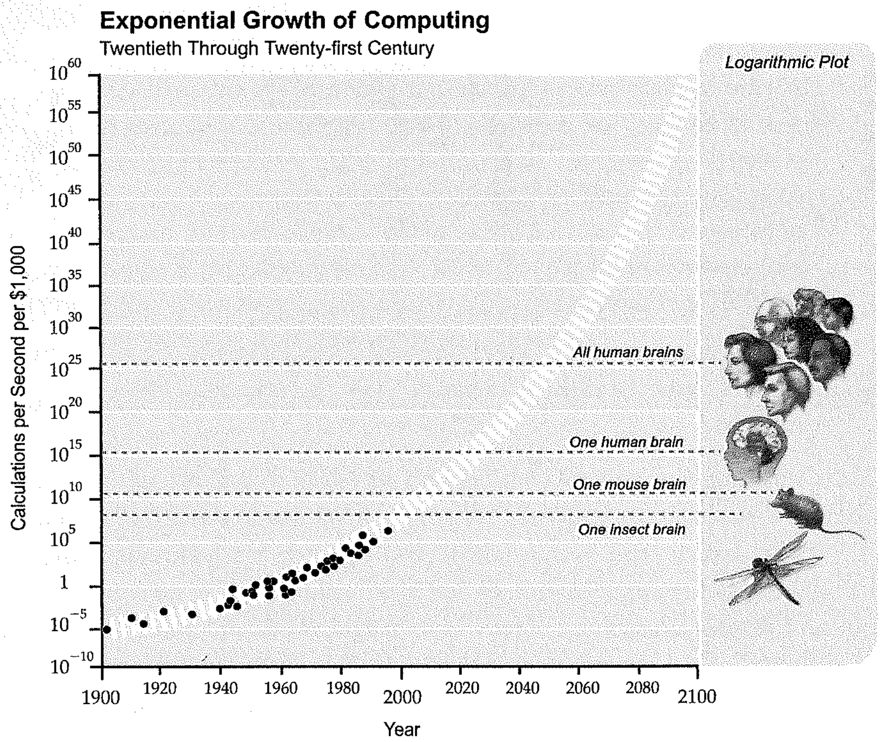

The exponential growth of computing is a marvelous quantitative example of the exponentially growing returns from an evolutionary process. We can express the exponential growth of computing in terms of its accelerating pace: it took ninety years to achieve the first MIPS per thousand dollars; now we add one MIPS per thousand dollars every five hours.[40]

To accommodate this exponential growth, the data transmission speed of the Internet backbone (as represented by the fastest announced backbone communication channels actually used for the Internet) has itself grown exponentially. Note that in the figure "Internet Backbone Bandwidth" below, we can actually see the progression of S-curves: the acceleration fostered by a new paradigm, followed by a leveling off as the paradigm runs out of steam, followed by renewed acceleration through paradigm shift.[53]

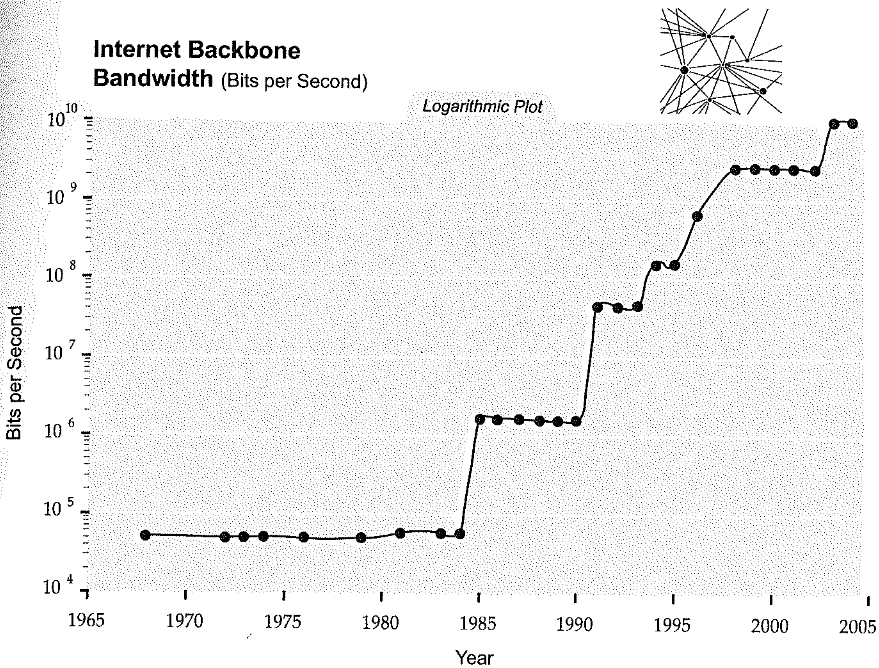

Exponential Software Price-Performance Improvement[89]

