

Age of Invention

Technological optimists maintain that the impact of innovation on our lives is increasing, but the evidence goes the other way. The author's grandmother [see photo] lived from the 1880s through the 1960s and witnessed the adoption of electricity, phonographs, telephones, radio, television, airplanes, antibiotics, vacuum tubes, transistors, and the automobile. In 1924 she became one of the first in her neighborhood to own a car. The author contends that the inventions unveiled in his own lifetime have made a far smaller difference.

1837
• Oceangoing steamship

1842
• Ether anesthesia

1843
• Vulcanized rubber

1844 ▲
• Telegraph

1846
• Cylinder printing press

1857
• Passenger elevator

1859
• Theory of evolution
• Oil well

1876 >
• Telephone

1876-1920 >
• Germ theory of disease

1877
Toilet paper

1877 >
Phonograph

1879
Lightbulb

1885
Motorbike

1887 >
Aspirin

1889 ▲
Automobile

1896 >
Flashlight

1906
Vacuum tube

▲ 1902
Teddy bear

▲ 1903
Airplane

1905
Theory of special relativity

1914 >
Zipper

1909 >
Toaster

1830 1840 1850 1860 1870 1880 1890 1900 1910

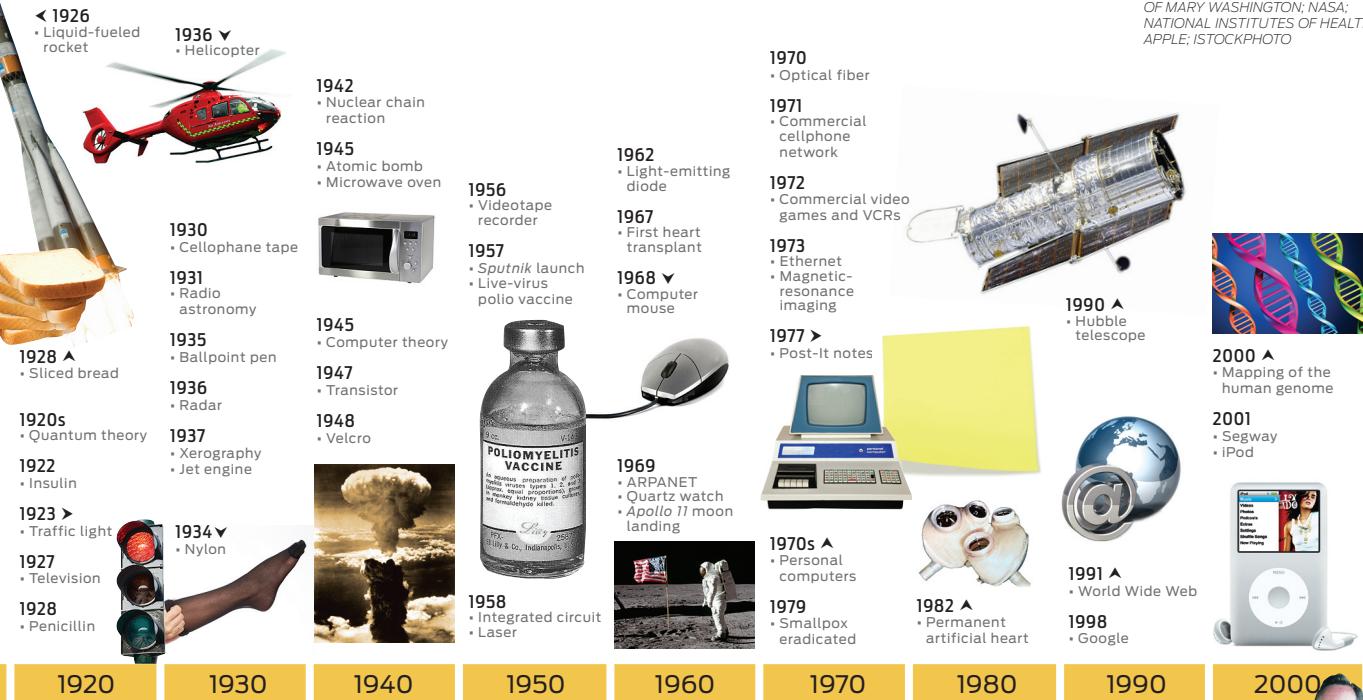
GRANDMOTHER'S LIFETIME

THE STORY OF THE SINGULARITY IS SWEEPING, DRAMATIC, SIMPLE—AND WRONG BY ALFRED NORDMANN

Singular Simplicity

TAKE THE IDEA of exponential technological growth, work it through to its logical conclusion, and there you have the singularity. Its bold incredibility pushes aside incredulity, as it challenges us to confront all the things we thought could never come true—the creation of superintelligent, conscious organisms, nanorobots that can swim in our bloodstreams and fix what ails us, and direct communication from mind to mind. And the pièce de résistance: a posthuman existence of disembodied uploaded minds, living on indefinitely without fear, sickness, or want in a virtual paradise ingeniously designed to delight, thrill, and stimulate.

PHOTO CREDITS: THE NATIONAL ARCHIVES; HONDA; NATIONAL MUSEUM OF AMERICAN HISTORY; RADIOMUSEUM KOELN; ESA; U.S. AIR FORCE; UNIVERSITY OF MARY WASHINGTON; NASA; NATIONAL INSTITUTES OF HEALTH; APPLE; ISTOCKPHOTO



AUTHOR'S LIFETIME



This vision argues that machines will become conscious and then perfect themselves, as described elsewhere in this issue. Yet for all its show of tough-minded audacity, the argument is shot through with sloppy reasoning, wishful thinking, and irresponsibility. Infatuated with statistics and seduced by the power of extrapolation, singularitarians abduct the moral imagination into a speculative no-man's-land. To be sure, they are hardly the first to spread fanciful technological prophecies, but among enthusiasts and doomsayers alike their proposition enjoys an inexplicable popularity. Perhaps the real question is how they have gotten away with it.

THE TROUBLE BEGINS with the singularitarians' assumption that technological advances have accelerated. I'd argue that I have seen less technological progress than my parents did, let alone my grandparents. Born in 1956, I can testify primarily to the development of the information age, fueled by the doubling of computing power every 18 to 24 months,

