PLEASE RETURN
=THANKS!=

What the Internet is doing to our brains

BY NICHOLAS CARR

Illustration by Guy Billout

Is Google Making Us Stupid?

"Dave, stop. Stop, will you? Stop, Dave. Will you stop, Dave?" So the supercomputer HAL pleads with the implacable astronaut Dave Bowman in a famous and weirdly poignant scene toward the end of Stanley Kubrick's *2001: A Space Odyssey*. Bowman, having nearly been sent to a deep-space death by the malfunctioning machine, is calmly, coldly disconnecting the memory circuits that control its artificial »

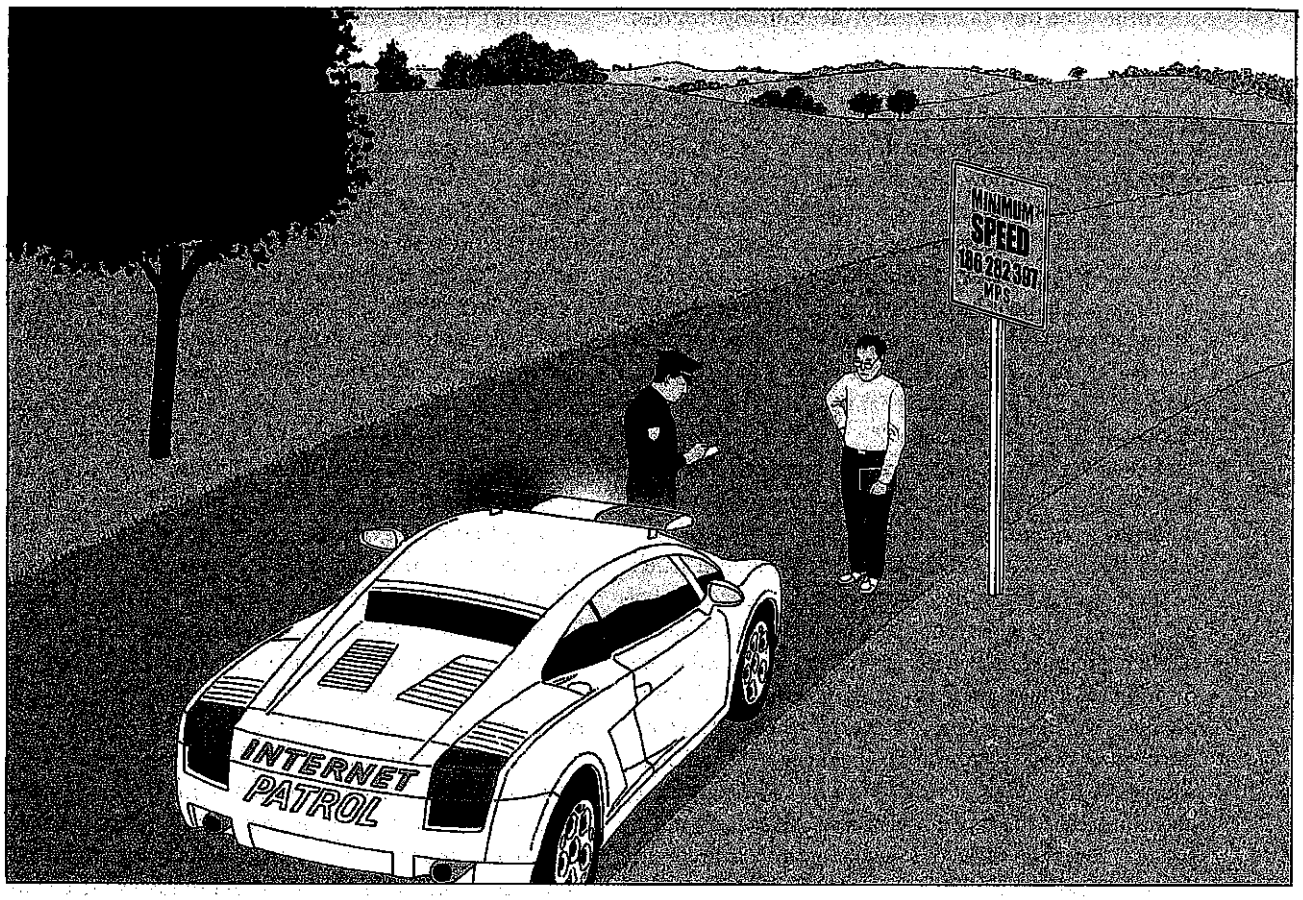

brain. "Dave, my mind is going," HAL says, forlornly. "I can feel it. I can feel it."

I can feel it, too. Over the past few years I've had an uncomfortable sense that someone, or something, has been tinkering with my brain, remapping the neural circuitry, reprogramming the memory. My mind isn't going—so far as I can tell—but it's changing. I'm not thinking the way I used to think. I can feel it most strongly when I'm reading. Immersing myself in a book or a lengthy article used to be easy. My mind would get caught up in the narrative or the turns of the argument, and I'd spend hours strolling through long stretches of prose. That's rarely the case anymore. Now my concentration often starts to drift after two or three pages. I get fidgety, lose the thread, begin looking for something else to do. I feel as if I'm always dragging my wayward brain back to the text. The deep reading that used to come naturally has become a struggle.

I think I know what's going on. For more than a decade now, I've been spending a lot of time online, searching and surfing and sometimes adding to the great databases of the Internet. The Web has been a godsend to me as a writer. Research that once required days in the stacks or periodical rooms of libraries can now be done in minutes. A few Google searches, some quick clicks on hyperlinks, and I've got the telltale fact or pithy quote I was after. Even when I'm not working, I'm as likely as not to be foraging in the Web's info-thickets—reading and writing e-mails, scanning headlines and blog posts, watching videos and listening to podcasts, or just tripping from link to link to link. (Unlike footnotes, to which they're sometimes likened, hyperlinks don't merely point to related works; they propel you toward them.)