Example: Automatic Speech-Recognition Software

	1985	1995	2000
Price	$5,000	$500	$50
Vocabulary Size (number of words)	1,000	10,000	100,000
Continuous Speech?	No	No	Yes
User Training Required (minutes)	180	60	5
Accuracy	Poor	Fair	Good

The impact of distributed and intelligent communications has been felt perhaps most intensely in the world of business. Despite dramatic mood swings on Wall Street, the extraordinary values ascribed to so-called e-companies during the 1990s boom era reflected a valid perception: the business models that have sustained businesses for decades are in the early phases of a radical transformation. New models based on direct personalized communication with the customer will transform every industry, resulting in massive disintermediation of the middle layers that have traditionally separated the customer from the ultimate source of products and services. There is, however, a pace to all revolutions, and the investments and stock market valuations in this area expanded way beyond the early phases of this economic S-curve.

The boom-and-bust cycle in these information technologies was strictly a capital-markets (stock-value) phenomenon. Neither boom nor bust is apparent in the actual business-to-consumer (B2C) and business-to-business (B2B) data (see the figure on the next page). Actual B2C revenues grew smoothly from $1.8 billion in 1997 to $70 billion in 2002. B2B had similarly smooth growth from $56 billion in 1999 to $482 billion in 2002.[90] In 2004 it is approaching $1 trillion. We certainly do not see any evidence of business cycles in the actual price-performance of the underlying technologies, as I discussed extensively above.

Another implication of the law of accelerating returns is exponential growth in education and learning. Over the past 120 years, we have increased our investment in K–12 education (per student and in constant dollars) by a factor of ten. There has been a hundredfold increase in the number of college students. Automation started by amplifying the power of our muscles and in recent times has been amplifying the power of our minds. So for the past two centuries, automation has been eliminating jobs at the bottom of the skill ladder while creating new (and better-paying) jobs at the top of the skill ladder. The ladder has been moving up, and thus we have been exponentially increasing investments in education at all levels (see the figure below).[93]

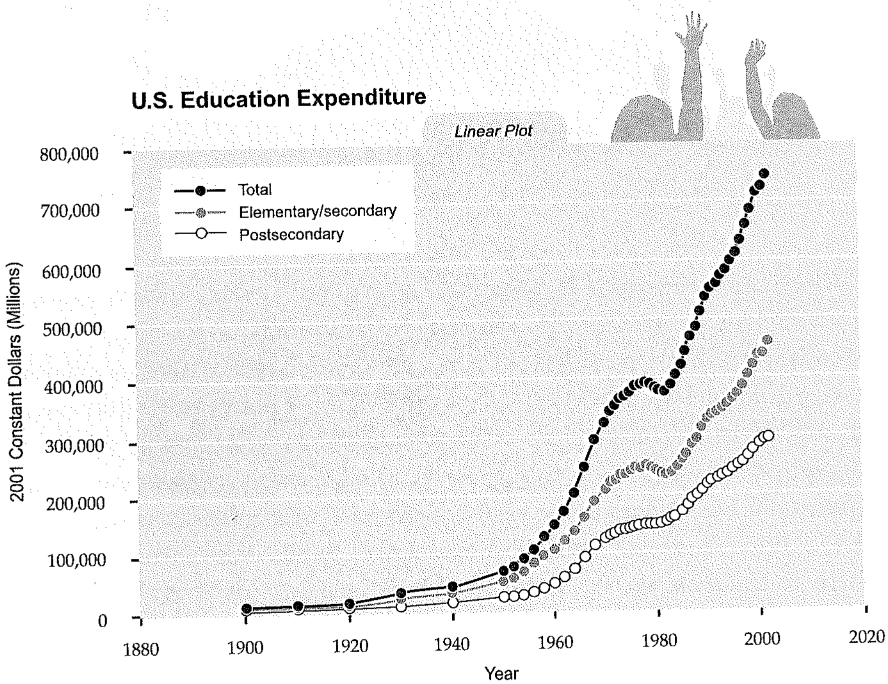

Oh, and about that "offer" at the beginning of this précis, consider that present stock values are based on future expectations. Given that the (literally) shortsighted linear intuitive view represents the ubiquitous outlook, the common wisdom in economic expectations is dramatically understated. Since stock prices reflect the consensus of a buyer-seller market, the prices reflect the underlying linear assumption that most people share regarding future economic growth. But the law of accelerating returns clearly implies that the growth rate will continue to grow exponentially, because the rate of progress will continue to accelerate.

