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 each 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 the fate of Tom and 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 his friends come up with to save the day?[2] The moral of these tales was simple: the right idea had the power to overcome a seemingly overwhelming challenge.

To this day, I remain convinced of this basic philosophy: no matter what quandaries we face—business problems, health issues, relationship difficulties, as well as the great scientific, social, and cultural challenges of our time—there is an idea that can enable us to prevail. Furthermore, we can find that idea. And when we find it, we need to implement it. My life has been shaped by this imperative. The power of an idea—this is itself an idea.

Around the same time that I was reading the Tom Swift Jr. series, I recall my grandfather, who had also fled Europe with my mother, coming back from his first return visit to Europe with two key memories. One was the gracious treatment he received from the Austrians and Germans, the same people who had forced him to flee in 1938. The other was a rare opportunity he had been given to touch with his own hands some original manuscripts of Leonardo da Vinci. Both recollections influenced me, but the latter is one I've returned to many times. He described the experience with reverence, as if he had touched the work of God himself. This, then, was the religion that I was raised with: veneration for human creativity and the power of ideas.

In 1960, at the age of twelve, I discovered the computer and became fascinated with its ability to model and re-create the world. I hung around the surplus electronics stores on Canal Street in Manhattan (they're still there!) and gathered parts to build my own computational devices. During the 1960s, I was as absorbed in the contemporary musical, cultural, and political movements as my peers, but I became equally engaged in a much more obscure trend: namely, the remarkable sequence of machines that IBM proffered during that decade, from their big "7000" series (7070, 7074, 7090, 7094) to their small 1620, effectively the first "minicomputer." The machines were introduced at yearly intervals, and each one was less expensive and more powerful than the last, a phenomenon familiar today. I got access to an IBM 1620 and began to write programs for statistical analysis and subsequently for music composition.

I still recall the time in 1968 when I was allowed into the secure, cavernous chamber housing what was then the most powerful computer in New England, a top-of-the-line IBM 360 Model 91, with a remarkable million bytes (one megabyte) of "core" memory, an impressive speed of one million instructions per second (one MIPS), and a rental cost of only one thousand dollars per hour. I had developed a computer program that matched high-school students to colleges, and I watched in fascination as the front-panel lights danced through a distinctive pattern as the machine processed each student's application.[3] Even though I was quite familiar with every line of code, it nonetheless seemed as if the computer were deep in thought when the lights dimmed for several seconds at the denouement of each such cycle. Indeed, it could do flawlessly in ten seconds what took us ten hours to do manually with far less accuracy.

As an inventor in the 1970s, I came to realize that my inventions needed to make sense in terms of the enabling technologies and market forces that would exist when the inventions were introduced, as that world would be a very different one from the one in which they were conceived. I began to develop models of how distinct technologies—electronics, communications, computer processors, memory, magnetic storage, and others—developed and how these changes rippled through markets and ultimately our social institutions. I realized that most inventions fail not because the R&D department can't get them to work but because the timing is wrong. Inventing is a lot like surfing: you have to anticipate and catch the wave at just the right moment.

My interest in technology trends and their implications took on a life of its own in the 1980s, and I began to use my models to project and anticipate future technologies, innovations that would appear in 2000, 2010, 2020, and beyond. This enabled me to invent with the capabilities of the future by conceiving and designing inventions using these future capabilities. In the mid-to-late 1980s, I wrote my first book, *The Age of Intelligent Machines*.[4] It included extensive (and reasonably accurate) predictions for the 1990s and 2000s, and ended with the specter of machine intelligence becoming indistinguishable from that of its human progenitors within the first half of the twenty-first century. It seemed like a poignant conclusion, and in any event I personally found it difficult to look beyond so transforming an outcome.

Over the last twenty years, I have come to appreciate an important meta-idea: that the power of ideas to transform the world is itself accelerating. Although people readily agree with this observation when it is simply stated, relatively few observers truly appreciate its profound implications. Within the next several decades, we will have the opportunity to apply ideas to conquer age-old problems—and introduce a few new problems along the way.