For me, as for others, the Net is becoming a universal medium, the conduit for most of the information that flows through my eyes and ears and into my mind. The advantages of having immediate access to such an incredibly rich store of information are many, and they've been widely described and duly applauded. "The perfect recall of silicon memory," *Wired*'s Clive Thompson has written, "can be an enormous boon to thinking." But that boon comes at a price. As the media theorist Marshall McLuhan pointed out in the 1960s, media are not just passive channels of information. They supply the stuff of thought, but they also shape the process of thought. And what the Net seems to be doing is chipping away my capacity for concentration and contemplation. My mind now expects to take in information the way the Net distributes it: in a swiftly moving stream of particles. Once I was a scuba diver in the sea of words. Now I zip along the surface like a guy on a Jet Ski.

I'm not the only one. When I mention my troubles with reading to friends and acquaintances—literary types, most of them—many say they're having similar experiences. The more they use the Web, the more they have to fight to stay focused on long pieces of writing. Some of the bloggers I follow have also begun mentioning the phenomenon. Scott Karp, who writes a blog about online media, recently confessed that he has stopped reading books altogether. "I was a lit major in college, and used to be [a] voracious book reader," he wrote. "What happened?" He speculates on the answer: "What if I do all my reading on the web not so much because the way I read has changed, i.e. I'm just seeking convenience, but because the way I THINK has changed?"

Bruce Friedman, who blogs regularly about the use of computers in medicine, also has described how the Internet has altered his mental habits. "I now have almost totally lost the ability to read and absorb a longish article on the web or in print," he wrote earlier this year. A pathologist who has long been on the faculty of the University of Michigan Medical School, Friedman elaborated on his comment in a telephone conversation with me. His thinking, he said, has taken on a "staccato" quality, reflecting the way he quickly scans short passages of text from many sources online. "I can't read *War and Peace* anymore," he admitted. "I've lost the ability to do that. Even a blog post of more than three or four paragraphs is too much to absorb. I skim it."

Anecdotes alone don't prove much. And we still await the long-term neurological and psychological experiments that will provide a definitive picture of how Internet use affects cognition. But a recently published study of online research habits, conducted by scholars from University College London, suggests that we may well be in the midst of a sea change in the way we read and think. As part of the five-year research program, the scholars examined computer logs documenting the behavior of visitors to two popular research sites, one operated by the British Library and one by a U.K. educational consortium, that provide access to journal articles, e-books, and other sources of written information. They found that people using the sites exhibited "a form of skimming activity," hopping from one source to another and rarely returning to any source they'd already visited. They typically read no more than one or two pages of an article or book before they would "bounce" out to another site. Sometimes they'd save a long article, but there's no evidence that they ever went back and actually read it. The authors of the study report:

> It is clear that users are not reading online in the traditional sense; indeed there are signs that new forms of "reading" are emerging as users "power browse" horizontally through titles, contents pages and abstracts going for quick wins. It almost seems that they go online to avoid reading in the traditional sense.

Thanks to the ubiquity of text on the Internet, not to mention the popularity of text-messaging on cell phones, we may well be reading more today than we did in the 1970s or 1980s, when television was our medium of choice. But it's a different kind of reading, and behind it lies a different kind of thinking—perhaps even a new sense of the self. "We are not only *what* we read," says Maryanne Wolf, a developmental psychologist at Tufts University and the author of *Proust and the Squid: The Story and Science of the Reading Brain*. "We are *how* we read." Wolf worries that the style of reading promoted by the Net, a style that puts "efficiency" and "immediacy" above all else, may be weakening our capacity for the kind of deep reading that emerged when an earlier technology, the printing press, made long and complex works of prose commonplace. When we read online, she says, we tend to become "mere decoders of information." Our ability to interpret text, to make the rich mental connections that form when we read deeply and without distraction, remains largely disengaged.

Reading, explains Wolf, is not an instinctive skill for human beings. It's not etched into our genes the way speech is. We have to teach our minds how to translate the symbolic characters we see into the language we understand. And the media or other technologies we use in learning and practicing the craft of reading play an important part in shaping the neural circuits inside our brains. Experiments demonstrate that readers of ideograms, such as the Chinese, develop a mental circuitry for reading that is very different from the circuitry found in those of us whose written language employs an alphabet. The variations extend across many regions of the brain, including those that govern such essential cognitive functions as memory and the interpretation of visual and auditory stimuli. We can expect as well that the circuits woven by our use of the Net will be different from those woven by our reading of books and other printed works.

Sometime in 1882, Friedrich Nietzsche bought a typewriter—a Malling-Hansen Writing Ball, to be precise. His vision was failing, and keeping his eyes focused on a page had become exhausting and painful, often bringing on crushing headaches. He had been forced to curtail his writing, and he feared that he would soon have to give it up. The typewriter rescued him, at least for a time. Once he had mastered touch-typing, he was able to write with his eyes closed, using only the tips of his fingers. Words could once again flow from his mind to the page.

But the machine had a subtler effect on his work. One of Nietzsche's friends, a composer, noticed a change

in the style of his writing. His already terse prose had become even tighter, more telegraphic. "Perhaps you will through this instrument even take to a new idiom," the friend wrote in a letter, noting that, in his own work, his "'thoughts' in music and language often depend on the quality of pen and paper."

"You are right," Nietzsche replied, "our writing equipment takes part in the forming of our thoughts." Under the sway of the machine, writes the German media scholar Friedrich A. Kittler, Nietzsche's prose "changed from arguments to aphorisms, from thoughts to puns, from rhetoric to telegram style."

The human brain is almost infinitely malleable. People used to think that our mental meshwork, the dense connections formed among the 100 billion or so neurons inside our skulls, was largely fixed by the time we reached adulthood. But brain researchers have discovered that that's not the case. James Olds, a professor of neuroscience who directs the Krasnow Institute for Advanced Study at George Mason University, says that even the adult mind "is very plastic." Nerve cells routinely break old connections and form new ones. "The brain," according to Olds, "has the ability to reprogram itself on the fly, altering the way it functions."

As we use what the sociologist Daniel Bell has called our "intellectual technologies"—the tools that extend our mental rather than our physical capacities—we inevitably begin to take on the qualities of those technologies. The mechanical clock, which came into common use in the 14th century, provides a compelling example. In *Technics and Civilization*, the historian and cultural critic Lewis Mumford described how the clock "disassociated time from human events and helped create the belief in an independent world of mathematically measurable sequences." The "abstract framework of divided time" became "the point of reference for both action and thought."

