

A Response to Bill Joy and the Doom-and-Gloom Technofuturists

John Seely Brown and Paul Duguid

July 2020

Contents

The Tunnel Ahead 3
The Failure of a “6-D” Vision 6

If you lived through the 1950s, you might remember President Eisenhower, orderly suburban housing tracts, backyard bomb shelters—and dreams of a nuclear power plant in every home. Plans for industrial nuclear generators had barely left the drawing board before futurists predicted that every house would have a miniature version. From there, technoenthusiasts predicted the end of power monopolies, the emergence of the “electronic cottage,” the death of the city and the decline of the corporation.

Pessimists and luddites, of course, envisioned nuclear apocalypse. Each side waited for nirvana, or Armageddon, so it could triumphantly tell the other, “I told you so.”

With “Why the Future Doesn’t Need Us” in the April issue of *Wired*, Bill Joy invokes those years gone by. No luddite, Joy is an awe-inspiring technologist—as cofounder and chief scientist of Sun Microsystems, he coauthored, among other things, the Java programming language. So when his article describes a technological juggernaut thundering toward society—bringing with it mutant genes, molecular-level nanotechnology machines and superintelligent robots—all need to listen. Like the nuclear prognosticators, Joy can see the juggernaut clearly. What he can’t see—which is precisely what makes his vision so scary—are any controls.

But it doesn’t follow that the juggernaut is uncontrollable. To understand why not, readers should note the publication in which this article appeared. For the better part of a decade, *Wired* has been a cheerleader for the digital age. Until now, *Wired* has rarely been a venue to which people have looked for a way to put a brake on innovation. Therefore its shift with Joy’s article from cheering to warning marks an important and surprising moment in the digital zeitgeist.

In an effort to locate some controls, let’s go back to the nuclear age. Innovation, the argument went back in the 1950s, would make nuclear power plants smaller and cheaper. They would enter mass production and quickly become available to all.

Even today the argument might appear inescapable until you notice what’s missing: The tight focus of this vision makes it almost impossible to see forces other than technology at work. In the case of nuclear development, a host of forces worked to dismantle the dream of a peaceful atom, including the environmental movement, antinuclear protests, concerned scientists, worried neighbors of Chernobyl and Three Mile Island, government regulators and antiproliferation treaties. Cumulatively, these forces slowed the nuclear juggernaut to a crawl.

Similar social forces are at work on technologies today. But because the digerati, like technoenthusiasts before them, look to the future with technological tunnel vision, they too have trouble bringing other forces into view.

The Tunnel Ahead

In Joy’s vision, as in the nuclear one, there’s a recognizable tunnel vision that leaves people out of the picture and focuses on technology in splendid isolation. This vision

leads not only to doom-and-gloom scenarios, but also to tunnel design: the design of “simple” technologies that are actually difficult to use.

To escape both trite scenarios and bad design, we have to widen our horizons and bring into view not only technological systems, but also social systems. Good designs look beyond the dazzling potential of the technology to social factors, such as the limited patience of most users.

Paying attention to the latter has, for example, allowed the PalmPilot and Nintendo Game Boy to sweep aside more complex rivals. Their elegant simplicity has made them readily usable. And their usability has in turn created an important social support system. The devices are so widely used that anyone having trouble with a Pilot or Game Boy rarely has to look far for advice from a more experienced user.

As this small example suggests, technological and social systems shape each other. The same is true on a larger scale. Technologies— such as gunpowder, the printing press, the railroad, the telegraph and the Internet—can shape society in profound ways. But, on the other hand, social systems—in the form of governments, the courts, formal and informal organizations, social movements, professional networks, local communities, market institutions and so forth—shape, moderate and redirect the raw power of technologies.

Given the crisp edges of technology and the fuzzy outlines of society, it certainly isn’t easy to use these two worldviews simultaneously. But if you want to see where we are going, or design the means to get there, you need to grasp both.

This perspective allows a more sanguine look at Joy’s central concerns: genetic engineering, nanotechnology and robotics. Undoubtedly, each deserves serious thought. But each should be viewed in the context of the social system in which it is inevitably embedded.

