

AP English Language and Composition
2023 Summer Assignment

In 2006, American author Michael Chabon published an essay about the Clock of the Long Now, a futurist project that sought (and in 2023, is still seeking) to build a timepiece/art installation that would keep time for the next 10,000 years. Ultimately, the Clock of the Long Now is just a starting point for Chabon to think about what the world might look like in his future, his children's future, his grandchildren's future.

The following five essays are interested in **futurism**: an understanding that contemporary progress will have a meaningful, often transformative impact on the future.

Please read the following:

"The Future Will Have to Wait" by Michael Chabon (2006)

"Is Google Making Us Stupid?" by Nicholas Carr (2008)

"Can Mark Zuckerberg Fix Facebook Before It Break's Democracy?" by Evan Osnos (2018)

"The Age of Instagram Face" by Jia Tolentino (2019)

"There is No A.I." by Jaron Lanier (2023)

Respond to the following writing prompts:

1. **Select a single sentence from each essay and explain, in one short paragraph apiece, why you find that sentence interesting.** It might be the sentence's structure, word choice, what the sentence is saying, its relationship to the rest of the essay, another essay, or your broader experience of the world, whatever. Pick a sentence and explain why you're drawn to it. (five paragraphs; one for each sentence)
2. Select **one** of the essays and do the following: **Write a short essay that a) identifies the author's position regarding their chosen topic** (Chabon and the abstract concept of the future; Carr and the Internet search engine, Osnos and social media platforms, Tolentino and plastic surgery, Lanier and generative AI) **and b) analyze how the author uses writing strategies (literary techniques or rhetorical strategies) to develop this position.** (1.5 pages)
3. Consider Chabon's optimism at the end of "The Future Will Have to Wait": While he observes a "shifting tension between the bright promise and the bleak menace of the Future," Chabon still believes in a Future built on humanity's benevolent attempts at progress and perfection. Now, consider the different sorts of futurism explored in the other four essays—Google's endless and accessible information repository, Zuckerberg's social media town square, Tolentino's cyborgian Instagram Face, Lanier's Data Dignity in the face of AI. Those essays offer a spectrum of perspectives on our attempts at improvement and perfection. **Write an essay where you present your perspective on whether humans will be here to see the Clock of the Long Now strike 10,000 years. While you may use the essays for support, your goal in this essay is to develop an authentic position supported by evidence from your reading, experiences, or observations.** (2 pages)

Your responses will be submitted to Turn It In on the first day of classes.

You are also, in accordance with the high school's summer reading program, expected to read two books of your choice (feel free to check the summer reading padlet for suggestions).

The Future Will Have to Wait

11-14 minutes

Published on Sunday, January 22, 02006 • 17 years, 3 months ago

Written by Michael Chabon for Details

I was reading, in a recent issue of *Discover*, about the Clock of the Long Now. Have you heard of this thing? It is going to be a kind of gigantic mechanical computer, slow, simple and ingenious, marking the hour, the day, the year, the century, the millennium, and the precession of the equinoxes, with a huge orrery to keep track of the immense ticking of the six naked-eye planets on their great orbital mainspring. The Clock of the Long Now will stand sixty feet tall, cost tens of millions of dollars, and when completed its designers and supporters, among them visionary engineer Danny Hillis, a pioneer in the concept of massively parallel processing; Whole Earth mahatma Stewart Brand; and British composer Brian Eno (one of my household gods), plan to hide it in a cave in the Great Basin National Park in Nevada, a day's hard walking from anywhere. Oh, and it's going to run for ten thousand years. That is about as long a span as separates us from the first makers of pottery, which is among the oldest technologies we have. Ten thousand years is twice as old as the pyramid of Cheops, twice as old as that mummified body found preserved in the Swiss Alps, which is one of the oldest mummies ever discovered. The Clock of the Long Now is being designed to thrive under regular human maintenance along the whole of that long span, though during periods when no one is around to tune it, the giant clock will contrive to adjust itself. But even if the Clock of the Long Now fails to last ten thousand years, even if it breaks down after half or a quarter or a tenth that span, this mad contraption will already have long since fulfilled its purpose. Indeed the Clock may have accomplished its greatest task before it is ever finished, perhaps without ever being built at all. The point of the Clock of the Long Now is not to measure out the passage, into their unknown future, of the race of creatures that built it. The point of the Clock is to revive and restore the whole idea of the Future, to get us thinking about the Future again, to the degree if not in quite the way same way that we used to do, and to reintroduce the notion that we don't just bequeath the future—though we do, whether we think about it or not. We also, in the very broadest sense of the first person plural pronoun, inherit it.