The clock's methodical ticking helped bring into being the scientific mind and the scientific man. But it also took something away. As the late MIT computer scientist Joseph Weizenbaum observed in his 1976 book, *Computer Power and Human Reason: From Judgment to Calculation*, the conception of the world that emerged from the widespread use of timekeeping instruments "remains an impoverished version of the older one, for it rests on a rejection of those direct experiences that formed the basis for, and indeed constituted, the old reality." In deciding when to eat, to work, to sleep, to rise, we stopped listening to our senses and started obeying the clock.

The process of adapting to new intellectual technologies is reflected in the changing metaphors we use to explain ourselves to ourselves. When the mechanical clock arrived, people began thinking of their brains as operating "like clockwork." Today, in the age of software, we have come to think of them as operating "like computers." But the changes, neuroscience tells us, go much deeper than metaphor. Thanks to our brain's plasticity, the adaptation occurs also at a biological level.

The Internet promises to have particularly far-reaching effects on cognition. In a paper published in 1936, the British mathematician Alan Turing proved that a digital computer, which at the time existed only as a theoretical machine, could be programmed to perform the function of any other information-processing device. And that's what we're seeing today. The Internet, an immeasurably powerful computing system, is subsuming most of our other intellectual technologies. It's becoming our map and our clock, our printing press and our typewriter, our calculator and our telephone, and our radio and TV.

When the Net absorbs a medium, that medium is recreated in the Net's image. It injects the medium's content with hyperlinks, blinking ads, and other digital gewgaws, and it surrounds the content with the content of all the other media it has absorbed. A new e-mail message, for instance, may announce its arrival as we're glancing over the latest headlines at a newspaper's site. The result is to scatter our attention and diffuse our concentration.

The Net's influence doesn't end at the edges of a computer screen, either. As people's minds become attuned to the crazy quilt of Internet media, traditional media have to adapt to the audience's new expectations. Television programs add text crawls and pop-up ads, and magazines and newspapers shorten their articles, introduce capsule summaries, and crowd their pages with easy-to-browse info-snippets. When, in March of this year, *The New York Times* decided to devote the second and third pages of every edition to article abstracts, its design director, Tom Bodkin, explained that the "shortcuts" would give harried readers a quick "taste" of the day's news, sparing them the "less efficient" method of actually turning the pages and reading the articles. Old media have little choice but to play by the new-media rules.

Never has a communications system played so many roles in our lives—or exerted such broad influence over our thoughts—as the Internet does today. Yet, for all that's been written about the Net, there's been little consideration of how, exactly, it's reprogramming us. The Net's intellectual ethic remains obscure.

About the same time that Nietzsche started using his typewriter, an earnest young man named Frederick Winslow Taylor carried a stopwatch into the Midvale Steel plant in Philadelphia and began a historic series of experiments aimed at improving the efficiency of the plant's machinists. With the approval of Midvale's owners, he recruited a group of factory hands, set them to work on various metalworking machines, and recorded and timed their every movement as well as the operations of the machines. By breaking down every job into a

sequence of small, discrete steps and then testing different ways of performing each one, Taylor created a set of precise instructions—an "algorithm," we might say today—for how each worker should work. Midvale's employees grumbled about the strict new regime, claiming that it turned them into little more than automatons, but the factory's productivity soared.

More than a hundred years after the invention of the steam engine, the Industrial Revolution had at last found its philosophy and its philosopher. Taylor's tight industrial choreography—his "system," as he liked to call it—was embraced by manufacturers throughout the country and, in time, around the world. Seeking maximum speed, maximum efficiency, and maximum output, factory owners used time-and-motion studies to organize their work and configure the jobs of their workers. The goal, as Taylor defined it in his celebrated 1911 treatise, *The Principles of Scientific Management*, was to identify and adopt, for every job, the "one best method" of work and thereby to effect "the gradual substitution of science for rule of thumb throughout the mechanic arts." Once his system was applied to all acts of manual labor, Taylor assured his followers, it would bring about a restructuring not only of industry but of society, creating a utopia of perfect efficiency. "In the past the man has been first," he declared; "in the future the system must be first."

Taylor's system is still very much with us; it remains the ethic of industrial manufacturing. And now, thanks to the growing power that computer engineers and software coders wield over our intellectual lives, Taylor's ethic is beginning to govern the realm of the mind as well. The Internet is a machine designed for the efficient and automated collection, transmission, and manipulation of information, and its legions of programmers are intent on finding the "one best method"—the perfect algorithm—to carry out every mental movement of what we've come to describe as "knowledge work."

Google's headquarters, in Mountain View, California—the Googleplex—is the Internet's high church, and the religion practiced inside its walls is Taylorism. Google, says its chief executive, Eric Schmidt, is "a company that's founded around the science of measurement," and it is striving to "systematize everything" it does. Drawing on the terabytes of behavioral data it collects through its search engine and other sites, it carries out thousands of experiments a day, according to the *Harvard Business Review*, and it uses the results to refine the algorithms that increasingly control how people find information and extract meaning from it. What Taylor did for the work of the hand, Google is doing for the work of the mind.

The company has declared that its mission is "to organize the world's information and make it universally accessible and useful." It seeks to develop "the perfect search engine," which it defines as something that "understands exactly what you mean and gives you back exactly what you want." In Google's view, information is a kind of commodity, a utilitarian resource that can be mined and processed with industrial efficiency. The more pieces of information we can "access" and the faster we can extract their gist, the more productive we become as thinkers.

Where does it end? Sergey Brin and Larry Page, the gifted young men who founded Google while pursuing doctoral degrees in computer science at Stanford, speak frequently of their desire to turn their search engine into an artificial intelligence, a HAL-like machine that might be connected directly to our brains. "The ultimate search engine is something as smart as people—or smarter," Page said in a speech a few years back. "For us, working on search is a way to work on artificial intelligence." In a 2004 interview with *Newsweek*, Brin said, "Certainly if you had all the world's information directly attached to your brain, or an artificial brain that was smarter than your brain, you'd be better off." Last year, Page told a convention of scientists that Google is "really trying to build artificial intelligence and to do it on a large scale."

