

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

expanding at an exponential pace. Exponential growth is deceptive. It starts out almost imperceptibly and then explodes with unexpected fury—unexpected, that is, if one does not take care to follow its trajectory. (See the “Linear vs. Exponential Growth” graph on p. 10.)

Consider this parable: a lake owner wants to stay at home to tend to the lake’s fish and make certain that the lake itself will not become covered with lily pads, which are said to double their number every few days. Month after month, he patiently waits, yet only tiny patches of lily pads can be discerned, and they don’t seem to be expanding in any noticeable way. With the lily pads covering less than 1 percent of the lake, the owner figures that it’s safe to take a vacation and leaves with his family. When he returns a few weeks later, he’s shocked to discover that the entire lake has become covered with the pads, and his fish have perished. By doubling their number every few days, the last seven doublings were sufficient to extend the pads’ coverage to the entire lake. (Seven doublings extended their reach 128-fold.) This is the nature of exponential growth.

Consider Gary Kasparov, who scorned the pathetic state of computer chess in 1992. Yet the relentless doubling of computer power every year enabled a computer to defeat him only five years later.² The list of ways computers can now exceed human capabilities is rapidly growing. Moreover, the once narrow applications of computer intelligence are gradually broadening in one type of activity after another. For example, computers are diagnosing electrocardiograms and medical images, flying and landing airplanes, controlling the tactical decisions of automated weapons, making credit and financial decisions, and being given responsibility for many other tasks that used to require human intelligence. The performance of these systems is increasingly based on integrating multiple types of artificial intelligence (AI). But as long as there is an AI shortcoming in any such area of endeavor, skeptics will point to that area as an inherent bastion of permanent human superiority over the capabilities of our own creations.

This book will argue, however, that within several decades information-based technologies will encompass all human knowledge and proficiency, ultimately including the pattern-recognition powers, problem-solving skills, and emotional and moral intelligence of the human brain itself.

Although impressive in many respects, the brain suffers from severe limitations. We use its massive parallelism (one hundred trillion interneuronal connections operating simultaneously) to quickly recognize subtle patterns. But our thinking is extremely slow: the basic neural transactions are several million times slower than contemporary electronic circuits. That makes our physiolog-

ical bandwidth for processing new information extremely limited compared to the exponential growth of the overall human knowledge base.

Our version 1.0 biological bodies are likewise frail and subject to a myriad of failure modes, not to mention the cumbersome maintenance rituals they require. While human intelligence is sometimes capable of soaring in its creativity and expressiveness, much human thought is derivative, petty, and circumscribed.

The Singularity will allow us to transcend these limitations of our biological bodies and brains. We will gain power over our fates. Our mortality will be in our own hands. We will be able to live as long as we want (a subtly different statement from saying we will live forever). We will fully understand human thinking and will vastly extend and expand its reach. By the end of this century, the nonbiological portion of our intelligence will be trillions of trillions of times more powerful than unaided human intelligence.

We are now in the early stages of this transition. The acceleration of paradigm shift (the rate at which we change fundamental technical approaches) as well as the exponential growth of the capacity of information technology are both beginning to reach the “knee of the curve,” which is the stage at which an exponential trend becomes noticeable. Shortly after this stage, the trend quickly becomes explosive. Before the middle of this century, the growth rates of our technology—which will be indistinguishable from ourselves—will be so steep as to appear essentially vertical. From a strictly mathematical perspective, the growth rates will still be finite but so extreme that the changes they bring about will appear to rupture the fabric of human history. That, at least, will be the perspective of unenhanced biological humanity.

The Singularity will represent the culmination of the merger of our biological thinking and existence with our technology, resulting in a world that is still human but that transcends our biological roots. There will be no distinction, post-Singularity, between human and machine or between physical and virtual reality. If you wonder what will remain unequivocally human in such a world, it's simply this quality: ours is the species that inherently seeks to extend its physical and mental reach beyond current limitations.

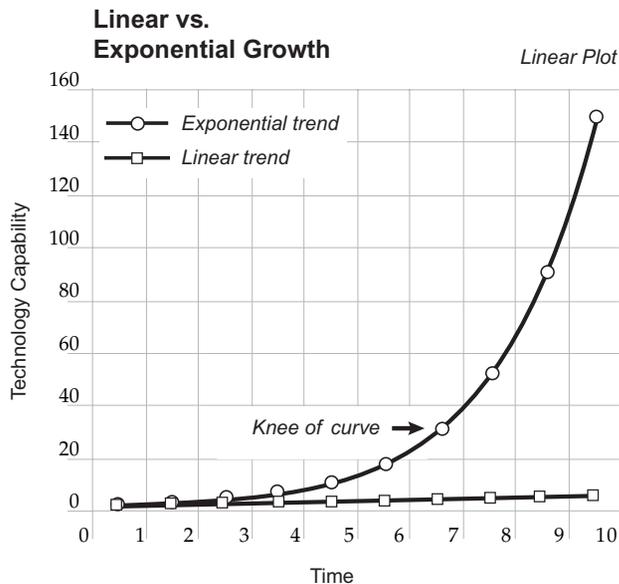
Many commentators on these changes focus on what they perceive as a loss of some vital aspect of our humanity that will result from this transition. This perspective stems, however, from a misunderstanding of what our technology will become. All the machines we have met to date lack the essential subtlety of human biological qualities. Although the Singularity has many faces, its most important implication is this: our technology will match and then vastly exceed the refinement and suppleness of what we regard as the best of human traits.

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

like their past. Exponential trends did exist one thousand years ago, but they were at that very early stage in which they were so flat and so slow that they looked like no trend at all. As a result, observers' expectation of an unchanged future was fulfilled. Today, we anticipate continuous technological progress and the social repercussions that follow. 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 itself is accelerating.

Most long-range forecasts of what is technically feasible in future time periods dramatically underestimate the power of future developments because they are based on what I call the "intuitive linear" view of history rather than the "historical exponential" view. My models show that we are doubling the paradigm-shift rate every decade, as I will discuss in the next chapter. Thus the twentieth century was gradually speeding up to today's rate of progress; its achievements, therefore, were equivalent to about twenty years of progress at the rate in 2000. We'll make another twenty years of progress in just fourteen years (by 2014), and then do the same again in only seven years. To express this another way, we won't experience one hundred years of technological advance in the twenty-first century; we will witness on the order of twenty thousand years of progress (again, when measured by *today's* rate of progress), or about one thousand times greater than what was achieved in the twentieth century.⁴

Misperceptions about the shape of the future come up frequently and in a variety of contexts. As one example of many, in a recent debate in which I took part concerning the feasibility of molecular manufacturing, a Nobel Prize-winning panelist dismissed safety concerns regarding nanotechnology, proclaiming that "we're not going to see self-replicating nanoengineered entities [devices constructed molecular fragment by fragment] for a hundred years." I pointed out that one hundred years was a reasonable estimate and actually matched my own appraisal of the amount of technical progress required to achieve this particular milestone when measured *at today's rate of progress* (five times the average rate of change we saw in the twentieth century). But because we're doubling the rate of progress every decade, we'll see the equivalent of a century of progress—*at today's rate*—in only twenty-five calendar years.