The Sex Pistols, strictly speaking, were right: there is no future, for you or for me. The future, by definition, does not exist. "The Future," whether you capitalize it or not, is always just an idea, a proposal, a scenario, a sketch for a mad contraption that may or may not work. "The Future" is a story we tell, a narrative of hope, dread or wonder. And it's a story that, for a while now, we've been pretty much living without.

Ten thousand years from now: can you imagine that day? Okay, but do you? Do you believe "the Future" is going to happen? If the Clock works the way that it's supposed to do—if it lasts—do you believe there will be a human being around to witness, let alone mourn its passing, to appreciate its accomplishment, its faithfulness, its immense antiquity? What about five thousand years from now, or even five hundred? Can you extend the horizon of your expectations for our world, for our complex of civilizations and cultures, beyond the lifetime of your own children, of the next two or three generations? Can you even imagine the survival of the world beyond the present presidential administration?

I was surprised, when I read about the Clock of the Long Now, at just how long it had been since I had given any thought to the state of the world ten thousand years hence. At one time I was a frequent visitor to that imaginary mental locale. And I don't mean merely that I regularly encountered "the Future" in the pages of science fiction novels or comic books, or when watching a TV show like *The Jetsons* (1962) or a movie like *Beneath the Planet of the Apes* (1970). The story of the Future was told to me, when I was growing up, not just by popular art and media but by public and domestic architecture, industrial design, school textbooks, theme parks, and by public institutions from museums to government agencies. I heard the story of the Future when I looked at the space-ranger profile of the Studebaker Avanti, at Tomorrowland through the portholes of the Disneyland monorail, in the tumbling plastic counters of my father's Seth Thomas Speed Read clock. I can remember writing a report in sixth grade on hydroponics; if you had tried to tell me then that by 2005 we would still be growing our vegetables in dirt, you would have broken my heart.

Even thirty years after its purest expression on the covers of pulp magazines like *Amazing Stories* and, supremely, at the New York World's Fair of 1939, the collective cultural narrative of the Future remained largely an optimistic one of the impending blessings of technology and the benevolent, computer-assisted meritocracy of Donald Fagen's "fellows with compassion and vision." But by the early seventies—indeed from early in the history of the Future—it was not all farms under the sea and family vacations on Titan. Sometimes the Future could be a total downer. If nuclear holocaust didn't wipe everything out, then humanity would be enslaved to computers, by the ineluctable syllogisms of "the Machine." My childhood dished up a series of grim cinematic prognostications best exemplified by the Hestonian trilogy that began with the first *Planet of the Apes* (1968) and continued through *The Omega Man* (1971) and *Soylent Green* (1973). Images of future dystopia were rife in rock albums of the day, as on David Bowie's *Diamond Dogs* (1974) and Rush's *2112* (1976), and the futures presented by seventies writers of science fiction such as John Brunner tended to be unrelentingly or wryly bleak.