The media or other technologies we use in learning the craft of reading play an important part in shaping the neural circuits inside our brains.

Such an ambition is a natural one, even an admirable one, for a pair of math whizzes with vast quantities of cash at their disposal and a small army of computer scientists in their employ. A fundamentally scientific enterprise, Google is motivated by a desire to use technology, in Eric Schmidt's words, "to solve problems that have never been solved before," and artificial intelligence is the hardest problem out there. Why wouldn't Brin and Page want to be the ones to crack it?

Still, their easy assumption that we'd all "be better off" if our brains were supplemented, or even replaced, by an artificial intelligence is unsettling. It suggests a belief that intelligence is the output of a mechanical process, a series of discrete steps that can be isolated, measured, and

optimized. In Google's world, the world we enter when we go online, there's little place for the fuzziness of contemplation. Ambiguity is not an opening for insight but a bug to be fixed. The human brain is just an outdated computer that needs a faster processor and a bigger hard drive.

The idea that our minds should operate as high-speed data-processing machines is not only built into the workings of the Internet, it is the network's reigning business model as well. The faster we surf across the Web—the more links we click and pages we view—the more opportunities Google and other companies gain to collect information about us and to feed us advertisements. Most of the proprietors of the commercial Internet have a financial stake in collecting the crumbs of data we leave behind as we flit from link to link—the more crumbs, the better. The last thing these companies want is to encourage leisurely reading or slow, concentrated thought. It's in their economic interest to drive us to distraction.

Maybe I'm just a worrywart. Just as there's a tendency to glorify technological progress, there's a countertendency to expect the worst of every new tool or machine. In Plato's *Phaedrus*, Socrates bemoaned the development of writing. He feared that, as people came to rely on the written word as a substitute for the knowledge they used to carry inside their heads, they would, in the words of one of the dialogue's characters, "cease to exercise their memory and become forgetful." And because they would be able to "receive a quantity of information without proper instruction," they would "be thought very knowledgeable when they are for the most part quite ignorant." They would be "filled with the conceit of wisdom instead of real wisdom." Socrates wasn't wrong—the new technology did often have the effects he feared—but he was shortsighted. He couldn't foresee the many ways that writing and reading would serve to spread information, spur fresh ideas, and expand human knowledge (if not wisdom).

The arrival of Gutenberg's printing press, in the 15th century, set off another round of teeth gnashing. The Italian humanist Hieronimo Squarciafico worried that the easy availability of books would lead to intellectual laziness, making men "less studious" and weakening their minds. Others argued that cheaply printed books and broadsheets would undermine religious authority, demean the work of scholars and scribes, and spread sedition and debauchery. As New York University professor Clay Shirky notes, "Most of the arguments made against the printing press were correct, even prescient." But, again, the doomsayers were unable to imagine the myriad blessings that the printed word would deliver.

So, yes, you should be skeptical of my skepticism. Perhaps those who dismiss critics of the Internet as Luddites or nostalgists will be proved correct, and from our hyperactive, data-stoked minds will spring a golden age of intellectual discovery and universal wisdom. Then again, the Net isn't the alphabet, and although it may replace the printing press, it produces something altogether different. The kind of deep reading that a sequence of printed pages promotes is valuable not just for the knowledge we acquire from the author's words but for the intellectual vibrations those words set off within our own minds. In the quiet spaces opened up by the sustained, undistracted reading of a book, or by any other act of contemplation, for that matter, we make our own associations, draw our own inferences and analogies, foster our own ideas. Deep reading, as Maryanne Wolf argues, is indistinguishable from deep thinking.

If we lose those quiet spaces, or fill them up with "content," we will sacrifice something important not only in our selves but in our culture. In a recent essay, the playwright Richard Foreman eloquently described what's at stake:

> I come from a tradition of Western culture, in which the ideal (my ideal) was the complex, dense and "cathedral-like" structure of the highly educated and articulate personality—a man or woman who carried inside themselves a personally constructed and unique version of the entire heritage of the West. [But now] I see within us all (myself included) the replacement of complex inner density with a new kind of self—evolving under the pressure of information overload and the technology of the "instantly available."

As we are drained of our "inner repertory of dense cultural inheritance," Foreman concluded, we risk turning into "'pancake people'—spread wide and thin as we connect with that vast network of information accessed by the mere touch of a button."

I'm haunted by that scene in *2001*. What makes it so poignant, and so weird, is the computer's emotional response to the disassembly of its mind: its despair as one circuit after another goes dark, its childlike pleading with the astronaut—"I can feel it. I can feel it. I'm afraid"—and its final reversion to what can only be called a state of innocence. HAL's outpouring of feeling contrasts with the emotionlessness that characterizes the human figures in the film, who go about their business with an almost robotic efficiency. Their thoughts and actions feel scripted, as if they're following the steps of an algorithm. In the world of *2001*, people have become so machinelike that the most human character turns out to be a machine. That's the essence of Kubrick's dark prophecy: as we come to rely on computers to mediate our understanding of the world, it is our own intelligence that flattens into artificial intelligence.

Nicholas Carr's most recent book, *The Big Switch: Rewiring the World, From Edison to Google*, was published earlier this year.

PANDEMICS. GLOBAL WARMING. *Food shortages. No more fossil fuels. What are humans to do? The same thing the species has done before: evolve to meet the challenge. But this time we don't have to rely on natural evolution to make us smart enough to survive. We can do it ourselves, right now, by harnessing technology and pharmacology to boost our intelligence. Is Google actually making us smarter?*

TECHNOLOGY

GET SMART

By Jamais Cascio

SEVENTY-FOUR THOUSAND YEARS ago, humanity nearly went extinct. A super-volcano at what's now Lake Toba, in Sumatra, erupted with a strength more than a thousand times that of Mount St. Helens in 1980. Some 800 cubic kilometers of ash filled the skies of the Northern Hemisphere, lowering global temperatures and pushing a climate already on the verge of an ice age over the edge. Some scientists speculate that as the Earth went into a deep freeze, the population of *Homo sapiens* may have dropped to as low as a few thousand families.