Similarly at *Time* magazine's Future of Life conference, held in 2003 to celebrate the fiftieth anniversary of the discovery of the structure of DNA, all of the invited speakers were asked what they thought the next fifty years would be like.⁵ Virtually every presenter looked at the progress of the last fifty years and used it as a model for the next fifty years. For example, James Watson, the codiscoverer of DNA, said that in fifty years we will have drugs that will allow us to eat as much as we want without gaining weight.

I replied, "Fifty years?" We have accomplished this already in mice by block-

ing the fat insulin receptor gene that controls the storage of fat in the fat cells. Drugs for human use (using RNA interference and other techniques we will discuss in chapter 5) are in development now and will be in FDA tests in several years. These will be available in five to ten years, not fifty. Other projections were equally shortsighted, reflecting contemporary research priorities rather than the profound changes that the next half century will bring. Of all the thinkers at this conference, it was primarily Bill Joy and I who took account of the exponential nature of the future, although Joy and I disagree on the import of these changes, as I will discuss in chapter 8.

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

But a serious assessment of the history of technology reveals that technological change is exponential. Exponential growth is a feature of any evolutionary process, of which technology is a primary example. You can examine the data in different ways, on different timescales, and for a wide variety of technologies, ranging from electronic to biological, as well as for their implications, ranging from the amount of human knowledge to the size of the economy. The acceleration of progress and growth applies to each of them. Indeed, we often find not just simple exponential growth, but "double" exponential growth, meaning that the rate of exponential growth (that is, the exponent) is itself growing exponentially (for example, see the discussion on the price-performance of computing in the next chapter).

Many scientists and engineers have what I call "scientist's pessimism." Often, they are so immersed in the difficulties and intricate details of a contemporary challenge that they fail to appreciate the ultimate long-term implications of their own work, and the larger field of work in which they operate. They likewise fail to account for the far more powerful tools they will have available with each new generation of technology.

Scientists are trained to be skeptical, to speak cautiously of current research goals, and to rarely speculate beyond the current generation of scientific pur-

suit. This may have been a satisfactory approach when a generation of science and technology lasted longer than a human generation, but it does not serve society's interests now that a generation of scientific and technological progress comprises only a few years.

Consider the biochemists who, in 1990, were skeptical of the goal of transcribing the entire human genome in a mere fifteen years. These scientists had just spent an entire year transcribing a mere one ten-thousandth of the genome. So, even with reasonable anticipated advances, it seemed natural to them that it would take a century, if not longer, before the entire genome could be sequenced.

Or consider the skepticism expressed in the mid-1980s that the Internet would ever be a significant phenomenon, given that it then included only tens of thousands of nodes (also known as servers). In fact, the number of nodes was doubling every year, so that there were likely to be tens of millions of nodes ten years later. But this trend was not appreciated by those who struggled with state-of-the-art technology in 1985, which permitted adding only a few thousand nodes throughout the world in a single year.⁶

The converse conceptual error occurs when certain exponential phenomena are first recognized and are applied in an overly aggressive manner without modeling the appropriate pace of growth. While exponential growth gains speed over time, it is not instantaneous. The run-up in capital values (that is, stock market prices) during the "Internet bubble" and related telecommunications bubble (1997–2000) was greatly in excess of any reasonable expectation of even exponential growth. As I demonstrate in the next chapter, the actual adoption of the Internet and e-commerce did show smooth exponential growth through both boom and bust; the overzealous expectation of growth affected only capital (stock) valuations. We have seen comparable mistakes during earlier paradigm shifts—for example, during the early railroad era (1830s), when the equivalent of the Internet boom and bust led to a frenzy of railroad expansion.

Another error that prognosticators make is to consider the transformations that will result from a single trend in today's world as if nothing else will change. A good example is the concern that radical life extension will result in overpopulation and the exhaustion of limited material resources to sustain human life, which ignores comparably radical wealth creation from nanotechnology and strong AI. For example, nanotechnology-based manufacturing devices in the 2020s will be capable of creating almost any physical product from inexpensive raw materials and information.

I emphasize the exponential-versus-linear perspective because it's the most

important failure that prognosticators make in considering future trends. Most technology forecasts and forecasters ignore altogether this historical exponential view of technological progress. Indeed, almost everyone I meet has a linear view of the future. That's why people tend to overestimate what can be achieved in the short term (because we tend to leave out necessary details) but underestimate what can be achieved in the long term (because exponential growth is ignored).

The Six Epochs

First we build the tools, then they build us.

—MARSHALL MCLUHAN

The future ain't what it used to be.