Genetic engineering presents the clearest example. Barely a year ago, the technology seemed to be an unstoppable force. Major chemical and agricultural interests were barreling down an open highway. In the past year, however, road conditions changed dramatically for the worse: Cargill faced Third World protests against its patents; Monsanto suspended research on sterile seeds; and champions of genetically modified foods, who once saw an unproblematic and lucrative future, are scurrying to counter consumer boycotts of their products.

Almost certainly, those who support genetic modification will have to look beyond the technology if they want to advance it. They need to address society directly—not just by putting labels on modified foods, but by educating people about the costs and the benefits of these new agricultural products. Having ignored social concerns, however, proponents have made the people they need to educate profoundly suspicious and hostile.

Nanotechnology offers a rather different example of how the future can frighten us. Because the technology involves engineering at a molecular level, both the promise and the threat seem immeasurable. But they are immeasurable for a good reason: The technology is still almost wholly on the drawing board.

Two of nanotechnology's main proponents, Ralph Merkle and Eric Drexler, worked with us at the Xerox Palo Alto Research Center in Palo Alto, Calif. The two built powerful nano-CAD tools and then ran simulations of the resulting molecular-level designs. These experiments showed definitively that nano devices are theoretically feasible. No one, however, has laid out a route from lab-based simulation to practical systems in any detail.

In the absence of a plan, it's important to ask the right questions: Can nanotechnology fulfill its great potential in tasks ranging from data storage to pollution control, all without spiraling out of control? If the lesson of genetic engineering is any guide, planners would do well to consult and educate the public early on, even though useful nano systems are probably decades away.

Worries about robotics appear premature, as well. Internet "bots" that search, communicate and negotiate for their human masters may appear to behave like Homo sapiens, but in fact, bots are often quite inept at functions that humans do well—functions that call for judgment, discretion, initiative or tacit understanding. They are good (and useful) for those tasks that humans do poorly. So they are better thought of as complementary systems, not rivals to humanity. Although bots will undoubtedly get better at what they do, such development will not necessarily make them more human.

Are more conventional clanking robots—the villains of science fiction—any great threat to society? We doubt it. Xerox PARC research on self-aware, reconfigurable "polybots" has pushed the boundaries of what robots can do, pointing the way to "morphing robots" that are able to move and change shape.

Nonetheless, for all their cutting-edge agility, these robots are a long way from making good dance partners. The chattiness of *Star Wars*' C-3PO still lies well beyond real-world machines. Indeed, what talk robots or computers achieve, though it may appear similar, is quite different from human talk. Talking machines travel routes designed specifically to avoid the full complexities of human language.

Robots may seem intelligent, but such intelligence is profoundly hampered by their inability to learn in any significant way. (This failing has apparently led Toyota, after heavy investment in robotics, to consider replacing robots with humans on many production lines.) And without learning, simple common sense will lie beyond robots for a long time to come.

Indeed, despite years of startling advances and innumerable successes like the chess-playing Big Blue, computer science is still about as far as it ever was from building a machine with the learning abilities, linguistic competence, common sense or social skills of a 5-year-old child.

As with Internet bots, real-world robots will no doubt become increasingly useful. But they will probably also become increasingly frustrating to use as a result of tunnel design. In that regard, they may indeed seem antisocial, but not in the way of *Terminator*-like fantasies of robot armies that lay waste to human society.

Indeed, the thing that handicaps robots most is their lack of a social existence. For it is our social existence as humans that shapes how we speak, learn, think and develop common sense. All forms of artificial life (whether bugs or bots) will remain primarily a metaphor for—rather than a threat to—society, at least until they manage to enter a debate, sing in a choir, take a class, survive a committee meeting, join a union, pass a law, engineer a cartel or summon a constitutional convention.

These critical social mechanisms allow society to shape its future. It is through planned, collective action that society forestalls expected consequences (such as Y2K) and responds to unexpected events (such as epidemics).

The Failure of a “6-D” Vision

Why does the threat of a cunning, replicating robot society look so close from one perspective, yet so distant from another? The difference lies in the well-known tendency of futurologists to count “1, 2, 3 . . . a million.” That is, once the first step on a path is taken, it’s very easy to assume that all subsequent steps are trivial.