CHAPTER THREE

Achieving the Computational Capacity of the Human Brain

As I discuss in *Engines of Creation*, if you can build genuine AI, there are reasons to believe that you can build things like neurons that are a million times faster. That leads to the conclusion that you can make systems that think a million times faster than a person. With AI, these systems could do engineering design. Combining this with the capability of a system to build something that is better than it, you have the possibility for a very abrupt transition. This situation may be more difficult to deal with even than nanotechnology, but it is much more difficult to think about it constructively at this point. Thus, it hasn't been the focus of things that I discuss, although I periodically point to it and say: "That's important too."

—ERIC DREXLER, 1989

The Sixth Paradigm of Computing Technology: Three-Dimensional Molecular Computing and Emerging Computational Technologies

In the April 19, 1965, issue of *Electronics*, Gordon Moore wrote, "The future of integrated electronics is the future of electronics itself. The advantages of integration will bring about a proliferation of electronics, pushing this science into many new areas."[1] With those modest words, Moore ushered in a revolution that is still gaining momentum. To give his readers some idea of how profound this new science would be, Moore predicted that "by 1975, economics may dictate squeezing as many as 65,000 components on a single silicon chip." Imagine that.

Moore's article described the repeated annual doubling of the number of transistors (used for computational elements, or gates) that could be fitted onto an integrated circuit. His 1965 "Moore's Law" prediction was criticized at the time because his logarithmic chart of the number of components on a chip

CHAPTER FOUR

Achieving the Software of Human Intelligence

How to Reverse Engineer the Human Brain

There are good reasons to believe that we are at a turning point, and that it will be possible within the next two decades to formulate a meaningful understanding of brain function. This optimistic view is based on several measurable trends, and a simple observation which has been proven repeatedly in the history of science: *Scientific advances are enabled by a technology advance that allows us to see what we have not been able to see before.* At about the turn of the twenty-first century, we passed a detectable turning point in both neuroscience knowledge and computing power. For the first time in history, we collectively know enough about our own brains, and have developed such advanced computing technology, that we can now seriously undertake the construction of a verifiable, real-time, high-resolution model of significant parts of our intelligence.

—LLOYD WATTS, NEUROSCIENTIST[1]

Now, for the first time, we are observing the brain at work in a global manner with such clarity that we should be able to discover the overall programs behind its magnificent powers.

—J. G. TAYLOR, B. HORWITZ, K. J. FRISTON, NEUROSCIENTISTS[2]

The brain is good: it is an existence proof that a certain arrangement of matter can produce mind, perform intelligent reasoning, pattern recognition, learning and a lot of other important tasks of engineering interest. Hence we can learn to build new systems by borrowing ideas from the brain. . . . The brain is bad: it is an evolved, messy system where a lot of interactions happen because of evolutionary contingencies. . . . On the other hand, it must also be robust (since we can survive with it) and be able to stand fairly major variations and environmental insults, so the truly valuable insight from the brain might be how to create resilient complex systems that self-organize well. . . . The interactions within a neuron are complex, but on the next level neurons seem to be somewhat simple objects that can

CHAPTER FIVE

GNR

Three Overlapping Revolutions

> There are few things of which the present generation is more justly proud than the wonderful improvements which are daily taking place in all sorts of mechanical appliances. . . . But what would happen if technology continued to evolve so much more rapidly than the animal and vegetable kingdoms? Would it displace us in the supremacy of earth? Just as the vegetable kingdom was slowly developed from the mineral, and as in like manner the animal supervened upon the vegetable, so now in these last few ages an entirely new kingdom has sprung up, of which we as yet have only seen what will one day be considered the antediluvian prototypes of the race. . . . We are daily giving [machines] greater power and supplying by all sorts of ingenious contrivances that self-regulating, self-acting power which will be to them what intellect has been to the human race.
>
> —SAMUEL BUTLER, 1863 (FOUR YEARS AFTER PUBLICATION OF DARWIN'S *THE ORIGIN OF SPECIES*)