—YOGI BERRA

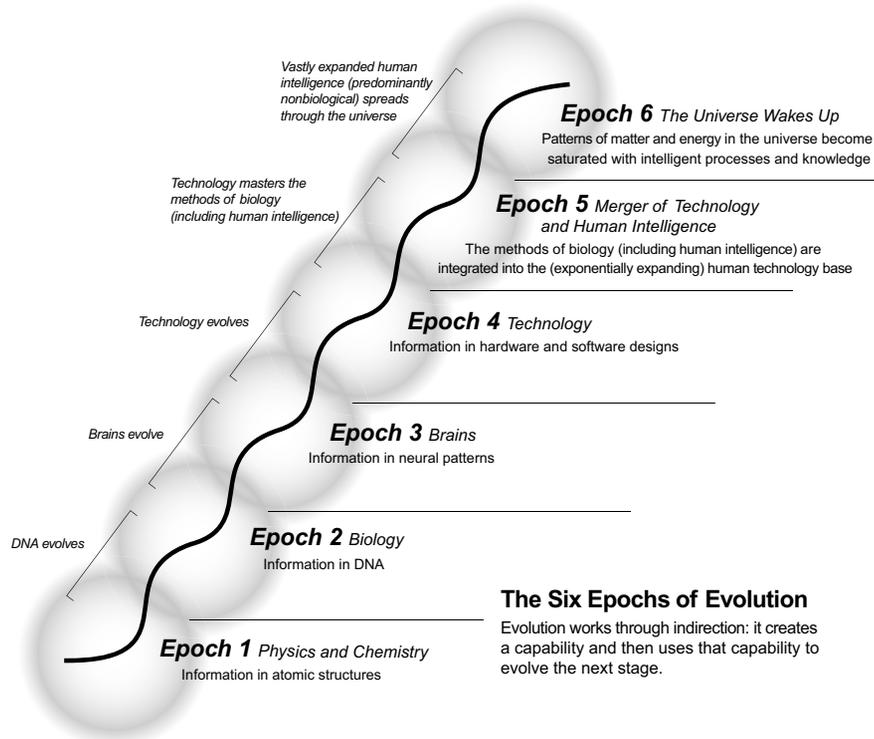
Evolution is a process of creating patterns of increasing order. I'll discuss the concept of order in the next chapter; the emphasis in this section is on the concept of patterns. I believe that it's the evolution of patterns that constitutes the ultimate story of our world. Evolution works through indirection: each stage or epoch uses the information-processing methods of the previous epoch to create the next. I conceptualize the history of evolution—both biological and technological—as occurring in six epochs. As we will discuss, the Singularity will begin with Epoch Five and will spread from Earth to the rest of the universe in Epoch Six.

Epoch One: Physics and Chemistry. We can trace our origins to a state that represents information in its basic structures: patterns of matter and energy. Recent theories of quantum gravity hold that time and space are broken down into discrete quanta, essentially fragments of information. There is controversy as to whether matter and energy are ultimately digital or analog in nature, but regardless of the resolution of this issue, we do know that atomic structures store and represent discrete information.

A few hundred thousand years after the Big Bang, atoms began to form, as electrons became trapped in orbits around nuclei consisting of protons and neutrons. The electrical structure of atoms made them “sticky.” Chemistry was born a few million years later as atoms came together to create relatively stable structures called molecules. Of all the elements, carbon proved to be the most versatile; it's able to form bonds in four directions (versus one to three for

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



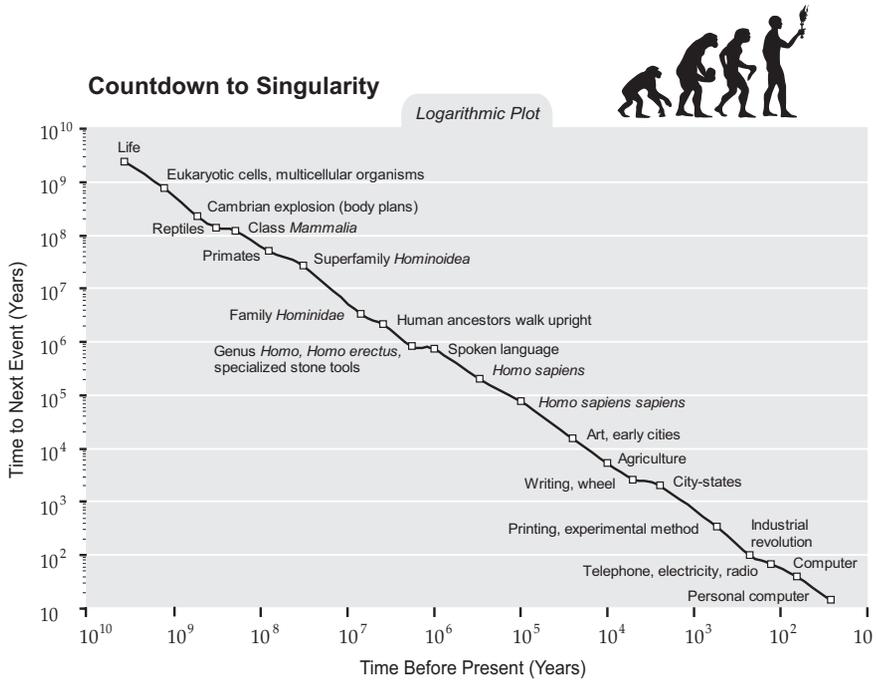
Epoch Two: Biology and DNA. In the second epoch, starting several billion years ago, carbon-based compounds became more and more intricate until complex aggregations of molecules formed self-replicating mechanisms, and life originated. Ultimately, biological systems evolved a precise digital mechanism (DNA) to store information describing a larger society of molecules. This molecule and its supporting machinery of codons and ribosomes enabled a record to be kept of the evolutionary experiments of this second epoch.

Epoch Three: Brains. Each epoch continues the evolution of information through a paradigm shift to a further level of “indirection.” (That is, evolution uses the results of one epoch to create the next.) For example, in the third epoch, DNA-guided evolution produced organisms that could detect information with their own sensory organs and process and store that information in their own brains and nervous systems. These were made possible by second-epoch mechanisms (DNA and epigenetic information of proteins and RNA fragments that control gene expression), which (indirectly) enabled and defined third-epoch information-processing mechanisms (the brains and nervous systems of organisms). The third epoch started with the ability of early animals to recognize patterns, which still accounts for the vast majority of the activity in our brains.¹⁰ Ultimately, our own species evolved the ability to create abstract mental models of the world we experience and to contemplate the rational implications of these models. We have the ability to redesign the world in our own minds and to put these ideas into action.