as described by Moore's Law. The birth-control pill and other reproductive technologies have had an equally profound impact, on the culture if not the economy, but they are not developing at an accelerating speed. Beyond that, I saw men walk on the moon, with little to come of it, and I am surrounded by bio- and nanotechnologies that so far haven't affected my life at all. Medical research has developed treatments that make a difference in our lives, particularly at the end of them. But despite daily announcements of one breakthrough or another, morbidity and mortality from cancer and stroke continue practically unabated, even in developed countries.

Now consider the life of someone who was born in the 1880s and died in the 1960s—my grandmother, for instance. She witnessed the introduction of electric light and telephones, of automobiles and airplanes, the atomic bomb and nuclear power, vacuum electronics and semiconductor electronics, plastics and the computer, most vaccines and all antibiotics. All of those things mattered greatly in human terms, as can be seen in a single statistic: child mortality in indus-

trialized countries dropped by 80 percent in those years.

So on what do intelligent people base the idea that technological progress is moving faster than ever before? It's simple: a chart of productivity from the dawn of humanity to the present day. It shows a line that inclines very gradually until around 1750, when it suddenly shoots almost straight up.

But that's hardly surprising. Since around 1750 the world has witnessed the spread of an economic system, by the name of capitalism, that is predicated on economic growth. And how the economy has grown since then! But surely the creation of new markets and the increasingly fine division of labor cannot be equated with technological progress, as every consumer knows.

Even if we were to accept, for the sake of argument, that technological innovation has truly accelerated, the line leading to the singularity would still be nothing but the simple-minded extrapolation of an existing pattern. Moore's Law has been remarkably successful at describing and predicting



EXPERT VIEW:
Jim Fruchterman

WHO HE IS
Founder and CEO of the Benetech Initiative, in Palo Alto, Calif., one of the first companies to focus on social entrepreneurship. Former rocket scientist and optical-character-recognition pioneer. Winner of a 2006 MacArthur Fellowship, the so-called genius grant.

SINGULARITY WILL OCCUR
Within 70 years

MACHINE CONSCIOUSNESS WILL OCCUR
Yes

MOORE'S LAW WILL CONTINUE FOR
30 more years

THOUGHTS
"I believe the singularity theory is plausible in that there will be a major shift in the rate of technology change. I am less convinced by projections of what it will mean to humans and humanity, such as human downloading in our lifetimes.