> Who will be man's successor? To which the answer is: We are ourselves creating our own successors. Man will become to the machine what the horse and the dog are to man; the conclusion being that machines are, or are becoming, animate.
>
> —SAMUEL BUTLER, 1863 LETTER, "DARWIN AMONG THE MACHINES"[1]

The first half of the twenty-first century will be characterized by three overlapping revolutions—in Genetics, Nanotechnology, and Robotics. These will usher in what I referred to earlier as Epoch Five, the beginning of the Singularity. We are in the early stages of the "G" revolution today. By understanding the information processes underlying life, we are starting to learn to reprogram our biology to achieve the virtual elimination of disease, dramatic expansion of human potential, and radical life extension. Hans Moravec points out, however, that no matter how successfully we fine-tune our DNA-based biology, humans will remain "second-class robots," meaning that our judgment. Another relates to our own ineptitude. We do arithmetic or keep records so painstakingly and externally that the small mechanical steps in a long calculation are obvious, while the big picture often escapes us. Like Deep Blue's builders, we see the process too much from the inside to appreciate the subtlety that it may have on the outside. But there is a nonobviousness in snowstorms or tornadoes that emerge from the repetitive arithmetic of weather simulations, or in rippling tyrannosaur skin from movie animation calculations. We rarely call it intelligence, but "artificial reality" may be an even more profound concept than artificial intelligence.

The mental steps underlying good human chess playing and theorem proving are complex and hidden, putting a mechanical interpretation out of reach. Those who can follow the play naturally describe it instead in mentalistic language, using terms like strategy, understanding and creativity. When a machine manages to be simultaneously meaningful and surprising in the same rich way, it too compels a mentalistic interpretation. Of course, somewhere behind the scenes, there are programmers who, in principle, have a mechanical interpretation. But even for them, that interpretation loses its grip as the working program fills its memory with details too voluminous for them to grasp.

As the rising flood reaches more populated heights, machines will begin to do well in areas a greater number can appreciate. The visceral sense of a thinking presence in machinery will become increasingly widespread. When the highest peaks are covered, there will be machines that can interact as intelligently as any human on any subject. The presence of minds in machines will then become self-evident.

—HANS MORAVEC[214]

Because of the exponential nature of progress in information-based technologies, performance often shifts quickly from pathetic to daunting. In many diverse realms, as the examples in the previous section make clear, the performance of narrow AI is already impressive. The range of intelligent tasks in which machines can now compete with human intelligence is continually expanding. In a cartoon I designed for *The Age of Spiritual Machines*, a defensive "human race" is seen writing out signs that state what only people (and not machines) can do.[215] Littered on the floor are the signs the human race has already discarded because machines can now perform these functions: diagnose an electrocardiogram, compose in the style of Bach, recognize faces, guide a missile, play Ping-Pong, play master chess, pick stocks, improvise jazz, prove important theorems, and understand continuous speech. Back in 1999 these tasks were no longer solely the province of human intelligence; machines could do them all.