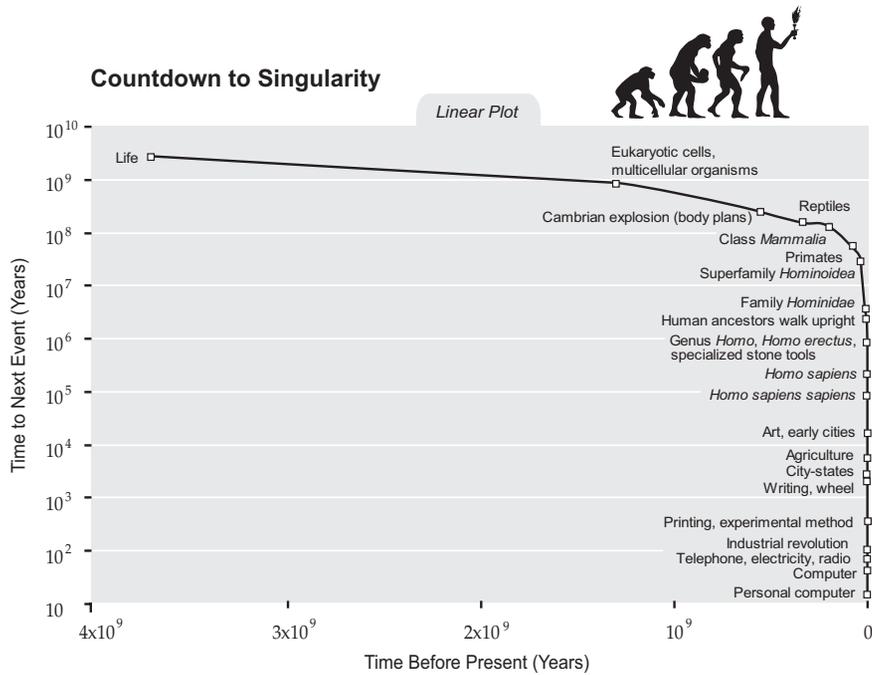
Epoch Four: Technology. Combining the endowment of rational and abstract thought with our opposable thumb, our species ushered in the fourth epoch and the next level of indirection: the evolution of human-created technology. This started out with simple mechanisms and developed into elaborate automata (automated mechanical machines). Ultimately, with sophisticated computational and communication devices, technology was itself capable of sensing, storing, and evaluating elaborate patterns of information. To compare the rate of progress of the biological evolution of intelligence to that of technological evolution, consider that the most advanced mammals have added about one cubic inch of brain matter every hundred thousand years, whereas we are roughly doubling the computational capacity of computers every year (see the next chapter). Of course, neither brain size nor computer capacity is the sole determinant of intelligence, but they do represent enabling factors.

If we place key milestones of both biological evolution and human technological development on a single graph plotting both the *x*-axis (number of

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

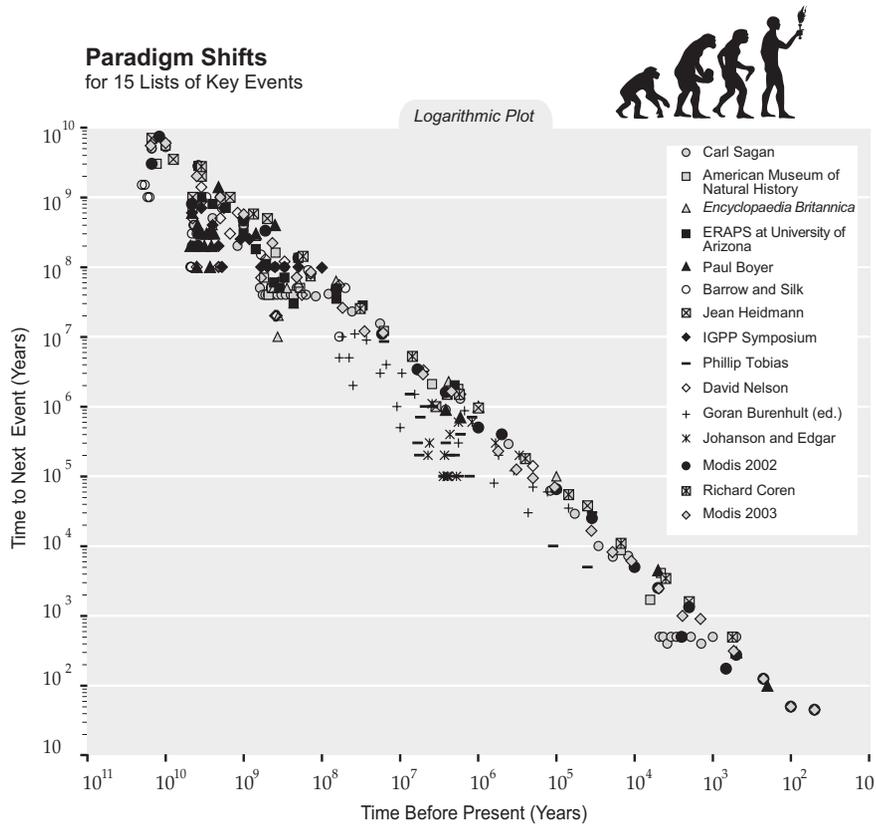


Countdown to Singularity: Biological evolution and human technology both show continual acceleration, indicated by the shorter time to the next event (two billion years from the origin of life to cells; fourteen years from the PC to the World Wide Web).



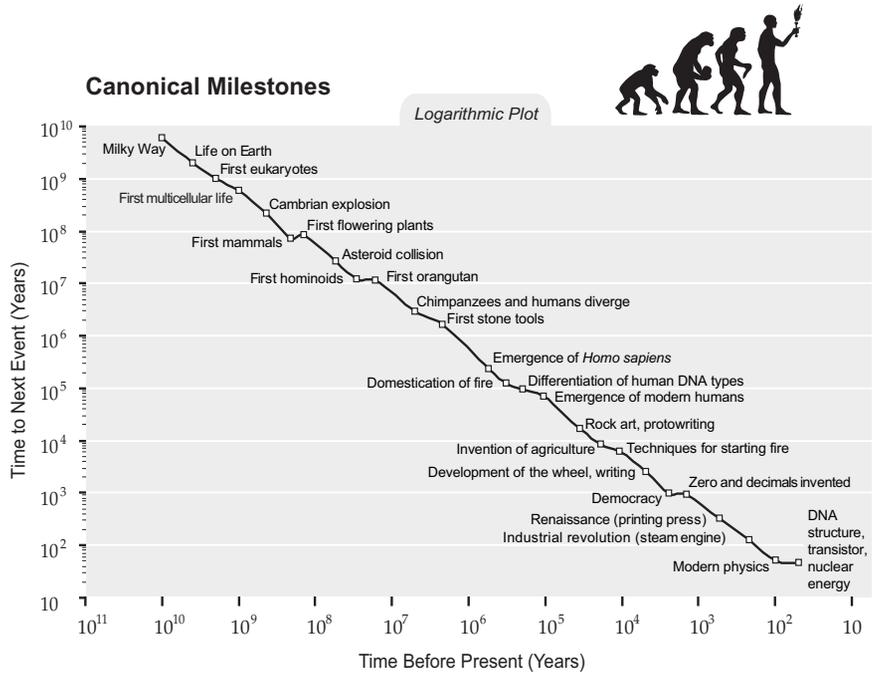
Linear view of evolution: This version of the preceding figure uses the same data but with a linear scale for time before present instead of a logarithmic one. This shows the acceleration more dramatically, but details are not visible. From a linear perspective, most key events have just happened “recently.”