"Two things that rarely come up are the bug and algorithm questions. Douglas Hofstadter has more or less proved that perfect programs are not practically possible. And algorithms don't scale as nicely as processing power does: $n \log(n)$ is not our friend. A Linux system that needs rebooting only every three years is a modern technological marvel. But do you want to reboot your brain regularly?"

the development of semiconductors, in part because it has molded that development, ever since the semiconductor manufacturing industry adopted it as its road map and began spending vast sums on R&D to meet its requirements. Yet researchers and developers in the semiconductor industry have never denied that Moore's Law will finally come up against physical limits—indeed, many fear that the day of reckoning is nigh—whereas singularitarians happily extrapolate the law indefinitely into the future. And just as the semiconductor industry wonders nervously whether nanotechnology really can give Moore's Law another lease on life, singularitarians accept that this will occur as a given and then appropriate the exponential growth curve of Moore's Law not only to all the nano- and biotechnologies but to the cognitive sciences as well.

A typical example is the therapeutic development of brain-machine interfaces. In 2002, people were able to transmit 2 bits per minute to a computer. Four years later that figure had risen to 40 bits—that is, five letters—per minute. If this rate of progress continues, the argument goes, then by 2020 brain communication with computers will be as fast as speech. This isn't just the breathless cant of a true believer. The idea that an enhanced communication of thoughts will exceed

speech can also be found in the 2002 report "Converging Technologies for Improving Human Performance," issued by the U.S. National Science Foundation and the Department of Commerce. It says that such methods "could complement verbal communication, sometimes replacing spoken language when speed is a priority or enhancing speech when needed to exploit maximum mental capabilities." Presumably, the singularity will be reached soon afterward, when transmission rates exceed the speed of thought itself, allowing the computer to transmit our thoughts before we think them.

This fantastic vision works only by ignoring the critical limit, which is the great concentration you have to muster to send the bits. It is a procedure far more tedious than speech. To ease that requirement—to make a brain-machine interface into a true mind-machine interface—we'd have to know a lot more than we do about the relation between specific thoughts and corresponding physical processes in the brain.

The seductive power of extrapolation has also been applied in ways less spectacular but no less foolish. The "lab on a chip" and other technologies for biochemical analysis have significantly increased the number of measurements—blood lipids, for instance—that can be obtained from a single drop of blood. It's a fine achieve-

MICHAEL CALLOPY/THE SINGULARITY FOUNDATION

Semiconductors and the Singularity

A COMMON BELIEF in singularity circles is that Moore's Law will not only continue indefinitely but will also apply to other areas of technology, including some yet to be invented. First expressed in 1965 by Intel cofounder Gordon E. Moore, Moore's Law describes the periodic doubling of the density of components on semiconductors. Sustained by enormous investments in R&D, the law still holds true, 43 years later, much to the surprise of Moore himself.

In his interesting book *The Singularity Is Near*, inventor and entrepreneur Ray Kurzweil accurately reviews the history of Moore's Law and describes some of the near-term challenges. Kurzweil then declares his belief that the industry is on the verge of a new paradigm that again increases the exponential growth rate of computational power, this time through the introduction of three-dimensional integrated circuits.

He may be right. That idea is very much in line with the views of industry leaders, who foresee not only stacked memories

in the 2010- to 2013 time frame but also hybrid devices that incorporate logic, sensors, and microelectromechanical devices in the same 3-D package. An international research effort known as More Than Moore is now looking for ways to build extremely high-performance ICs for many different applications by combining existing technologies on silicon or in the same package. Nanowires in 3-D chips will help keep Moore's Law going by letting designers pack in transistors vertically as well as horizontally, in much the same way that skyscrapers in New York City pack in more people by housing them above and below one another.

Kurzweil goes on to speculate that the advanced 3-D molecular-computing devices of the future will more likely be grown in a test tube than made in a silicon fab. I find this forecast much less plausible. There is no such process that shows any promise of coming anywhere near silicon's manufacturing capabilities and economies of scale. There are research efforts that aim to build biological structures on silicon

ment, no doubt, but visionaries stretch the imagination when they assume that a second Moore's Law is about to produce astounding success stories and a transformation of all medical diagnostics.

Yet that assumption, which extrapolates an extrapolation—Moore's Law—to another field, is precisely what lies behind the now commonly expressed fear that increasing diagnostic powers are creating ethical problems in medicine. Physicians, we are told, will routinely inform patients of impending diseases for which they can offer no cure.