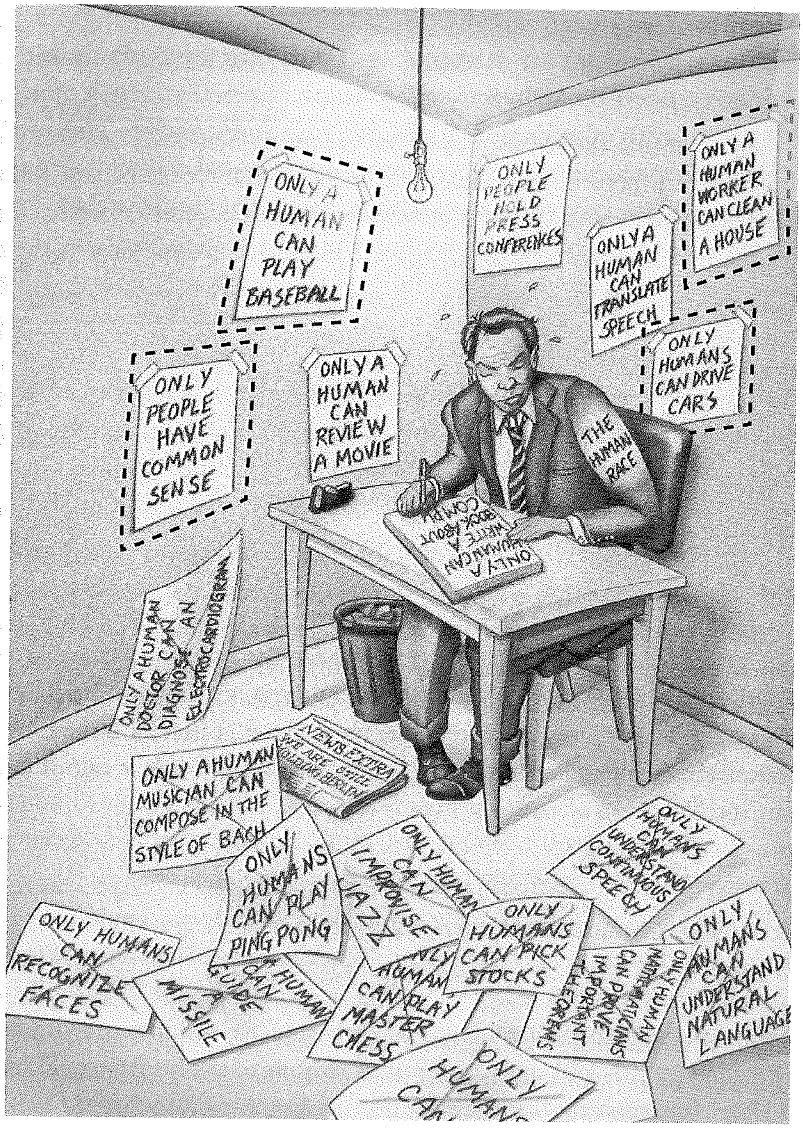

On the wall behind the man symbolizing the human race are signs he has written out describing the tasks that were still the sole province of humans: have common sense, review a movie, hold press conferences, translate speech, clean a house, and drive cars. If we were to redesign this cartoon in a few years, some of these signs would also be likely to end up on the floor. When CYC reaches one

CHAPTER SIX

The Impact . . .

The future enters into us in order to transform itself in us long before it happens.

—Rainer Maria Rilke

One of the biggest flaws in the common conception of the future is that the future is something that happens to us, not something we create.

—Michael Anissimov

"Playing God" is actually the highest expression of human nature. The urges to improve ourselves, to master our environment, and to set our children on the best path possible have been the fundamental driving forces of all of human history. Without these urges to "play God," the world as we know it wouldn't exist today. A few million humans would live in savannahs and forests, eking out a hunter-gatherer existence, without writing or history or mathematics or an appreciation of the intricacies of their own universe and their own inner workings.

—Ramez Naam

A **Panoply of Impacts.** What will be the nature of human experience once nonbiological 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.

As the Singularity approaches we will have to reconsider our ideas about the nature of human life and redesign our human institutions. We will explore a few of these ideas and institutions in this chapter.

For example, the intertwined revolutions of G, N, and R will transform our frail version 1.0 human bodies into their far more durable and capable version