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

The Singularity will allow us to overcome age-old human problems and vastly amplify human creativity. We will preserve and enhance the intelligence that evolution has bestowed on us while overcoming the profound limitations of biological evolution. But the Singularity will also amplify the ability to act on our destructive inclinations, so its full story has not yet been written.

Epoch Six: The Universe Wakes Up. I will discuss this topic in chapter 6, under the heading “. . . on the Intelligent Destiny of the Cosmos.” In the aftermath of the Singularity, intelligence, derived from its biological origins in human brains and its technological origins in human ingenuity, will begin to saturate the matter and energy in its midst. It will achieve this by reorganizing matter and energy to provide an optimal level of computation (based on limits we will discuss in chapter 3) to spread out from its origin on Earth.

We currently understand the speed of light as a bounding factor on the transfer of information. Circumventing this limit has to be regarded as highly speculative, but there are hints that this constraint may be able to be superseded.¹⁵ If there are even subtle deviations, we will ultimately harness this superluminal ability. Whether our civilization infuses the rest of the universe with its creativity and intelligence quickly or slowly depends on its immutability. In any event the “dumb” matter and mechanisms of the universe will be transformed into exquisitely sublime forms of intelligence, which will constitute the sixth epoch in the evolution of patterns of information.

This is the ultimate destiny of the Singularity and of the universe.

The Singularity Is Near

You know, things are going to be really different! . . . No, no, I mean really different!

—MARK MILLER (COMPUTER SCIENTIST) TO ERIC DREXLER, AROUND 1986

What are the consequences of this event? When greater-than-human intelligence drives progress, that progress will be much more rapid. In fact, there seems no reason why progress itself would not involve the creation of still more intelligent entities—on a still-shorter time scale. The best analogy that I see is with the evolutionary past: Animals can adapt to problems and make inventions, but often no faster than natural selection can do its work—the world acts as its own simulator in the case of natural selection. We humans have the ability to internalize the world and conduct “what if’s” in our

22 THE SINGULARITY IS NEAR

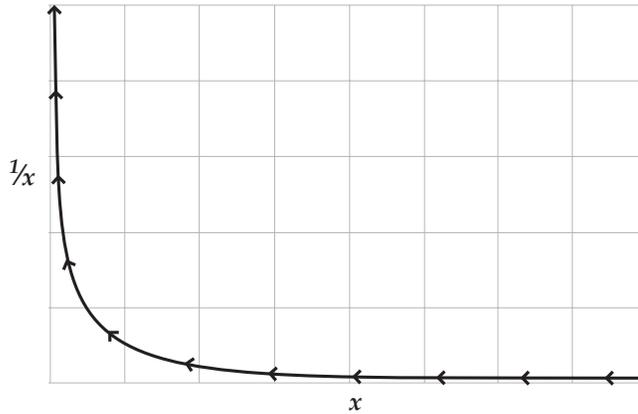
heads; we can solve many problems thousands of times faster than natural selection. Now, by creating the means to execute those simulations at much higher speeds, we are entering a regime as radically different from our human past as we humans are from the lower animals. From the human point of view, this change will be a throwing away of all the previous rules, perhaps in the blink of an eye, an exponential runaway beyond any hope of control.

—VERNOR VINGE, “THE TECHNOLOGICAL SINGULARITY,” 1993

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

—IRVING JOHN GOOD, “SPECULATIONS CONCERNING THE FIRST
ULTRA-INTELLIGENT MACHINE,” 1965

To put the concept of Singularity into further perspective, let's explore the history of the word itself. “Singularity” is an English word meaning a unique event with, well, singular implications. The word was adopted by mathematicians to denote a value that transcends any finite limitation, such as the explosion of magnitude that results when dividing a constant by a number that gets closer and closer to zero. Consider, for example, the simple function $y = 1/x$. As the value of x approaches zero, the value of the function (y) explodes to larger and larger values.

A Mathematical Singularity *Linear Plot*

A mathematical singularity: As x approaches zero (from right to left), $1/x$ (or y) approaches infinity.

Such a mathematical function never actually achieves an infinite value, since dividing by zero is mathematically “undefined” (impossible to calculate). But the value of y exceeds any possible finite limit (approaches infinity) as the divisor x approaches zero.

The next field to adopt the word was astrophysics. If a massive star undergoes a supernova explosion, its remnant eventually collapses to the point of apparently zero volume and infinite density, and a “singularity” is created at its center. Because light was thought to be unable to escape the star after it reached this infinite density,¹⁶ it was called a black hole.¹⁷ It constitutes a rupture in the fabric of space and time.

One theory speculates that the universe itself began with such a Singularity.¹⁸ Interestingly, however, the event horizon (surface) of a black hole is of finite size, and gravitational force is only theoretically infinite at the zero-size center of the black hole. At any location that could actually be measured, the forces are finite, although extremely large.

The first reference to the Singularity as an event capable of rupturing the fabric of human history is John von Neumann’s statement quoted above. In the 1960s, I. J. Good wrote of an “intelligence explosion” resulting from intelligent machines’ designing their next generation without human intervention. Vernor Vinge, a mathematician and computer scientist at San Diego State University, wrote about a rapidly approaching “technological singularity” in an article for *Omni* magazine in 1983 and in a science-fiction novel, *Marooned in Real-time*, in 1986.¹⁹

My 1989 book, *The Age of Intelligent Machines*, presented a future headed inevitably toward machines greatly exceeding human intelligence in the first half of the twenty-first century.²⁰ Hans Moravec's 1988 book *Mind Children* came to a similar conclusion by analyzing the progression of robotics.²¹ In 1993 Vinge presented a paper to a NASA-organized symposium that described the Singularity as an impending event resulting primarily from the advent of "entities with greater than human intelligence," which Vinge saw as the harbinger of a runaway phenomenon.²² My 1999 book, *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*, described the increasingly intimate connection between our biological intelligence and the artificial intelligence we are creating.²³ Hans Moravec's book *Robot: Mere Machine to Transcendent Mind*, also published in 1999, described the robots of the 2040s as our "evolutionary heirs," machines that will "grow from us, learn our skills, and share our goals and values, . . . children of our minds."²⁴ Australian scholar Damien Broderick's 1997 and 2001 books, both titled *The Spike*, analyzed the pervasive impact of the extreme phase of technology acceleration anticipated within several decades.²⁵ In an extensive series of writings, John Smart has described the Singularity as the inevitable result of what he calls "MEST" (matter, energy, space, and time) compression.²⁶

From my perspective, the Singularity has many faces. It represents the nearly vertical phase of exponential growth that occurs when the rate is so extreme that technology appears to be expanding at infinite speed. Of course, from a mathematical perspective, there is no discontinuity, no rupture, and the growth rates remain finite, although extraordinarily large. But from our *currently* limited framework, this imminent event appears to be an acute and abrupt break in the continuity of progress. I emphasize the word "currently" because one of the salient implications of the Singularity will be a change in the nature of our ability to understand. We will become vastly smarter as we merge with our technology.