Yet in fact the path is very long from quicker blood analysis to instantaneous detection of the near certainty of a dread disease in a patient's future. A lab on a chip may provide mountains of data, but without great advances in many other fields—notably systems biology, pathology, and physiology—no one will be able to do much with it. Doctors already have more physiological information than they can profitably use.

Both examples of mindless extrapolation constitute wishful thinking. And in both cases, public debate is diverted from the real moral issues and quandaries that technology raises.

Rather than dream about how technology will soon effect an almost magical transformation of human life, societies need to debate the many real problems

connected with technological changes that are already under way. These problems belong to the here and now.

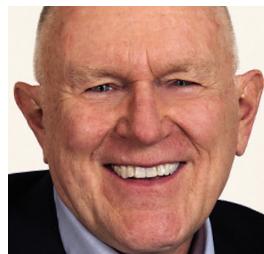
WHY, THEN, ARE SO many people captivated by the simple story of exponential growth that culminates in a life-altering singularity? Part of the appeal lies in simplicity itself, part in technological optimism—yet both of these tendencies are very old. What's new, though, is the changing role of technical expertise.

Plainly put, it is getting harder than ever to know whom to believe. Policy makers and members of the public have always had to put a degree of trust in experts. But now, when considering highly complex phenomena—in cellular processes, in chips containing billions of transistors, or in programs numbering hundreds of thousands of lines of code—even the experts must take a great deal on trust. That is because they have no choice but to study such phenomena using a cross-disciplinary approach.

These experts greet extraordinary claims made from within their own disciplines with skepticism and even indignation. But they can find it very hard to maintain such methodological vigilance in the hothouse atmosphere of a high-stakes collaboration in which researchers want desperately to believe that their own contributions can have wonderfully synergistic effects when combined with those of experts in other fields. And so, modest researchers recruit one another into immodest funding schemes.

The electronics engineer and the physiologist, the cognitive scientist and the physicist, the economist and the manufacturing specialist—all must take one another's statements on trust. They must trust in the contributions from other disciplines, trust in the power of visions to motivate the cooperation, trust in techniques and instruments that remain somewhat opaque to their users, trust in the trajectories of technical development.

Where trust has become a virtue even for scientists, there is little incentive to challenge outrageous claims or to hold singularitarians accountable. They describe the progressive realization of technical possibility, after all, and their story has a pleasant ring to it. Indeed, there is nothing wrong with the singular simplicity of the singularitarian myth—unless you have something against sloppy reasoning, wishful thinking, and an invitation to irresponsibility. □



EXPERT VIEW: Gordon Bell

WHO HE IS

Principal researcher at Microsoft Research, Silicon Valley. Led the development of or helped design a long list of time-share computers and minicomputers at Digital Equipment Corp., including the PDP-6 and the VAX. A founder of Encore Computer; Ardent Computer; the Computer Museum, in Boston; and the Computer History Museum, in Mountain View, Calif.

SINGULARITY WILL OCCUR

Someday in the distant future

MACHINE CONSCIOUSNESS WILL OCCUR

No

MOORE'S LAW WILL CONTINUE FOR 20 years

THOUGHTS

"Singularity is that point in time when computing is able to know all human and natural-systems knowledge and exceed it in problem-solving capability with the diminished need for humankind as we know it. I basically support the notion, but I have trouble seeing the specific transitions or break points that let the exponential take over and move to the next transition. [If it does,] there'll be a hierarchy of machines versus having a separate race. [But] it is unlikely to happen, because the population will destroy itself before the technology singularity."



By Bill Arnold, *chief scientist, ASML*

using photolithography to define the lattice on which the structures would grow, like vines clinging to a lattice. But these are in a very primitive state today.

Self-replication of superintelligent machines is another pillar of the singularity hypothesis. That kind of reproduction would require the machines to be able to fabricate superadvanced chips because those chips would form the basis of the machines' brains. It's possible that a lithography scanner itself could be the prototype of a self-replicating intelligent robot. I should add here, though, that I have no idea whether truly human-level artificial intelligence is possible, even if the device density and parallelism or the number of computations per second achieved by chips comes to rival those of the human brain.

Researchers will have to meet many, many challenges to keep lithography on the astounding curve it's been on for the past four decades. I have no doubt that they will. Lithography will continue to be the core technology that drives our future. □