In the aggregate, then, stories of the Future presented an enchanting ambiguity. The other side of the marvelous *Jetsons* future might be a story of worldwide corporate-authoritarian technotyranny, but the other side of a post-apocalyptic mutational nightmare landscape like that depicted in *The Omega Man* was a landscape of semi-barbaric splendor and unfettered (if dangerous) freedom to roam, such as I found in the pages of Jack Kirby's classic adventure comic book *Kamandi, The Last Boy on Earth* (1972-76). That ambiguity and its enchantment, the shifting tension between the bright promise and the bleak menace of the Future, was in itself a kind of story about the ways, however freakish or tragic, in which humanity (and by implication American culture and its values however freakish and tragic) would, in spite of it all, continue. Eed plebnista, intoned the devolved Yankees, in the *Star Trek* episode "The Omega Glory," who had somehow managed to hold on to and venerate as sacred gobbledygook the Preamble to the Constitution, norkon forden perfectunun. All they needed was a Captain Kirk to come and add a little interpretive water to the freeze-dried document, and the American way of life would flourish again.

I don't know what happened to the Future. It's as if we lost our ability, or our will, to envision anything beyond the next hundred years or so, as if we lacked the fundamental faith that there will in fact be any future at all beyond that not-too-distant date. Or maybe we stopped talking about the Future around the time that, with its microchips and its twenty-four-hour news cycles, it arrived. Some days when you pick up the newspaper it seems to have been co-written by J. G. Ballard, Isaac Asimov, and Philip K. Dick. Human sexual reproduction without male genetic material, digital viruses, identity theft, robot firefighters and minesweepers, weather

control, pharmaceutical mood engineering, rapid species extinction, US Presidents controlled by little boxes mounted between their shoulder blades, air-conditioned empires in the Arabian desert, transnational corporatocracy, reality television—some days it feels as if the imagined future of the mid-twentieth century was a kind of checklist, one from which we have been too busy ticking off items to bother with extending it. Meanwhile, the dwindling number of items remaining on that list—interplanetary colonization, sentient computers, quasi-immortality of consciousness through brain-download or transplant, a global government (fascist or enlightened)—have been represented and re-represented so many hundreds of times in films, novels and on television that they have come to seem, paradoxically, already attained, already known, lived with, and left behind. Past, in other words.

This is the paradox that lies at the heart of our loss of belief or interest in the Future, which has in turn produced a collective cultural failure to imagine that future, any Future, beyond the rim of a couple of centuries. The Future was represented so often and for so long, in the terms and characteristic styles of so many historical periods from, say, Jules Verne forward, that at some point the idea of the Future—along with the cultural appetite for it—came itself to feel like something historical, outmoded, no longer viable or attainable.

If you ask my eight-year-old about the Future, he pretty much thinks the world is going to end, and that's it. Most likely global warming, he says—floods, storms, desertification—but the possibility of viral pandemic, meteor impact, or some kind of nuclear exchange is not alien to his view of the days to come. Maybe not tomorrow, or a year from now. The kid is more than capable of generating a full head of optimistic steam about next week, next vacation, his tenth birthday. It's only the world a hundred years on that leaves his hopes a blank. My son seems to take the end of everything, of all human endeavor and creation, for granted. He sees himself as living on the last page, if not in the last paragraph, of a long, strange and bewildering book. If you had told me, when I was eight, that a little kid of the future would feel that way—and that what's more, he would see a certain justice in our eventual extinction, would think the world was better off without human beings in it—that would have been even worse than hearing that in 2006 there are no hydroponic megafarms, no human colonies on Mars, no personal jetpacks for everyone. That would truly have broken my heart.

When I told my son about the Clock of the Long Now, he listened very carefully, and we looked at the pictures on the Long Now Foundation's website. "Will there really be people then, Dad?" he said. "Yes," I told him without hesitation, "there will." I don't know if that's true, any more than do Danny Hillis and his colleagues, with the beating clocks of their hopefulness and the orreries of their imaginations. But in having children—in engendering them, in loving them, in teaching them to love and care about the world—parents are betting, whether they know it or not, on the Clock of the Long Now. They are betting on their children, and their children after them, and theirs beyond them, all the way down the line from now to 12,006. If you don't believe in the Future, unreservedly and dreamingly, if you aren't willing to bet that somebody will be there to cry when the Clock finally, ten thousand years from now, runs down, then I don't see how you can have children. If you have children, I don't see how you can fail to do everything in your power to ensure that you win your bet, and that they, and their grandchildren, and their grandchildren's grandchildren, will inherit a world whose perfection can never be accomplished by creatures whose imagination for perfecting it is limitless and free. And I don't see how anybody can force me to pay up on my bet if I turn out, in the end, to be wrong.