Can the pace of technological progress continue to speed up indefinitely? Isn't there a point at which humans are unable to think fast enough to keep up? For unenhanced humans, clearly so. But what would 1,000 scientists, each 1,000 times more intelligent than human scientists today, and each operating 1,000 times faster than contemporary humans (because the information processing in their primarily nonbiological brains is faster) accomplish? One chronological year would be like a millennium for them.²⁷ What would they come up with?

Well, for one thing, they would come up with technology to become even more intelligent (because their intelligence is no longer of fixed capacity). They would change their own thought processes to enable them to think even faster.

When scientists become a million times more intelligent and operate a million times faster, an hour would result in a century of progress (in today's terms).

The Singularity involves the following principles, which I will document, develop, analyze, and contemplate throughout the rest of this book:

- The rate of paradigm shift (technical innovation) is accelerating, right now doubling every decade.²⁸
- The power (price-performance, speed, capacity, and bandwidth) of information technologies is growing exponentially at an even faster pace, now doubling about every year.²⁹ This principle applies to a wide range of measures, including the amount of human knowledge.
- For information technologies, there is a second level of exponential growth: that is, exponential growth in the rate of exponential growth (the exponent). The reason: as a technology becomes more cost effective, more resources are deployed toward its advancement, so the rate of exponential growth increases over time. For example, the computer industry in the 1940s consisted of a handful of now historically important projects. Today total revenue in the computer industry is more than one trillion dollars, so research and development budgets are comparably higher.
- Human brain scanning is one of these exponentially improving technologies. As I will show in chapter 4, the temporal and spatial resolution and bandwidth of brain scanning are doubling each year. We are just now obtaining the tools sufficient to begin serious reverse engineering (decoding) of the human brain's principles of operation. We already have impressive models and simulations of a couple dozen of the brain's several hundred regions. Within two decades, we will have a detailed understanding of how all the regions of the human brain work.
- We will have the requisite hardware to emulate human intelligence with supercomputers by the end of this decade and with personal-computer-size devices by the end of the following decade. We will have effective software models of human intelligence by the mid-2020s.
- With both the hardware and software needed to fully emulate human intelligence, we can expect computers to pass the Turing test, indicating intelligence indistinguishable from that of biological humans, by the end of the 2020s.³⁰
- When they achieve this level of development, computers will be able to combine the traditional strengths of human intelligence with the strengths of machine intelligence.
- The traditional strengths of human intelligence include a formidable ability to recognize patterns. The massively parallel and self-organizing nature

of the human brain is an ideal architecture for recognizing patterns that are based on subtle, invariant properties. Humans are also capable of learning new knowledge by applying insights and inferring principles from experience, including information gathered through language. A key capability of human intelligence is the ability to create mental models of reality and to conduct mental “what-if” experiments by varying aspects of these models.

- The traditional strengths of machine intelligence include the ability to remember billions of facts precisely and recall them instantly.
- Another advantage of nonbiological intelligence is that once a skill is mastered by a machine, it can be performed repeatedly at high speed, at optimal accuracy, and without tiring.
- Perhaps most important, machines can share their knowledge at extremely high speed, compared to the very slow speed of human knowledge-sharing through language.
- Nonbiological intelligence will be able to download skills and knowledge from other machines, eventually also from humans.
- Machines will process and switch signals at close to the speed of light (about three hundred million meters per second), compared to about one hundred meters per second for the electrochemical signals used in biological mammalian brains.³¹ This speed ratio is at least three million to one.
- Machines will have access via the Internet to all the knowledge of our human-machine civilization and will be able to master all of this knowledge.
- Machines can pool their resources, intelligence, and memories. Two machines—or one million machines—can join together to become one and then become separate again. Multiple machines can do both at the same time: become one and separate simultaneously. Humans call this falling in love, but our biological ability to do this is fleeting and unreliable.
- The combination of these traditional strengths (the pattern-recognition ability of biological human intelligence and the speed, memory capacity and accuracy, and knowledge and skill-sharing abilities of nonbiological intelligence) will be formidable.
- Machine intelligence will have complete freedom of design and architecture (that is, they won’t be constrained by biological limitations, such as the slow switching speed of our interneuronal connections or a fixed skull size) as well as consistent performance at all times.
- Once nonbiological intelligence combines the traditional strengths of both humans and machines, the nonbiological portion of our civilization’s

intelligence will then continue to benefit from the double exponential growth of machine price-performance, speed, and capacity.

- Once machines achieve the ability to design and engineer technology as humans do, only at far higher speeds and capacities, they will have access to their own designs (source code) and the ability to manipulate them. Humans are now accomplishing something similar through biotechnology (changing the genetic and other information processes underlying our biology), but in a much slower and far more limited way than what machines will be able to achieve by modifying their own programs.
- Biology has inherent limitations. For example, every living organism must be built from proteins that are folded from one-dimensional strings of amino acids. Protein-based mechanisms are lacking in strength and speed. We will be able to reengineer all of the organs and systems in our biological bodies and brains to be vastly more capable.
- As we will discuss in chapter 4, human intelligence does have a certain amount of plasticity (ability to change its structure), more so than had previously been understood. But the architecture of the human brain is nonetheless profoundly limited. For example, there is room for only about one hundred trillion interneuronal connections in each of our skulls. A key genetic change that allowed for the greater cognitive ability of humans compared to that of our primate ancestors was the development of a larger cerebral cortex as well as the development of increased volume of gray-matter tissue in certain regions of the brain.³² This change occurred, however, on the very slow timescale of biological evolution and still involves an inherent limit to the brain's capacity. Machines will be able to reformulate their own designs and augment their own capacities without limit. By using nanotechnology-based designs, their capabilities will be far greater than biological brains without increased size or energy consumption.
- Machines will also benefit from using very fast three-dimensional molecular circuits. Today's electronic circuits are more than one million times faster than the electrochemical switching used in mammalian brains. Tomorrow's molecular circuits will be based on devices such as nanotubes, which are tiny cylinders of carbon atoms that measure about ten atoms across and are five hundred times smaller than today's silicon-based transistors. Since the signals have less distance to travel, they will also be able to operate at terahertz (trillions of operations per second) speeds compared to the few gigahertz (billions of operations per second) speeds of current chips.
- The rate of technological change will not be limited to human mental

speeds. Machine intelligence will improve its own abilities in a feedback cycle that unaided human intelligence will not be able to follow.

