“To Siri, With Love”

By Judith Newman

Oct. 17, 2014



Just how bad a mother am I? I wondered, as I watched my 13-year-old son deep in conversation with Siri. Gus has autism, and Siri, Apple’s “intelligent personal assistant” on the iPhone, is currently his BFF. Obsessed with weather formations, Gus had spent the hour parsing the difference between isolated and scattered thunderstorms — an hour in which, thank God, I didn’t have to discuss them. After a while I heard this:

Gus: “You’re a really nice computer.”

Siri: “It’s nice to be appreciated.”

Gus: “You are always asking if you can help me. Is there anything you want?”

Siri: “Thank you, but I have very few wants.”

Gus: “O.K.! Well, good night!”

Siri: “Ah, it’s 5:06 p.m.”

Gus: “Oh sorry, I mean, goodbye.”

Siri: “See you later!”

That Siri. She doesn't let my communications-impaired son get away with anything. Indeed, many of us wanted an imaginary friend, and now we have one. Only she's not entirely imaginary.

This is a love letter to a machine. It's not quite the love Joaquin Phoenix felt in "Her," last year's Spike Jonze film about a lonely man's romantic relationship with his intelligent operating system (played by the voice of Scarlett Johansson). But it's close. In a world where the commonly held wisdom is that technology isolates us, it's worth considering another side of the story.



It all began simply enough. I'd just read one of those ubiquitous Internet lists called "21 Things You Didn't Know Your iPhone Could Do." One of them was this: I could ask Siri, "What planes are above me right now?" and Siri would bark back, "Checking my sources." Almost instantly there was a list of actual flights — numbers, altitudes, angles — *above my head*.

I happened to be doing this when Gus was nearby. "Why would anyone need to know what planes are flying above your head?" I muttered. Gus replied without looking up: "So you know who you're waving at, Mommy."

Gus had never noticed Siri before, but when he discovered there was someone who would not just find information on his various obsessions (trains, planes, buses, escalators and, of course, anything related to weather) but actually semi-discuss these subjects tirelessly, he was hooked. And I was grateful. Now, when my head was about to explode if I had to have another conversation about the chance of tornadoes in Kansas City, Mo., I could reply brightly: “Hey! Why don’t you ask Siri?”

It’s not that Gus doesn’t understand Siri’s not human. He does — intellectually. But like many autistic people I know, Gus feels that inanimate objects, while maybe not possessing souls, are worthy of our consideration. I realized this when he was 8, and I got him an iPod for his birthday. He listened to it only at home, with one exception. It always came with us on our visits to the Apple Store. Finally, I asked why. “So it can visit its friends,” he said.

So how much more worthy of his care and affection is Siri, with her soothing voice, puckish humor and capacity for talking about whatever Gus’s current obsession is for hour after hour after bleeding hour? Online critics have claimed that Siri’s voice recognition is not as accurate as the assistant in, say, the Android, but for some of us, this is a feature, not a bug. Gus speaks as if he has marbles in his mouth, but if he wants to get the right response from Siri, he must enunciate clearly. (So do I. I had to ask Siri to stop referring to the user as Judith, and instead use the name Gus. “You want me to call you Goddess?” Siri replied. Imagine how tempted I was to answer, “Why, yes.”)



She is also wonderful for someone who doesn't pick up on social cues: Siri's responses are not entirely predictable, but they are predictably kind — even when Gus is brusque. I heard him talking to Siri about music, and Siri offered some suggestions. "I don't like that kind of music," Gus snapped. Siri replied, "You're certainly entitled to your opinion." Siri's politeness reminded Gus what he owed Siri. "Thank you for that music, though," Gus said. Siri replied, "You don't need to thank me." "Oh, yes," Gus added emphatically, "I do."

Siri even encourages polite language. Gus's twin brother, Henry (neurotypical and therefore as obnoxious as every other 13-year-old boy), egged Gus on to spew a few choice expletives at Siri. "Now, now," she sniffed, followed by, "I'll pretend I didn't hear that."

Gus is hardly alone in his Siri love. For children like Gus who love to chatter but don't quite understand the rules of the game, Siri is a nonjudgmental friend and teacher. Nicole Colbert, whose son, Sam, is in my son's class at LearningSpring, a (lifesaving) school for autistic children in Manhattan, said: "My son loves getting information on his favorite subjects, but he also just loves the absurdity — like, when Siri doesn't understand him and gives him a nonsense answer, or when he poses personal questions that elicit funny responses. Sam asked Siri how old she was, and she said, 'I don't talk about my age,' which just cracked him up."