Written by Michael Chabon for Details. Originally published in 02006.

Is Google Making Us Stupid?

WHAT THE INTERNET IS DOING TO OUR BRAINS

By Nicholas Carr

Illustration by Guy Billout



"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.”

Also see:

Living With a Computer (July 1982)

"The process works this way. When I sit down to write a letter or start the first draft of an article, I simply type on the keyboard and the words appear on the screen..." By James Fallows

"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.

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.”

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.

This article available online at:

<http://www.theatlantic.com/magazine/archive/2008/07/is-google-making-us-stupid/306868/>

Copyright © 2013 by The Atlantic Monthly Group. All Rights Reserved.

CAN MARK ZUCKERBERG FIX FACEBOOK BEFORE IT BREAKS DEMOCRACY?

*The most famous entrepreneur of his generation is facing a
public reckoning with the power of Big Tech.*

By Evan Osnos

September 10, 2018

Zuckerberg is at the center of a debate about the moral character of Silicon Valley and the conscience of its leaders. Illustration by Javier Jaén; photograph by David Yellen / Corbis / Getty



Listen to this story

At ten o'clock on a weekday morning in August, Mark Zuckerberg, the chairman and C.E.O. of Facebook, opened the front door of his house in Palo Alto, California, wearing the tight smile of obligation. He does not enjoy interviews, especially after two years of ceaseless controversy. Having got his start as a programmer with a nocturnal bent, he is also not a morning person. Walking toward the kitchen, which has a long farmhouse table and cabinets painted forest green, he said, "I haven't eaten breakfast yet. Have you?"

Since 2011, Zuckerberg has lived in a century-old white clapboard Craftsman in the Crescent Park neighborhood, an enclave of giant oaks and historic homes not far from Stanford University. The house, which cost seven million dollars, affords him a sense of sanctuary. It's set back from the road, shielded by hedges, a wall, and mature trees. Guests enter through an arched wooden gate and follow a long gravel path to a front lawn with a saltwater pool in the center. The year after Zuckerberg bought the house, he and his longtime girlfriend, Priscilla Chan, held their wedding in the back yard, which encompasses gardens, a pond, and a shaded pavilion. Since then, they have had two children, and acquired a seven-hundred-acre estate in Hawaii, a ski retreat in Montana, and a four-story town house on Liberty Hill, in San Francisco. But the family's full-time residence is here, a ten-minute drive from Facebook's headquarters.

Occasionally, Zuckerberg records a Facebook video from the back yard or the dinner table, as is expected of a man who built his fortune exhorting employees to keep “pushing the world in the direction of making it a more open and transparent place.” But his appetite for personal openness is limited. Although Zuckerberg is the most famous entrepreneur of his generation, he remains elusive to everyone but a small circle of family and friends, and his efforts to protect his privacy inevitably attract attention. The local press has chronicled his feud with a developer who announced plans to build a mansion that would look into Zuckerberg’s master bedroom. After a legal fight, the developer gave up, and Zuckerberg spent forty-four million dollars to buy the houses surrounding his. Over the years, he has come to believe that he will always be the subject of criticism. “We’re not—pick your noncontroversial business—selling dog food, although I think that people who do that probably say there is controversy in that, too, but this is an inherently cultural thing,” he told me, of his business. “It’s at the intersection of technology and psychology, and it’s very personal.”