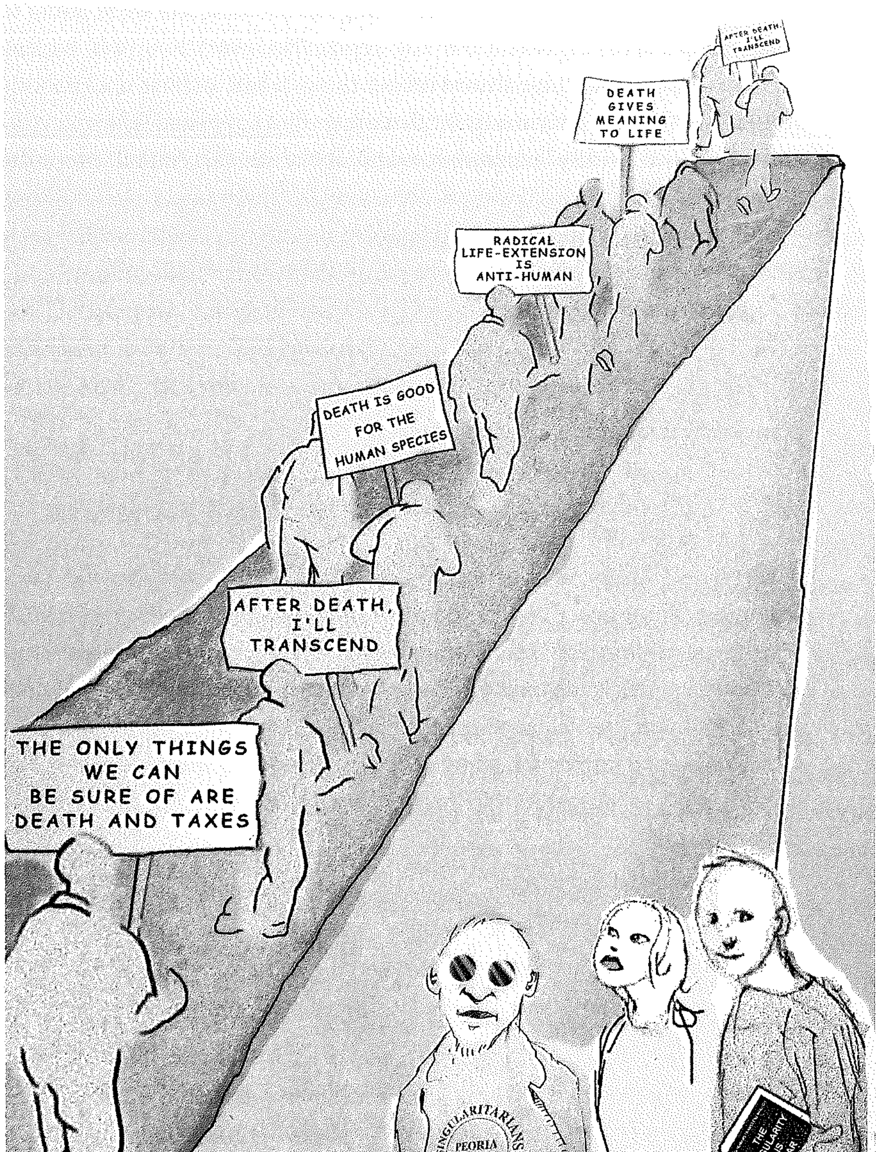
AFTER DEATH, I'LL TRANSCEND
DEATH GIVES MEANING TO LIFE
RADICAL LIFE-EXTENSION IS ANTI-HUMAN
DEATH IS GOOD FOR THE HUMAN SPECIES
AFTER DEATH, I'LL TRANSCEND
THE ONLY THINGS WE CAN BE SURE OF ARE DEATH AND TAXES
INTERNATIONAL SINGULARITARIANS
PEORIA CHAPTER
THE SINGULARITY IS NEAR

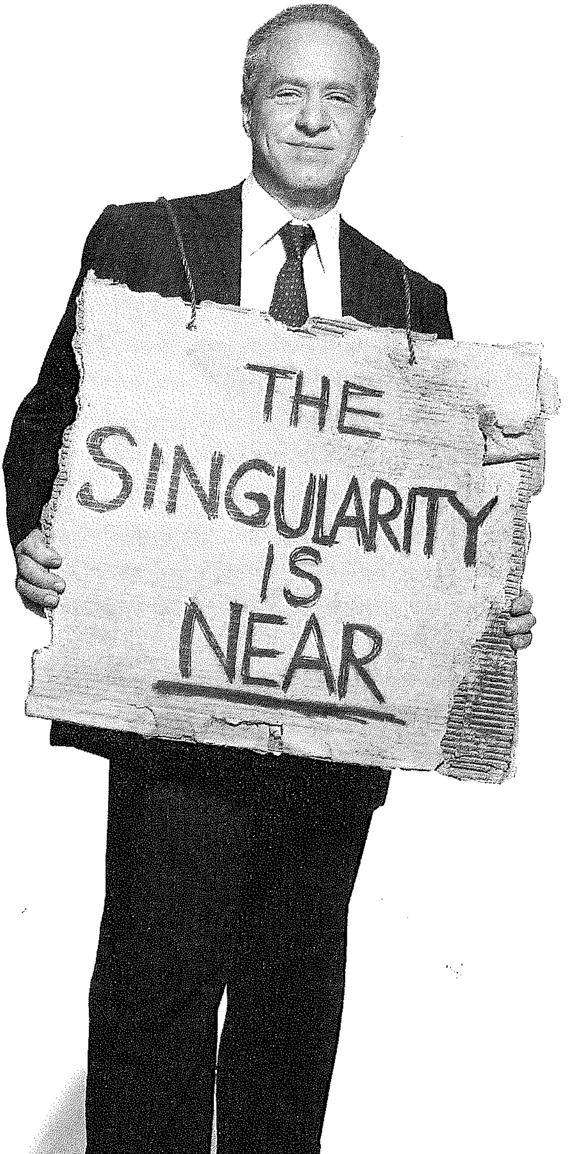
THE
SINGULARITY
IS
NEAR

CHAPTER SEVEN

Ich bin ein Singularitarian

The most common of all follies is to believe passionately in the palpably not true.