The Mount Toba incident, although unprecedented in magnitude, was part of a broad pattern. For a period of 2 million years, ending with the last ice age around 10,000 B.C., the Earth experienced a series of convulsive glacial events. This rapid-fire climate change meant that humans couldn't rely on consistent patterns to know which animals to hunt, which plants to gather, or even which predators might be waiting around the corner.

How did we cope? By getting smarter. The neurophysiologist William Calvin argues persuasively that modern human cognition—including sophisticated language and the capacity to plan ahead—evolved in response to the demands of this long age of turbulence. According to Calvin, the reason we survived is that our brains changed to meet the challenge: we transformed the ability to target a moving animal with a thrown rock into a capability for foresight and long-term planning. In the process, we may have developed syntax and formal structure from our simple language.

Our present century may not be quite as perilous for the human race as an ice age in the aftermath of a super-volcano eruption, but the next few decades will pose enormous hurdles that go beyond the climate crisis. The end of the fossil-fuel era, the fragility of the global food web, growing population density, and the spread of pandemics, as well as the emergence of radically transformative bio- and nanotechnologies—each of these threatens us with broad disruption or even devastation. And as good as our brains have become at planning ahead, we're still biased toward looking for near-term, simple threats. Subtle, long-term risks, particularly those involving complex, global processes, remain devilishly hard for us to manage.

But here's an optimistic scenario for you: if the next several decades are as bad as some of us fear they could be, we can respond, and survive, the way our species has done time and again: by getting smarter. But this time, we don't have to rely solely on natural evolutionary processes to boost our intelligence. We can do it ourselves.

Most people don't realize that this process is already under way. In fact, it's happening all around us, across the full spectrum of how we understand intelligence. It's visible in the hive mind of the Internet, in the powerful tools for simulation and visualization that are jump-starting new scientific disciplines, and in the development of drugs that some people (myself included) have discovered let them study harder, focus better, and stay awake longer with full clarity. So far, these augmentations have largely been outside of our bodies, but they're very much part of who we are

ANASTASIA VASILAKIS

today: they're physically separate from us, but we and they are becoming cognitively inseparable. And advances over the next few decades, driven by breakthroughs in genetic engineering and artificial intelligence, will make today's technologies seem primitive. The nascent jargon of the field describes this as "intelligence augmentation." I prefer to think of it as "You+."

Scientists refer to the 12,000 years or so since the last ice age as the Holocene epoch. It encompasses the rise of human civilization and our co-evolution with tools and technologies that allow us to grapple with our physical environment. But if intelligence augmentation has the kind of impact I expect, we may soon have to start thinking of ourselves as living in an entirely new era. The focus of our technological evolution would be less on how we manage and adapt to our physical world, and more on how we manage and adapt to the immense amount of knowledge we've created. We can call it the Nöocene epoch, from Pierre Teilhard de Chardin's concept of the Nöosphere, a collective consciousness created by the deepening interaction of human minds. As that epoch draws closer, the world is becoming a very different place.

OF COURSE, WE'VE been augmenting our ability to think for millennia. When we developed written language, we significantly increased our functional memory and our ability to share insights and knowledge across time and space. The same thing happened with the invention of the printing press, the telegraph, and the radio. The rise of urbanization allowed a fraction of the populace to focus on more-cerebral tasks—a fraction that grew inexorably as more-complex economic and social practices demanded more knowledge work, and industrial technology reduced the demand for manual labor. And caffeine and nicotine, of course, are both classic cognitive-enhancement drugs, primitive though they may be.

With every technological step forward, though, has come anxiety about the possibility that technology harms our natural ability to think. These anxieties were given eloquent expression in these pages by Nicholas Carr, whose essay "Is Google Making Us Stupid?" (July/August 2008 *Atlantic*) argued that the information-dense, hyperlink-rich, spastically churning Internet medium is effectively rewiring our brains, making it harder for us to engage in deep, relaxed contemplation.

Carr's fears about the impact of wall-to-wall connectivity on the human intellect echo cyber-theorist Linda Stone's description of "continuous partial attention," the modern phenomenon of having multiple activities and connections under way simultaneously. We're becoming so accustomed to interruption that we're starting to find focusing difficult, even when we've achieved a bit of quiet. It's an induced form of ADD—a "continuous partial attention-deficit disorder," if you will.

There's also just more information out there—because unlike with previous information media, with the Internet, creating material is nearly as easy as consuming it. And it's easy to mistake more voices for more noise. In reality, though, the proliferation of diverse voices may actually improve our overall ability to think. In *Everything Bad Is Good for You*, Steven Johnson argues that the increasing complexity and range of media we engage with have, over the past century, made us smarter, rather than dumber, by providing a form of cognitive calisthenics. Even pulp-television shows and video games have become extraordinarily dense with detail, filled with subtle references to broader subjects, and more open to interactive engagement. They reward the capacity to make connections and to see patterns—precisely the kinds of skills we need for managing an information glut.

Scientists describe these skills as our "fluid intelligence"—the ability to find meaning in confusion and to solve new problems, independent of acquired knowledge. Fluid intelligence doesn't look much like the capacity to memorize and recite facts, the skills that people have traditionally associated with brainpower. But building it up may improve the capacity to think deeply that Carr and others fear we're losing for good. And we shouldn't let the stresses associated with a transition to a new era blind us to that era's astonishing potential. We swim in an ocean of data, accessible from nearly anywhere, generated by billions of devices. We're only beginning to explore what we can do with this knowledge-at-a-touch.

Moreover, the technology-induced ADD that's associated with this new world may be a short-term problem. The trouble isn't that we have too much information at our fingertips, but that our tools for managing it are still in their infancy. Worries about "information overload" predate the rise of the Web (Alvin Toffler coined the phrase in 1970), and many of the technologies that Carr worries about were developed precisely to help us get some control over a flood of data and ideas. Google isn't the problem; it's the beginning of a solution.

In any case, there's no going back. The information sea isn't going to dry up, and relying on cognitive habits evolved and perfected in an era of limited information flow—and limited information access—is futile. Strengthening our fluid intelligence is the only viable approach to navigating the age of constant connectivity.