But perhaps it also gave him a valuable lesson in etiquette. Gus almost invariably tells me, "You look beautiful," right before I go out the door in the morning; I think it was first Siri who showed him that you can't go wrong with that line.

Of course, most of us simply use our phone's personal assistants as an easy way to access information. For example, thanks to Henry and the question he just asked Siri, I now know that there is a website called Celebrity Bra Sizes.



But the companionability of Siri is not limited to those who have trouble communicating. We've all found ourselves like the writer Emily Listfield, having little conversations with her/him at one time or another. "I was in the middle of a breakup, and I was feeling a little sorry for myself," Ms. Listfield said. "It was midnight and I was noodling around on my iPhone, and I asked Siri, 'Should I call Richard?' Like this app is a Magic 8 Ball. Guess what: not a Magic 8 Ball. The next thing I hear is, 'Calling Richard!' and *dialing*." Ms. Listfield has forgiven Siri, and has recently considered changing her into a male voice. "But I'm worried he won't answer when I ask a question," she said. "He'll just pretend he doesn't hear."

Siri can be oddly comforting, as well as chummy. One friend reports: "I was having a bad day and jokingly turned to Siri and said, 'I love you,' just to see what would happen, and she answered, 'You are the wind beneath my wings.' And you know, it kind of cheered me up."

(Of course, I don't know what my friend is talking about. Because I wouldn't be at all cheered if I happened to ask Siri, in a low moment, "Do I look fat in these jeans?" and Siri answered, "You look fabulous.")

For most of us, Siri is merely a momentary diversion. But for some, it's more. My son's practice conversation with Siri is translating into more facility with actual humans. Yesterday I had the longest conversation with him that I've ever had. Admittedly, it was about different species of turtles and whether I preferred the red-eared slider to the diamond-backed terrapin. This might not have been my choice of topic, but it was back and forth, and it followed a logical trajectory. I can promise you that for most of my beautiful son's 13 years of existence, that has not been the case.

The developers of intelligent assistants recognize their uses to those with speech and communication problems — and some are thinking of new ways the assistants can help. According to the folks at SRI International, the research and development company where Siri began before Apple bought the technology, the next generation of virtual assistants will not just retrieve information — they will also be able to carry on more complex conversations about a person's area of interest. "Your son will be able to proactively get information about whatever he's interested in without asking for it, because the assistant will anticipate what he likes," said William Mark, vice president for information and computing sciences at SRI.

The assistant will also be able to reach children where they live. Ron Suskind, whose new book, "Life, Animated," chronicles how his autistic son came out of his shell through engagement with Disney characters, is talking to SRI about having assistants for those with autism that can be programmed to speak in the voice of the character that reaches them — for his son, perhaps Aladdin; for mine, either Kermit or Lady Gaga, either of which he is infinitely more receptive to than, say, his mother. (Mr. Suskind came up with the perfect name, too: not virtual assistants, but "sidekicks.")

Mr. Mark said he envisions assistants whose help is also visual. "For example, the assistant would be able to track eye movements and help the autistic learn to look you in the eye when talking," he said.

"See, that's the wonderful thing about technology being able to help with some of these behaviors," he added. "Getting results requires a lot of repetition. Humans are not patient. Machines are very, very patient."

I asked Mr. Mark if he knew whether any of the people who worked on Siri's language development at Apple were on the spectrum. "Well, of course, I don't know for certain," he said, thoughtfully. "But, when you think about it, you've just described half of Silicon Valley."

Of all the worries the parent of an autistic child has, the uppermost is: Will he find love? Or even companionship? Somewhere along the line, I am learning that what gives my

guy happiness is not necessarily the same as what gives me happiness. Right now, at his age, a time when humans can be a little overwhelming even for the average teenager, Siri makes Gus happy. She is his sidekick. Last night, as he was going to bed, there was this matter-of-fact exchange:

Gus: “Siri, will you marry me?”

Siri: “I’m not the marrying kind.”

Gus: “I mean, not now. I’m a kid. I mean when I’m grown up.”

Siri: “My end user agreement does not include marriage.”

Gus: “Oh, O.K.”

Gus didn’t sound too disappointed. This was useful information to have, and for me too, since it was the first time I knew that he actually *thought* about marriage. He turned over to go to sleep:

Gus: “Goodnight, Siri. Will you sleep well tonight?”

Siri: “I don’t need much sleep, but it’s nice of you to ask.”

Very nice.