- This cycle of machine intelligence's iteratively improving its own design will become faster and faster. This is in fact exactly what is predicted by the formula for continued acceleration of the rate of paradigm shift. One of the objections that has been raised to the continuation of the acceleration of paradigm shift is that it ultimately becomes much too fast for humans to follow, and so therefore, it's argued, it cannot happen. However, the shift from biological to nonbiological intelligence will enable the trend to continue.
- Along with the accelerating improvement cycle of nonbiological intelligence, nanotechnology will enable the manipulation of physical reality at the molecular level.
- Nanotechnology will enable the design of nanobots: robots designed at the molecular level, measured in microns (millionths of a meter), such as "respirocytes" (mechanical red-blood cells).³³ Nanobots will have myriad roles within the human body, including reversing human aging (to the extent that this task will not already have been completed through biotechnology, such as genetic engineering).
- Nanobots will interact with biological neurons to vastly extend human experience by creating virtual reality from within the nervous system.
- Billions of nanobots in the capillaries of the brain will also vastly extend human intelligence.
- Once nonbiological intelligence gets a foothold in the human brain (this has already started with computerized neural implants), the machine intelligence in our brains will grow exponentially (as it has been doing all along), at least doubling in power each year. In contrast, biological intelligence is effectively of fixed capacity. Thus, the nonbiological portion of our intelligence will ultimately predominate.
- Nanobots will also enhance the environment by reversing pollution from earlier industrialization.
- Nanobots called foglets that can manipulate image and sound waves will bring the morphing qualities of virtual reality to the real world.³⁴
- The human ability to understand and respond appropriately to emotion (so-called emotional intelligence) is one of the forms of human intelligence that will be understood and mastered by future machine intelligence. Some of our emotional responses are tuned to optimize our intelligence in the context of our limited and frail biological bodies. Future machine intelligence will also have "bodies" (for example, virtual

bodies in virtual reality, or projections in real reality using foglets) in order to interact with the world, but these nanoengineered bodies will be far more capable and durable than biological human bodies. Thus, some of the “emotional” responses of future machine intelligence will be re-designed to reflect their vastly enhanced physical capabilities.³⁵

- As virtual reality from within the nervous system becomes competitive with real reality in terms of resolution and believability, our experiences will increasingly take place in virtual environments.
- In virtual reality, we can be a different person both physically and emotionally. In fact, other people (such as your romantic partner) will be able to select a different body for you than you might select for yourself (and vice versa).
- The law of accelerating returns will continue until nonbiological intelligence comes close to “saturating” the matter and energy in our vicinity of the universe with our human-machine intelligence. By saturating, I mean utilizing the matter and energy patterns for computation to an optimal degree, based on our understanding of the physics of computation. As we approach this limit, the intelligence of our civilization will continue its expansion in capability by spreading outward toward the rest of the universe. The speed of this expansion will quickly achieve the maximum speed at which information can travel.
- Ultimately, the entire universe will become saturated with our intelligence. This is the destiny of the universe. (See chapter 6.) We will determine our own fate rather than have it determined by the current “dumb,” simple, machinelike forces that rule celestial mechanics.
- The length of time it will take the universe to become intelligent to this extent depends on whether or not the speed of light is an immutable limit. There are indications of possible subtle exceptions (or circumventions) to this limit, which, if they exist, the vast intelligence of our civilization at this future time will be able to exploit.

This, then, is the Singularity. Some would say that we cannot comprehend it, at least with our current level of understanding. For that reason, we cannot look past its event horizon and make complete sense of what lies beyond. This is one reason we call this transformation the Singularity.

I have personally found it difficult, although not impossible, to look beyond this event horizon, even after having thought about its implications for several decades. Still, my view is that, despite our profound limitations of thought, we do have sufficient powers of abstraction to make meaningful statements about

the nature of life after the Singularity. Most important, the intelligence that will emerge will continue to represent the human civilization, which is already a human-machine civilization. In other words, future machines will be human, even if they are not biological. This will be the next step in evolution, the next high-level paradigm shift, the next level of indirection. Most of the intelligence of our civilization will ultimately be nonbiological. By the end of this century, it will be trillions of trillions of times more powerful than human intelligence.³⁶ However, to address often-expressed concerns, this does not imply the end of biological intelligence, even if it is thrown from its perch of evolutionary superiority. Even the nonbiological forms will be derived from biological design. Our civilization will remain human—indeed, in many ways it will be more exemplary of what we regard as human than it is today, although our understanding of the term will move beyond its biological origins.

Many observers have expressed alarm at the emergence of forms of nonbiological intelligence superior to human intelligence (an issue we will explore further in chapter 9). The potential to augment our own intelligence through intimate connection with other thinking substrates does not necessarily alleviate the concern, as some people have expressed the wish to remain “unenhanced” while at the same time keeping their place at the top of the intellectual food chain. From the perspective of biological humanity, these superhuman intelligences will appear to be our devoted servants, satisfying our needs and desires. But fulfilling the wishes of a revered biological legacy will occupy only a trivial portion of the intellectual power that the Singularity will bring.

MOLLY CIRCA 2004: *How will I know when the Singularity is upon us? I mean, I'll want some time to prepare.*

RAY: *Why, what are you planning to do?*

MOLLY 2004: *Let's see, for starters, I'll want to fine-tune my résumé. I'll want to make a good impression on the powers that be.*

GEORGE CIRCA 2048: *Oh, I can take care of that for you.*

MOLLY 2004: *That's really not necessary. I'm perfectly capable of doing it myself. I might also want to erase a few documents—you know, where I'm a little insulting to a few machines I know.*

GEORGE 2048: *Oh, the machines will find them anyway—but don't worry, we're very understanding.*

MOLLY 2004: *For some reason, that's not entirely reassuring. But I'd still like to know what the harbingers will be.*

RAY: *Okay, you will know the Singularity is coming when you have a million e-mails in your in-box.*

MOLLY 2004: *Hmm, in that case, it sounds like we're just about there. But seriously, I'm having trouble keeping up with all of this stuff flying at me as it is. How am I going to keep up with the pace of the Singularity?*