WHEN PEOPLE HEAR the phrase *intelligence augmentation*, they tend to envision people with computer chips plugged into their brains, or a genetically engineered race of post-human super-geniuses. Neither of these visions is likely to be realized, for reasons familiar to any Best Buy shopper. In a world of ongoing technological acceleration, today's cutting-edge brain implant would be tomorrow's obsolete junk—and good luck if the protocols change or you're on the wrong side of a "format war" (anyone want a Betamax implant?). And then there's the question of stability: Would you want a chip in your head made by the same folks that made your cell phone, or your PC?

Likewise, the safe modification of human genetics is still years away. And even after genetic modification of adult neurobiology becomes possible, the science will remain in flux; our understanding of how augmentation works, and what kinds of genetic modifications are possible, would still change rapidly. As with digital implants, the brain modification you might undergo one week could become obsolete the next. Who would want a 2025-vintage brain when you're competing against hotshots with Model 2026?

GALLERY Power *by Stephen Doyle*

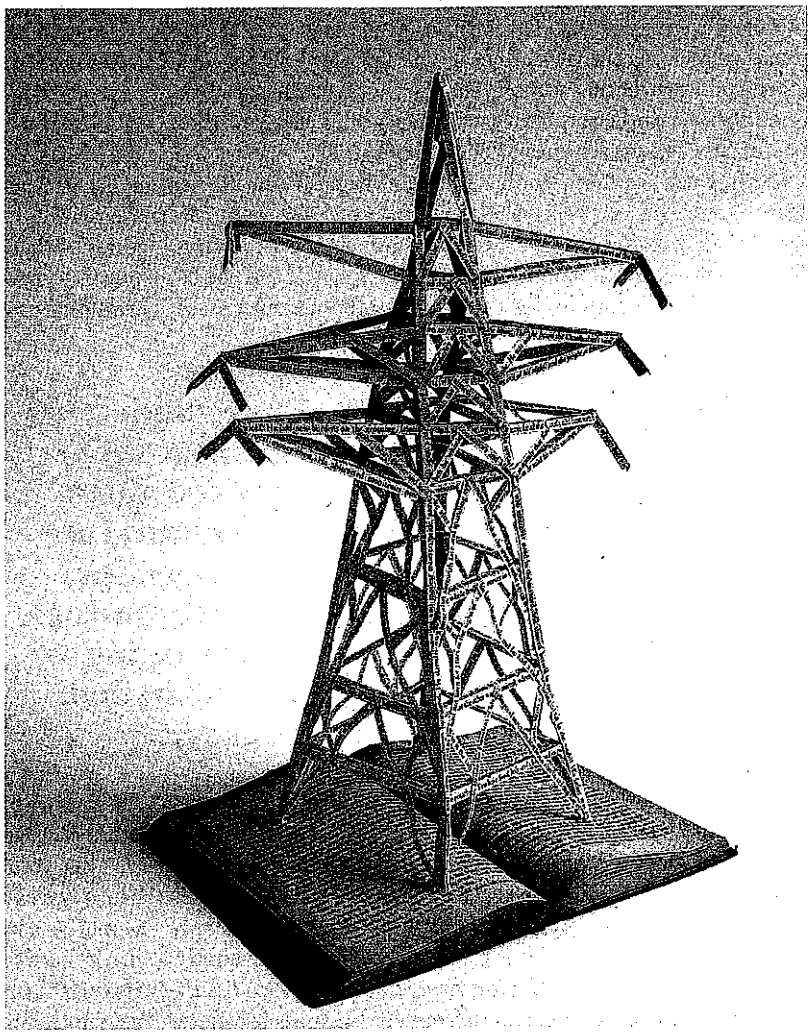

Yet in one sense, the age of the cyborg and the super-genius has already arrived. It just involves external information and communication devices instead of implants and genetic modification. The bioethicist James Hughes of Trinity College refers to all of this as "exocortical technology," but you can just think of it as "stuff you already own." Increasingly, we buttress our cognitive functions with our computing systems, no matter that the connections are mediated by simple typing and pointing. These tools enable our brains to do things that would once have been almost unimaginable:

- powerful simulations and massive data sets allow physicists to visualize, understand, and debate models of an 11-dimension universe;
- real-time data from satellites, global environmental databases, and high-resolution models allow geophysicists to recognize the subtle signs of long-term changes to the planet;
- cross-connected scheduling systems allow anyone to assemble, with a few clicks, a complex, multimodal travel itinerary that would have taken a human travel agent days to create.

If that last example sounds prosaic, it simply reflects how embedded these kinds of augmentation have become. Not much more than a decade ago, such a tool was outrageously impressive—and it destroyed the travel-agent industry.

That industry won't be the last one to go. Any occupation requiring pattern-matching and the ability to find obscure connections will quickly morph from the domain of experts to that of ordinary people whose intelligence has been augmented by cheap digital tools. Humans won't be taken out of the loop—in fact, many, many *more* humans will have the capacity to do something that was once limited to a hermetic priesthood. Intelligence augmentation decreases the need for specialization and increases participatory complexity.

As the digital systems we rely upon become faster, more sophisticated, and (with the usual hiccups) more capable, we're becoming more sophisticated and capable too. It's a form of co-evolution: we learn to adapt our thinking and expectations to these digital systems, even as the system designs become more complex and powerful to meet more of our needs—and eventually come to adapt to *us*.

Consider the Twitter phenomenon, which went from nearly invisible to nearly ubiquitous (at least among the online crowd) in early 2007. During busy periods, the user can easily be overwhelmed by the volume of incoming messages, most of which are of only passing interest. But there is a tiny minority of truly valuable posts. (Sometimes they have extreme value, as they did during the October 2007 wildfires in California and the November 2008 terrorist attacks in Mumbai.) At present, however, finding the most-useful bits requires wading through messages like "My kitty sneezed!" and "I hate this taco!"

But imagine if social tools like Twitter had a way to learn what kinds of messages you pay attention to, and which ones you discard. Over time, the messages that you don't really care about might start to fade in the display, while the ones that you do want to see could get brighter. Such attention filters—or focus assistants—are likely to become important parts of how we handle our daily lives. We'll move from a world of "continuous partial attention" to one we might call "continuous augmented awareness."