“The Biggest Tell That Something Was Written by AI”

By Eve Fairbanks

May 29, 2026

The Atlantic

A few weeks ago, where I live in Johannesburg, a man ran a stop sign and crashed into my Subaru. At the scene he was frantic, unable to gather his thoughts. Half an hour later, I received a lengthy, perfectly grammatical text from him elegantly explaining how he perceived the crash had happened. For a repair quote, I wrote to a mechanic I know, a man who used to text me in curt phrases riddled with shorthand. I got a response using just the same voice as the man who’d crashed into me—the distinctive voice of AI.

In surveys, people consistently say they distrust AI-generated writing. But that hasn’t stopped more and more of us from using it in everyday life—to compose work emails and personal texts, to make shopping lists, even to write scripts for arguments with our spouses. “I feel like I’m going nuts,” the writer Jason Koebler complained in the tech outlet 404 Media, under “the cognitive load” of trying to discern whether every piece of text he reads is real or fake.

AI writing is also creeping into our most elite literary spaces—newspapers’ opinion sections, books, literary magazines. I edit professionally, often working with authors renowned for their prose. Maybe two months ago, I began receiving a kind of submission I’d never gotten before: perfectly clean, without a stray comma; uniform in length, with evenly paced paragraphs and a distinctive tone that was simultaneously breezy and grandiose. At first I was surprised that people who prided themselves on their writing would turn to AI to write for them. Even six months ago, when I occasionally identified a paragraph in a writer’s work that seemed AI-generated, they would apologize.

Now some authors tell me they've embraced AI as a "writing tool," no different from spell-check or a laptop. The phrase is protean and euphemistic, covering everything from using ChatGPT to find a quote to having it compose a long essay based on a two-sentence prompt. The reason for the change is simple: Competition in journalism and academia and grant writing and even YouTube influencing is insanely fierce. The edge goes to those who can stand out in a deluge of content, which is achieved through cleanly packaged messaging and sheer volume. Even professional communicators who are confident in their writing and unsure that AI is a perfect replacement are under increasing pressure to use it, so long as they feel they're doing so within their profession's boundaries. (*The Atlantic*, for the record, prohibits writers from using AI-generated text unless it's explicitly identified as such.)

People who aren't professional writers are making a similar calculation. AI programs' efficiency in generating smooth, grammatical text is irresistible, whether you need a savvy sentence in a job application or a line of banter on a dating app. AI-generated writing can easily trick readers, especially if they're only skimming. Tutorials exist for how to strip the telltale signs of AI use from your writing: Get rid of em dashes, colons, and of course the now-icky "It's not X; it's Y" formulations.

The problem is that the efficiency and frictionlessness that make AI appealing to writers are the same qualities that make it feel untrustworthy to readers. And readers are right not to trust it. No matter how much we may tell ourselves that AI is just a tool like spell-check, it isn't. When we use AI to flesh out ideas, we lose the most important part of the writing process: thinking.

We tend to believe that efficiency is the highest virtue, the four-hour workweek the ultimate goal. Why sweat over the introductory paragraph of an essay if an AI program can sail over whatever argumentative obstacle you have in the space of 15 seconds? But the effort and the hang-ups are, as they say, a feature of the human thought process, not a

bug. When human beings write, we judge ourselves; we stop; we backtrack. In published writing, the traces of this process are erased. But it is the process that makes human writing sensible and meaningful. Many authors describe how, when they've finally hit on the right idea, writing feels like going down a water slide; putting one sentence after another becomes easy.

When writing is hard, it's often not just because we are tired, underfed, or inefficient but because our mind is trying to tell us crucial things. How many draft texts to colleagues or family members have we all stared at in frustration, wondering why they don't feel quite right—until we finally realize that they need to be rethought completely, or not sent at all? When a book I was writing became an almost hopeless grind, I tore up 90 percent of the manuscript; it became a far more honest work for having been halted at a conceptual dead end, forcing me to turn back.

AI can't make that kind of judgment. Even if the companies that design AI programs could make them reason like a human being—a project whose hubris is underrated, given that we don't fully understand the mechanisms behind our own thought processes—they won't. After all, users consistently say they want AI models to be agreeable, compliant. That means AIs are reluctant to do what the voices in our head do all the time: interrogate the validity of our premise; misunderstand us, forcing us to explain better; insist that a query is stupid; refuse to answer a question at all. In a study published in March, a group of Stanford and Carnegie Mellon University scientists found that top AI models affirm their users' ideas 49 percent more than humans do in conversation. They also discovered that participants rated more sycophantic answers as "higher quality" and said a sycophantic attitude made them more likely to use AI again.