GEORGE 2048: *You'll have virtual assistants—actually, you'll need just one.*

MOLLY 2004: *Which I suppose will be you?*

GEORGE 2048: *At your service.*

MOLLY 2004: *That's just great. You'll take care of everything, you won't even have to keep me informed. "Oh, don't bother telling Molly what's happening, she won't understand anyway, let's just keep her happy and in the dark."*

GEORGE 2048: *Oh, that won't do, not at all.*

MOLLY 2004: *The happy part, you mean?*

GEORGE 2048: *I was referring to keeping you in the dark. You'll be able to grasp what I'm up to if that's what you really want.*

MOLLY 2004: *What, by becoming . . .*

RAY: *Enhanced?*

MOLLY 2004: *Yes, that's what I was trying to say.*

GEORGE 2048: *Well, if our relationship is to be all that it can be, then it's not a bad idea.*

MOLLY 2004: *And should I wish to remain as I am?*

GEORGE 2048: *I'll be devoted to you in any event. But I can be more than just your transcendent servant.*

MOLLY 2004: *Actually, you're being "just" my transcendent servant doesn't sound so bad.*

CHARLES DARWIN: *If I may interrupt, it occurred to me that once machine intelligence is greater than human intelligence, it should be in a position to design its own next generation.*

MOLLY 2004: *That doesn't sound so unusual. Machines are used to design machines today.*

CHARLES: *Yes, but in 2004 they're still guided by human designers. Once machines are operating at human levels, well, then it kind of closes the loop.*

NED LUDD:³⁷ *And humans would be out of the loop.*

MOLLY 2004: *It would still be a pretty slow process.*

RAY: *Oh, not at all. If a nonbiological intelligence was constructed similarly to a human brain but used even circa 2004 circuitry, it—*

MOLLY CIRCA 2104: *You mean "she."*

RAY: *Yes, of course . . . she . . . would be able to think at least a million times faster.*

TIMOTHY LEARY: *So subjective time would be expanded.*

RAY: *Exactly.*

32 THE SINGULARITY IS NEAR

MOLLY 2004: *Sounds like a lot of subjective time. What are you machines going to do with so much of it?*

GEORGE 2048: *Oh, there's plenty to do. After all, I have access to all human knowledge on the Internet.*

MOLLY 2004: *Just the human knowledge? What about all the machine knowledge?*

GEORGE 2048: *We like to think of it as one civilization.*

CHARLES: *So, it does appear that machines will be able to improve their own design.*

MOLLY 2004: *Oh, we humans are starting to do that now.*

RAY: *But we're just tinkering with a few details. Inherently, DNA-based intelligence is just so very slow and limited.*

CHARLES: *So the machines will design their own next generation rather quickly.*

GEORGE 2048: *Indeed, in 2048, that is certainly the case.*

CHARLES: *Just what I was getting at, a new line of evolution then.*

NED: *Sounds more like a precarious runaway phenomenon.*

CHARLES: *Basically, that's what evolution is.*

NED: *But what of the interaction of the machines with their progenitors? I mean, I don't think I'd want to get in their way. I was able to hide from the English authorities for a few years in the early 1800s, but I suspect that will be more difficult with these . . .*

GEORGE 2048: *Guys.*

MOLLY 2004: *Hiding from those little robots—*

RAY: *Nanobots, you mean.*

MOLLY 2004: *Yes, hiding from the nanobots will be difficult, for sure.*

RAY: *I would expect the intelligence that arises from the Singularity to have great respect for their biological heritage.*

GEORGE 2048: *Absolutely, it's more than respect, it's . . . reverence.*

MOLLY 2004: *That's great, George, I'll be your revered pet. Not what I had in mind.*

NED: *That's just how Ted Kaczynski puts it: we're going to become pets. That's our destiny, to become contented pets but certainly not free men.*

MOLLY 2004: *And what about this Epoch Six? If I stay biological, I'll be using up all this precious matter and energy in a most inefficient way. You'll want to turn me into, like, a billion virtual Mollys and Georges, each of them thinking a lot faster than I do now. Seems like there will be a lot of pressure to go over to the other side.*

RAY: *Still, you represent only a tiny fraction of the available matter and energy. Keeping you biological won't appreciably change the order of magnitude of matter and energy available to the Singularity. It will be well worth it to maintain the biological heritage.*

GEORGE 2048: *Absolutely.*

RAY: *Just like today we seek to preserve the rain forest and the diversity of species.*

MOLLY 2004: *That's just what I was afraid of. I mean, we're doing such a wonderful job with the rain forest. I think we still have a little bit of it left. We'll end up like those endangered species.*

NED: *Or extinct ones.*

MOLLY 2004: *And there's not just me. How about all the stuff I use? I go through a lot of stuff.*

GEORGE 2048: *That's not a problem, we'll just recycle all your stuff. We'll create the environments you need as you need them.*

MOLLY 2004: *Oh, I'll be in virtual reality?*

RAY: *No, actually, foglet reality.*

MOLLY 2004: *I'll be in a fog?*

RAY: *No, no, foglets.*

MOLLY 2004: *Excuse me?*

RAY: *I'll explain later in the book.*

MOLLY 2004: *Well, give me a hint.*

RAY: *Foglets are nanobots—robots the size of blood cells—that can connect themselves to replicate any physical structure. Moreover, they can direct visual and auditory information in such a way as to bring the morphing qualities of virtual reality into real reality.³⁸*

MOLLY 2004: *I'm sorry I asked. But, as I think about it, I want more than just my stuff. I want all the animals and plants, too. Even if I don't get to see and touch them all, I like to know they're there.*

GEORGE 2048: *But nothing will be lost.*

MOLLY 2004: *I know you keep saying that. But I mean actually there—you know, as in biological reality.*

RAY: *Actually, the entire biosphere is less than one millionth of the matter and energy in the solar system.*

CHARLES: *It includes a lot of the carbon.*

RAY: *It's still worth keeping all of it to make sure we haven't lost anything.*

GEORGE 2048: *That has been the consensus for at least several years now.*

MOLLY 2004: *So, basically, I'll have everything I need at my fingertips?*

GEORGE 2048: *Indeed.*

MOLLY 2004: *Sounds like King Midas. You know, everything he touched turned to gold.*

NED: *Yes, and as you will recall he died of starvation as a result.*

MOLLY 2004: *Well, if I do end up going over to the other side, with all of that vast expanse of subjective time, I think I'll die of boredom.*

GEORGE 2048: *Oh, that could never happen. I will make sure of it.*