—H. L. Mencken

Philosophies of life rooted in centuries-old traditions contain much wisdom concerning personal, organizational, and social living. Many of us also find shortcomings in those traditions. How could they not reach some mistaken conclusions when they arose in pre-scientific times? At the same time, ancient philosophies of life have little or nothing to say about fundamental issues confronting us as advanced technologies begin to enable us to change our identity as individuals and as humans and as economic, cultural, and political forces change global relationships.

—Max More, "Principles of Extropy"

The world does not need another totalistic dogma.

—Max More, "Principles of Extropy"

Yes, we have a soul. But it's made of lots of tiny robots.

—Giulio Giorelli

Substrate is morally irrelevant, assuming it doesn't affect functionality or consciousness. It doesn't matter, from a moral point of view, whether somebody runs on silicon or biological neurons (just as it doesn't matter whether you have dark or pale skin). On the same grounds, that we reject racism and speciesism, we should also reject carbon-chauvinism, or *bioism*.

—Nick Bostrom, "Ethics for Intelligent Machines: A Proposal, 2001"

Philosophers have long noted that their children were born into a more complex world than that of their ancestors. This early and perhaps even

I discuss below (see the section "A Program for GNR Defense," p. 422) steps we can take to address these grave risks, but we cannot have complete assurance in any strategy that we devise today. These risks are what Nick Bostrom calls "existential risks," which he defines as the dangers in the upper-right quadrant of the following table:[18]

Bostrom's Categorization of Risks

		Intensity of Risk	
		Moderate	**Profound**
	Global	Ozone Thinning	***Existential Risks***
Scope	**Local**	Recession	Genocide
	Personal	Stolen Car	Death
		Endurable	**Terminal**

Biological life on Earth encountered a human-made existential risk for the first time in the middle of the twentieth century with the advent of the hydrogen bomb and the subsequent cold-war buildup of thermonuclear forces. President Kennedy reportedly estimated that the likelihood of an all-out nuclear war during the Cuban missile crisis was between 33 and 50 percent.[19] The legendary information theorist John von Neumann, who became the chairman of the Air Force Strategic Missiles Evaluation Committee and a government adviser on nuclear strategies, estimated the likelihood of nuclear Armageddon (prior to the Cuban missile crisis) at close to 100 percent.20 Given the perspective of the 1960s what informed observer of those times would have predicted that the world would have gone through the next forty years without another nontest nuclear explosion?

Despite the apparent chaos of international affairs we can be grateful for the successful avoidance thus far of the employment of nuclear weapons in war. But we clearly cannot rest easily, since enough hydrogen bombs still exist to destroy all human life many times over.[21] Although attracting relatively little public discussion, the massive opposing ICBM arsenals of the United States and Russia remain in place, despite the apparent thawing of relations.

Nuclear proliferation and the widespread availability of nuclear materials and know-how is another grave concern, although not an existential one for our civilization. (That is, only an all-out thermonuclear war involving the ICBM arsenals poses a risk to survival of all humans.) Nuclear proliferation

CHAPTER NINE

Response to Critics

The human mind likes a strange idea as little as the body likes a strange protein and resists it with a similar energy.

—W. I. Beveridge

If a . . . scientist says that something is possible he is almost certainly right, but if he says that it is impossible he is very probably wrong.

—Arthur C. Clarke

A Panoply of Criticisms

In *The Age of Spiritual Machines*, I began to examine some of the accelerating trends that I have sought to explore in greater depth in this book. *ASM* inspired a broad variety of reactions, including extensive discussions of the profound, imminent changes it considered (for example, the promise-versus-peril debate prompted by Bill Joy's *Wired* story, "Why the Future Doesn't Need Us," as I reviewed in the previous chapter). The response also included attempts to argue on many levels why such transformative changes would not, could not, or should not happen. Here is a summary of the critiques I will be responding to in this chapter:

- The "criticism from Malthus": *It's a mistake to extrapolate exponential trends indefinitely, since they inevitably run out of resources to maintain the exponential growth. Moreover, we won't have enough energy to power the extraordinarily dense computational platforms forecast, and even if we did they would be as hot as the sun.* Exponential trends do reach an asymptote, but the matter and energy resources needed for computation and communication are so small per compute and per bit that these trends can

Epilogue

I do not know what I may appear to the world, but to myself I seem to have been only like a boy playing on the seashore, and diverting myself in now and then finding a smoother pebble or a prettier shell than ordinary, whilst the great ocean of truth lay undiscovered before me.