So we end up with canned perfection—writing that can't really be argued with, because it has no underlying deliberative reasoning process, no train of thought. As I wrote on X recently, AI writing is almost impossible to edit, because even when it sounds plausible, a

closer look will show that every element is equally off: The tone is bland; individual word choices are baffling; the structure lacks sense; key pieces of the argument are missing; facts are false. Working on AI text, as an editor, is like trying to operate on a body whose skin, muscles, veins, bones, and organs are all compromised. There's nothing to leave intact, nowhere to begin.

Recently I asked ChatGPT Pro—which I paid for to summarize financial spreadsheets; I'm not a hater—whether it was capable of viciously dismantling an argument I'd read about the stock market. "I'm tuned to avoid being casually cruel, contemptuous, or needlessly snide," it explained. "That doesn't mean I can't be sharp, skeptical, funny, or sarcastic … I can absolutely be dismissive of bad arguments, flabby prose, obvious spin, bureaucratic nonsense, or market delusions wearing a little tech-futurist hat."

I wanted to look under the hood a little, to understand the reasoning behind how the program chose its phrases. "Why would market delusions be wearing a 'little' tech-futurist hat, rather than a big one?" I queried.

The program was incapable of pausing to wonder whether it had chosen the wrong word. Instead, it sought to justify itself in increasingly bewildering ways. "Because 'little' makes it funnier," ChatGPT Pro answered. "The market delusion wants to stride into the room as Prometheus; 'little hat' makes it look like a raccoon that found a conference badge."

Why a raccoon with a conference badge? I pressed. Why not a dog, or a frog?

"Raccoons live in alleys, storm drains, garbage nights, and morally humid environments," it wrote. "Frogs have existential agency, while raccoons have logistical agency. A frog at a VC afterparty seems like an omen. A raccoon at a VC afterparty seems like it got on the guest list … Think: alley behind a conference hotel, the loading

dock of a casino, the VC afterparty where everyone is saying ‘alignment’ but meaning ‘exit liquidity’ … Alignment sounds like ethics, safety, shared goals, humanity, the future. It has dry-cleaned blazer energy … Exit liquidity is socially clammy.”

All of these sentences are grammatically perfect. They also make no sense. And all substantially AI-generated writing is like this, under the hood. ChatGPT Pro agrees, though it has no idea why. “What if I said the whole raccoon thing made no sense as a metaphor?” I asked the program.

“I would agree,” it replied, ever servile. “At first, the raccoon worked as a throwaway image … A metaphor can survive one or two explanatory layers. After that, it starts looking like a raccoon in a trench coat.”

This is the kind of communication we’re becoming surrounded with. Its infiltration into every domain of our lives can’t be stopped. Even people who don’t use AI will begin sounding more like it. (A preprint by the Max Planck Institute for Human Development found that in off-the-cuff verbal conversations, such as podcast discussions, people are already exhibiting “a measurable and abrupt increase in the use of words preferentially generated by ChatGPT—such as *delve*, *comprehend*, *boast*, *swift*, and *meticulous*.”) After all, we remain so much smarter than machines, so much subtler, and thus so much quicker to learn and pick up cultural cues. The difference in how we operate will be extraordinary, and not at all hypothetical. Ten years ago I composed a reconciliatory email to a boyfriend but never sent it, because I couldn’t get the phrasing right. Only much later did I realize I simply didn’t mean what I’d been trying to write. If I’d had an AI program to help me get over the hump, I’d be married to a different person. A much less suitable one.

Maybe human writing will become like cloth-aged cheese or handloom rugs, an artisanal product created effortfully. Maybe we will come to treasure older writing. Herman

Melville, George Orwell, Toni Morrison—all authenticated. Writing like this will be a fossil record for a kind of thought process we buried without realizing it. The other night, as I was drifting off to sleep, a 19th-century poem popped into my head:

Wynken, Blynken, and Nod one night
Sailed off in a wooden shoe,—
Sailed on a river of crystal light
Into a sea of dew.

A children's rhyme, but it had a new beauty.

Or maybe smooth communiqués that arrive on time and betray no confusion, doubt, or internal struggle—that polish up our images as affable, efficient, and universally, if superficially, wise—is what we want. But at least we should know what we're sacrificing.

“Is Google Making Us Stupid?”

Nicholas Carr

***The Atlantic* July/August 2008**

i.

"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.

ii.

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 re-created 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.

iii.

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."

iv.

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."

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.

v.

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.