As processor power increases, tools like Twitter may be able to draw on the complex simulations and massive data sets that have unleashed a revolution in science. They could become individualized systems that augment our capacity for planning and foresight, letting us play "what-if" with our life choices: where to live, what to study, maybe even where to go for dinner. Initially crude and clumsy, such a system would get better with more data and more experience; just

as important, we'd get better at asking questions. These systems, perhaps linked to the cameras and microphones in our mobile devices, would eventually be able to pay attention to what we're doing, and to our habits and language quirks, and learn to interpret our sometimes ambiguous desires. With enough time and complexity, they would be able to make useful suggestions without explicit prompting.

And such systems won't be working for us alone. Intelligence has a strong social component; for example, we already provide crude cooperative information-filtering for each other. In time, our interactions through the use of such intimate technologies could dovetail with our use of collaborative knowledge systems (such as Wikipedia), to help us not just to build better data sets, but to filter them with greater precision. As our capacity to provide that filter gets faster and richer, it increasingly becomes something akin to collaborative intuition—in which everyone is effectively augmenting everyone else.

The trouble isn't that we have too much information at our fingertips, but that our tools for managing it are still in their infancy.

IN PHARMACOLOGY, TOO, the future is already here. One of the most prominent examples is a drug called modafinil. Developed in the 1970s, modafinil—sold in the U.S. under the brand name Provigil—appeared on the cultural radar in the late 1990s, when the American military began to test it for long-haul pilots. Extended use of modafinil can keep a person awake and alert for well over 32 hours on end, with only a full night's sleep required to get back to a normal schedule.

While it is FDA-approved only for a few sleep disorders, like narcolepsy and sleep apnea, doctors increasingly prescribe it to those suffering from depression, to "shift workers" fighting fatigue, and to frequent business travelers dealing with time-zone shifts. I'm part of the latter group: like more and more professionals, I have a prescription for modafinil in order to help me overcome jet lag when I travel internationally. When I started taking the drug, I expected it to keep me awake; I didn't expect it to make me feel smarter, but that's exactly what happened. The change was subtle but clear, once I recognized it: within an hour of taking a standard 200-mg tablet, I was much more alert, and thinking with considerably more clarity and focus than usual. This isn't just a subjective conclusion. A University of Cambridge study, published in 2003, concluded that modafinil confers a measurable cognitive-enhancement effect across a variety of mental tasks, including pattern recognition and spatial planning, and sharpens focus and alertness.

I'm not the only one who has taken advantage of this effect. The Silicon Valley insider webzine *Tech Crunch* reported in July 2008 that some entrepreneurs now see modafinil as an important competitive tool. The tone of the piece was judgmental, but the implication was clear: everybody's doing it, and if you're not, you're probably falling behind.

This is one way a world of intelligence augmentation emerges. Little by little, people who don't know about drugs like modafinil or don't want to use them will face stiffer competition from the people who do. From the perspective of a culture immersed in athletic doping wars, the use of such drugs may seem like cheating. From the perspective of those who find that they're much more productive using this form of enhancement, it's no more cheating than getting a faster computer or a better education.

Modafinil isn't the only example; on college campuses, the use of ADD drugs (such as Ritalin and Adderall) as study aids has become almost ubiquitous. But these enhancements are primitive. As the science improves, we could see other kinds of cognitive-modification drugs that boost recall, brain plasticity, even empathy and emotional intelligence. They would start as therapeutic treatments, but end up being used to make us "better than normal." Eventually, some of these may become over-the-counter products at your local pharmacy, or in the juice and snack aisles at the supermarket. Spam e-mail would be full of offers to make your brain bigger, and your idea production more powerful.

Such a future would bear little resemblance to *Brave New World* or similar narcomantic nightmares; we may fear the idea of a population kept doped and placated, but we're more likely to see a populace stuck in overdrive, searching out the last bits of competitive advantage, business insight, and radical innovation. No small amount of that innovation would be directed toward inventing the next, more powerful cognitive-enhancement technology.

This would be a different kind of nightmare, perhaps, and cause waves of moral panic and legislative restriction. Safety would be a huge issue. But as we've found with athletic doping, if there's a technique for beating out rivals (no matter how risky), shutting it down is nearly impossible. This would be yet another pharmacological arms race—and in this case, the competitors on one side would just keep getting smarter.

THE MOST RADICAL form of superhuman intelligence, of course, wouldn't be a mind augmented by drugs or exocortical technology; it would be a mind that isn't human at all. Here we move from the realm of extrapolation to the realm of speculation, since solid predictions about artificial intelligence are notoriously hard: our understanding of how the brain creates the mind remains far from good enough to tell us how to construct a mind in a machine.

But while the concept remains controversial, I see no good argument for why a mind running on a machine platform instead of a biological platform will forever be impossible; whether one might appear in five years or 50 or 500, however, is uncertain. I lean toward 50, myself. That's enough time to develop computing hardware able to run a high-speed neural network as sophisticated as that of a human brain, and enough time for the kids who will have grown up surrounded by virtual-world software and household robots—that is, the people who see this stuff not as "Technology," but as everyday tools—to come to dominate the field.

Many proponents of developing an artificial mind are sure that such a breakthrough will be the biggest change in human history. They believe that a machine mind would soon

modify itself to get smarter—and with its new intelligence, then figure out how to make itself smarter still. They refer to this intelligence explosion as "the Singularity," a term applied by the computer scientist and science-fiction author Vernor Vinge. "Within thirty years, we will have the technological means to create superhuman intelligence," Vinge wrote in 1993. "Shortly after, the human era will be ended." The Singularity concept is a secular echo of Teilhard de Chardin's "Omega Point," the culmination of the Nöosphere at the end of history. Many believers in Singularity—which one wag has dubbed "the Rapture for nerds"—think that building the first real AI will be the last thing humans do. Some imagine this moment with terror, others with a bit of glee.

My own suspicion is that a stand-alone artificial mind will be more a tool of narrow utility than something especially apocalyptic. I don't think the theory of an explosively self-improving AI is convincing—it's based on too many assumptions about behavior and the nature of the mind. Moreover, AI researchers, after years of talking about this prospect, are already ultra-conscious of the risk of runaway systems.