During the 1990s, I gathered empirical data on the apparent acceleration of all information-related technologies and sought to refine the mathematical models underlying these observations. I developed a theory I call the law of accelerating returns, which explains why technology and evolutionary processes in general progress in an exponential fashion.[5] In *The Age of Spiritual Machines* (*ASM*), which I wrote in 1998, I sought to articulate the nature of human life as it would exist past the point when machine and human cognition blurred. Indeed, I've seen this epoch as an increasingly intimate collaboration between our biological heritage and a future that transcends biology.

Since the publication of *ASM*, I have begun to reflect on the future of our civilization and its relationship to our place in the universe. Although it may seem difficult to envision the capabilities of a future civilization whose intelligence vastly outstrips our own, our ability to create models of reality in our mind enables us to articulate meaningful insights into the implications of this impending merger of our biological thinking with the nonbiological intelligence we are creating. This, then, is the story I wish to tell in this book. The story is predicated on the idea that we have the ability to understand our own intelligence—to access our own source code, if you will—and then revise and expand it.

Some observers question whether we are capable of applying our own thinking to understand our own thinking. AI researcher Douglas Hofstadter muses that "it could be simply an accident of fate that our brains are too weak to understand themselves. Think of the lowly giraffe, for instance, whose brain is obviously far below the level required for self-understanding—yet it is remarkably similar to our brain."[6] However, we have already succeeded in modeling portions of our brain—neurons and substantial neural regions—and the complexity of such models is growing rapidly. Our progress in reverse engineering the human brain, a key issue that I will describe in detail in this book, demonstrates that we do indeed have the ability to understand, to model, and to extend our own intelligence. This is one aspect of the uniqueness of our species: our intelligence is just sufficiently above the critical threshold necessary for us to scale our own ability to unrestricted heights of creative power—and we have the opposable appendage (our thumbs) necessary to manipulate the universe to our will.

A word on magic: when I was reading the Tom Swift Jr. books, I was also an avid magician. I enjoyed the delight of my audiences in experiencing apparently impossible transformations of reality. In my teen years, I replaced my parlor magic with technology projects. I discovered that unlike mere tricks, technology does not lose its transcendent power when its secrets are revealed. I am often reminded of Arthur C. Clarke's third law, that "any sufficiently advanced technology is indistinguishable from magic."

Consider J. K. Rowling's Harry Potter stories from this perspective. These tales may be imaginary, but they are not unreasonable visions of our world as it will exist only a few decades from now. Essentially all of the Potter "magic" will be realized through the technologies I will explore in this book. Playing quidditch and transforming people and objects into other forms will be feasible in full-immersion virtual-reality environments, as well as in real reality, using nanoscale devices. More dubious is the time reversal (as described in *Harry Potter and the Prisoner of Azkaban*), although serious proposals have even been put forward for accomplishing something along these lines (without giving rise to causality paradoxes), at least for bits of information, which essentially is what we comprise. (See the discussion in chapter 3 on the ultimate limits of computation.)

Consider that Harry unleashes his magic by uttering the right incantation. Of course, discovering and applying these incantations are no simple matters. Harry and his colleagues need to get the sequence, procedures, and emphasis exactly correct. That process is precisely our experience with technology. Our incantations are the formulas and algorithms underlying our modern-day magic. With just the right sequence, we can get a computer to read a book out loud, understand human speech, anticipate (and prevent) a heart attack, or predict the movement of a stock-market holding. If an incantation is just slightly off mark, the magic is greatly weakened or does not work at all.

One might object to this metaphor by pointing out that Hogwartian incantations are brief and therefore do not contain much information compared to, say, the code for a modern software program. But the essential methods of modern technology generally share the same brevity. The principles of operation of software advances such as speech recognition can be written in just a few pages of formulas. Often a key advance is a matter of applying a small change to a single formula.

The same observation holds for the "inventions" of biological evolution: consider that the genetic difference between chimpanzees and humans, for example, is only a few hundred thousand bytes of information. Although chimps are capable of some intellectual feats, that tiny difference in our genes was sufficient for our species to create the magic of technology.

Muriel Rukeyser says that "the universe is made of stories, not of atoms." In chapter 7, I describe myself as a "patternist," someone who views patterns of information as the fundamental reality. For example, the particles composing my brain and body change within weeks, but there is a continuity to the patterns that these particles make. A story can be regarded as a meaningful pattern of information, so we can interpret Muriel Rukeyser's aphorism from this perspective. This book, then, is the story of the destiny of the human-machine civilization, a destiny we have come to refer to as the Singularity.