Several of the steps Joy asks us to take—the leap from genetic engineering to a “white plague”; from simulations to out-of-control nanotechnology; from replicating peptides to a “robot species”—are extremely large. And they are certainly not steps that will be taken without diversions, regulations or controls.

One of the lessons of Joy’s article, then, is that the path to the future can look simple (and sometimes downright terrifying) if you look at it through what we call “6-D lenses.” We coined this phrase having so often in our research hit up against upon such “de-” or “di-” words as demassification, decentralization, disintermediation, despacialization, disaggregation and demarketization in the canon of futurology.

If you take any one of these words in isolation, it’s easy to follow their relentless logic to its evident conclusion. Because firms are getting smaller, for example, it’s easy to assume that companies and other intermediaries are simply disintegrating into markets. And because communication is growing cheaper and more powerful, it’s easy to believe in the “death of distance.”

But things rarely work in such linear fashion. Other forces are often at work, such as those driving firms into larger and larger mergers to take advantage of social, rather than merely technological, networks. Similarly, even though communications technology has killed distance, people curiously can’t stay away from the social hotbed of modern communications technology, Silicon Valley.

Importantly, these d-words indicate that the old ties that once bound communities, organizations and institutions are being picked apart by technologies. A simple, linear reading, then, suggests that these communities, organizations and institutions will now simply fall apart. A more complex reading, taking into account the multiple forces at work, offers another picture.

While many powerful national corporations have grown insignificant, some have transformed into more powerful transnational firms. While some forms of community may be dying, others, bolstered by technology, are growing stronger.

Technology and society are constantly forming and reforming new dynamic equilibriums with far-reaching implications. The challenge for futurology (and for all of us) is to see beyond the hype and past the oversimplifications to the full import of these new sociotechnical formations.

Two hundred years ago, Thomas Malthus, assuming that human society and agricultural technology developed on separate paths, predicted that society was growing so fast that it would starve itself to death, the so-called Malthusian trap.

A hundred years later, H.G. Wells similarly assumed that society and technology were developing independently. Like many people today, Wells saw the advance of technology outstripping the evolution of society, leading him to predict that technology's relentless juggernaut would unfeelingly crush society. Like Joy, both Malthus and Wells issued important warnings, alerting society to the dangers it faced. But by their actions, Malthus and Wells helped prevent the very future they were so certain would come about.

These self-*un*fulfilling prophecies failed to see that, once warned, society could galvanize itself into action. Of course, this social action in the face of threats showed that Malthus and Wells were most at fault in their initial assumption. Social and technological systems do not develop independently; the two evolve together in complex feedback loops, wherein each drives, restrains and accelerates change in the other. Malthus and Wells—and now Joy—are, indeed, critical parts of these complex loops. Each knew when and how to sound the alarm. But each thought little about how to respond to that alarm.

Once the social system is factored back into the equation like this, the road ahead becomes harder to navigate. Ultimately we should be grateful to Joy for saying, at the least, that there could be trouble ahead when so many of his fellow digerati will only tell us complacently that the road is clear.

John Seely Brown is chief scientist of the Xerox Corporation, and director of the Xerox Palo Alto Research Center (PARC). Paul Duguid is a research specialist in the division of Social and Cultural Studies in Education at the University of California, Berkeley, and a consultant at the Xerox PARC. This article reprinted by permission of The Industry Standard; www.thestandard.com, April 13, 2000. Copyright 2000 Standard Media International.

The Ted K Archive

John Seely Brown and Paul Duguid
A Response to Bill Joy and the Doom-and-Gloom Technofuturists
July 2020

Emerging Technologies: Ethics, Law and Governance (pp.65-71).
DOI:10.4324/9781003074960-6. <[https://web.cs.ucdavis.edu/~koehl/Teaching/
ECS188_W16/Reprints/Response_to_BillJoy.pdf](https://web.cs.ucdavis.edu/~koehl/Teaching/ECS188_W16/Reprints/Response_to_BillJoy.pdf)>

www.thetedkarchive.com