More important, though, is that the same advances in processor and process that would produce a machine mind would also increase the power of our own cognitive-enhancement technologies. As intelligence augmentation allows us to make *ourselves* smarter, and then smarter still, AI may turn out to be just a sideshow: we could always be a step ahead.

So what's life like in a world of brain doping, intuition networks, and the occasional artificial mind?

Banal.

Not from our present perspective, of course. For us, now, looking a generation ahead might seem surreal and dizzying. But remember: people living in, say, 2030 will have lived every moment from now until then—we won't jump into the future. For someone going from 2009 to 2030 day by day, most of these changes wouldn't be jarring; instead, they'd be incremental, almost overdetermined, and the occasional surprises would quickly blend into the flow of inevitability.

By 2030, then, we'll likely have grown accustomed to (and perhaps even complacent about) a world where sophisticated foresight, detailed analysis and insight, and augmented awareness are commonplace. We'll have developed a better capacity to manage both partial attention and laser-like focus, and be able to slip between the two with ease—perhaps by popping the right pill, or eating the right snack. Sometimes, our augmentation assistants will handle basic interactions on our behalf; that's okay, though, because we'll increasingly see those assistants as extensions of ourselves.

The amount of data we'll have at our fingertips will be staggering, but we'll finally have gotten over the notion that accumulated information alone is a hallmark of intelligence. The power of all of this knowledge will come from its ability to inform difficult decisions, and to support complex analysis. Most professions will likely use simulation and modeling in their day-to-day work, from political decisions to hairstyle options. In a world of augmented intelligence, we will have a far greater appreciation of the consequences of our actions.

This doesn't mean we'll all come to the same conclusions. We'll still clash with each other's emotions, desires, and beliefs. If anything, our arguments will be more intense, buttressed not just by strongly held opinions but by intricate reasoning. People in 2030 will look back aghast at how ridiculously unsubtle the political and cultural disputes of our present were, just as we might today snicker at simplistic advertising from a generation ago.

Conversely, the debates of the 2030s would be remarkable for us to behold. Nuance and multiple layers will characterize even casual disputes; our digital assistants will be there to catch any references we might miss. And all of this will be everyday, banal reality. Today, it sounds mind-boggling; by then, it won't even merit comment.

What happens if such a complex system collapses? Disaster, of course. But don't forget that we already depend upon enormously complex systems that we no longer even think of as technological. Urbanization, agriculture, and trade were at one time huge innovations. Their collapse (and all of them are now at risk, in different ways, as we have seen in recent months) would be an even greater catastrophe than the collapse of our growing webs of interconnected intelligence.

A less apocalyptic but more likely danger derives from the observation made by the science-fiction author William Gibson: "The future is already here, it's just unevenly distributed." The rich, whether nations or individuals, will inevitably gain access to many augmentations before anyone else. We know from history, though, that a world of limited access wouldn't last forever, even as the technology improved: those who sought to impose limits would eventually face angry opponents with newer, better systems.

Even as competition provides access to these kinds of technologies, though, development paths won't be identical. Some societies may be especially welcoming to biotech boosts; others may prefer to use digital tools. Some may readily adopt collaborative approaches; others may focus on individual enhancement. And around the world, many societies will reject the use of intelligence-enhancement technology entirely, or adopt a cautious wait-and-see posture.

The bad news is that these divergent paths may exacerbate cultural divides created by already divergent languages and beliefs. National rivalries often emphasize cultural differences, but for now we're all still standard human beings. What happens when different groups quite literally think in very, very different ways?

The good news, though, is that this diversity of thought can also be a strength. Coping with the various world-historical dangers we face will require the greatest possible insight, creativity, and innovation. Our ability to build the future that we want—not just a future we can survive—depends on our capacity to understand the complex relationships of the world's systems, to take advantage of the diversity of knowledge and experience our civilization embodies, and to fully appreciate the implications of our choices. Such an ability is increasingly within our grasp. The Nöocene awaits. ■

Jamais Cascio is an affiliate at the Institute for the Future and a senior fellow at the Institute for Ethics and Emerging Technologies.

Technology, Learning, and Mental Development

These two articles were published a year apart in *The Atlantic*. No, the editors did not have a change of heart. Rather, this periodical is known for its thought-provoking pieces exploring different sides of contemporary issues. As you read, contemplate the strengths and weaknesses of our society's growing reliance upon technology.

"Is Google Making Us Stupid?" –Nicholas Carr

1. According to Carr and others, how has the Internet changed our capacity to think?
2. What does developmental psychologist Maryanne Wolf mean when she asserts that Internet reveals, "We are not only *what* we read, we are *how* we read"?
3. What might the following anecdotes reveal about how we learn and think?
   -Chinese ideograms
   -Nietzsche's typewriter
   -the mechanical clock
4. How have various media adapted to the influence of the Internet?
5. What connections does the author attempt to make between the mind and Taylor's "scientific management"?
6. What does Google assert as its mission and larger purpose? (aside from making money)
7. According to Carr, why do companies like Google actively discourage methodical, concentrated use of their search engines?
8. Why does Carr include details from Plato's *Phaedrus* and the printing press in his closing remarks? What is his larger message?
9. What does Foreman mean when he explains his fear that we may become "pancake people"?

"Get Smart"- Jamais Cascio

10. Why does this article open with a description of Mount Toba's eruption and its aftermath?
11. How does Cascio counter Carr's (from the previous article) claims that the Internet has reduced our ability to think and concentrate?
12. Evaluate (describe possible benefits and drawbacks) Cascio's assertion that drugs should be considered as a means for cognitive development.
13. In terms of mental development and issue analysis, what sort of future does Cascio foresee?
14. Cascio concedes at least two major concerns: the possibility of system collapse and uneven distribution. Describe these concerns and his reaction to them.

Your thoughts:

Obviously, these articles just scratch the surface in the debate about technology and learning. The long-term impact, good or bad, may never be fully known. What are your thoughts on these issues? What questions, guidelines, and/or concerns should be considered?