—Isaac Newton[1]

The meaning of life is creative love. Not love as an inner feeling, as a private sentimental emotion, but love as a dynamic power moving out into the world and doing something original.

—Tom Morris, *If Aristotle Ran General Motors*

No exponential is forever . . . but we can delay "forever."

—Gordon E. Moore, 2004

How Singular? How singular is the Singularity? Will it happen in an instant? Let's consider again the derivation of the word. In mathematics a singularity is a value that is beyond any limit—in essence, infinity. (Formally the value of a function that contains such a singularity is said to be undefined at the singularity point, but we can show that the value of the function at nearby points exceeds any specific finite value).[2]

The Singularity, as we have discussed it in this book, does not achieve infinite levels of computation, memory, or any other measurable attribute. But it certainly achieves vast levels of all of these qualities, including intelligence. With the reverse engineering of the human brain we will be able to apply the parallel, self-organizing, chaotic algorithms of human intelligence to enormously powerful computational substrates. This intelligence will then be in a position to improve its own design, both hardware and software, in a rapidly accelerating iterative process.

But there still appears to be a limit. The capacity of the universe to support intelligence appears to be only about 10^{90} calculations per second, as I discussed in chapter 6. There are theories such as the holographic universe that suggest the possibility of higher numbers (such as 10^{120}), but these levels are all decidedly finite.

Of course, the capabilities of such an intelligence may appear infinite for all practical purposes to our current level of intelligence. A universe saturated with intelligence at 10^{90} cps would be one trillion trillion trillion trillion trillion times more powerful than all biological human brains on Earth today.[3] Even a one-kilogram "cold" computer has a peak potential of 10^{42} cps, as I reviewed in chapter 3, which is ten thousand trillion (10^{16}) times more powerful than all biological human brains.[4]

Given the power of exponential notation, we can easily conjure up bigger numbers, even if we lack the imagination to contemplate all of their implications. We can imagine the possibility of our future intelligence spreading into other universes. Such a scenario is conceivable given our current understanding of cosmology, although speculative. This could potentially allow our future intelligence to go beyond any limits. If we gained the ability to create and colonize other universes (and if there is a way to do this, the vast intelligence of our future civilization is likely to be able to harness it), our intelligence would ultimately be capable of exceeding any specific finite level. That's exactly what we can say for singularities in mathematical functions.

How does our use of "singularity" in human history compare to its use in physics? The word was borrowed from mathematics by physics, which has always shown a penchant for anthropomorphic terms (such as "charm" and "strange" for names of quarks). In physics "singularity" theoretically refers to a point of zero size with infinite density of mass and therefore infinite gravity. But because of quantum uncertainty there is no actual point of infinite density, and indeed quantum mechanics disallows infinite values.

Just like the Singularity as I have discussed it in this book, a singularity in physics denotes unimaginably large values. And the area of interest in physics is not actually zero in size but rather is an event horizon around the theoretical singularity point inside a black hole (which is not even black). Inside the event horizon particles and energy, such as light, cannot escape because gravity is too strong. Thus from outside the event horizon, we cannot see easily inside the event horizon with certainty.

However, there does appear to be a way to see inside a black hole, because black holes give off a shower of particles. Particle-antiparticle pairs are created near the event horizon (as happens everywhere in space), and for some of these pairs, one of the pair is pulled into the black hole while the other manages to escape. These escaping particles form a glow called Hawking radiation, named after its discoverer, Stephen Hawking. The current thinking is that this radiation does reflect (in a coded fashion, and as a result of a form of quantum entanglement with the particles inside) what is happening inside the black hole. Hawking initially resisted this explanation but now appears to agree.

So, we find our use of the term "Singularity" in this book to be no less appropriate than the deployment 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 brains each limited to 10^{16} to 10^{19} cps, imagine what our future civilization in 2099 with its 10^{60} cps will 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.

Human Centrality. 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."[5]

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, 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 will continue until the entire universe is at our fingertips.