He carried a plate of banana bread and a carafe of water into the living room, and settled onto a navy-blue velvet sofa. Since co-founding Facebook, in 2004, his uniform has evolved from hoodies and flip-flops to his current outfit, a gray sweater, indigo jeans, and black Nikes. At thirty-four, Zuckerberg, who has very fair skin, a tall forehead, and large eyes, is leaner than when he first became a public figure, more than a decade ago. On the porch, next to the front door, he keeps a Peloton stationary bike, a favorite accessory in the tech world, which live-streams a personal trainer to your home. Zuckerberg uses the machine, but he does not love cycling. A few years ago, on his first attempt to use a road bike with racing pedals, he forgot to unclip, tipped over, and broke his arm. Except for cycling on his porch, he said, “I haven’t clipped in since.”

He and his wife prefer board games to television, and, within reach of the couch, I noticed a game called Ricochet Robots. “It gets extremely competitive,” Zuckerberg said. “We play with these friends, and one of them is a genius at this. Playing with him is just infuriating.” Dave Morin, a former Facebook employee who is the founder and C.E.O. of Sunrise Bio, a startup seeking cures for depression, used to play Risk with Zuckerberg at the office. “He’s not playing you in a game of Risk. He’s playing you in a game of games,” Morin told me. “The first game, he might amass all his armies on one property, and the next game he might spread them all over the place. He’s trying to figure out the psychological way to beat you in all the games.”

Across the tech industry, the depth of Zuckerberg’s desire to win is often remarked upon. Dick Costolo, the former C.E.O. of Twitter, told me, “He’s a ruthless execution machine, and if he has decided to come after you, you’re going to take a beating.” Reid Hoffman, the founder of LinkedIn, said, “There are a number of people in the Valley who have a perception of Mark that he’s really aggressive and competitive. I think some people are a little hesitant about him from that perspective.” Hoffman has been an investor in Facebook since its early days, but for a long time he sensed that Zuckerberg kept his distance because they were both building social networks. “For many years, it was, like, ‘Your LinkedIn thing is going to be crushed, so even though we’re friendly, I don’t want to get too close to you personally, because I’m going to crush you.’ Now, of course, that’s behind us and we’re good friends.”

When I asked Zuckerberg about this reputation, he framed the dynamic differently. The survival of any social-media business rests on “network effects,” in which the value of the network grows only by finding new users. As a result, he said, “there’s a natural zero-sumness. If we’re going to achieve what we want to, it’s not just about building the best features. It’s about building the best community.” He added, “I care about succeeding. And, yes, sometimes you have to beat someone to something, in order to get to the next thing. But that’s not primarily the way that I think I roll.”

For many years, Zuckerberg ended Facebook meetings with the half-joking exhortation "Domination!" Although he eventually stopped doing this (in European legal systems, "dominance" refers to corporate monopoly), his discomfort with losing is undimmed. A few years ago, he played Scrabble on a corporate jet with a friend's daughter, who was in high school at the time. She won. Before they played a second game, he wrote a simple computer program that would look up his letters in the dictionary so that he could choose from all possible words. Zuckerberg's program had a narrow lead when the flight landed. The girl told me, "During the game in which I was playing the program, everyone around us was taking sides: Team Human and Team Machine."

If Facebook were a country, it would have the largest population on earth. More than 2.2 billion people, about a third of humanity, log in at least once a month. That user base has no precedent in the history of American enterprise. Fourteen years after it was founded, in Zuckerberg's dorm room, Facebook has as many adherents as Christianity.

A couple of years ago, the company was still revelling in its power. By collecting vast quantities of information about its users, it allows advertisers to target people with precision—a business model that earns Facebook more ad revenue in a year than all American newspapers combined. Zuckerberg was spending much of his time conferring with heads of state and unveiling plans of fantastical ambition, such as building giant drones that would beam free Internet (including Facebook) into developing countries. He enjoyed extraordinary control over his company; in addition to his positions as chairman and C.E.O., he controlled about sixty per cent of shareholder votes, thanks to a special class of stock with ten times the power of ordinary shares. His personal fortune had grown to more than sixty billion dollars. Facebook was one of four companies (along with Google, Amazon, and Apple) that dominated the Internet; the combined value of their stock is larger than the G.D.P. of France.



"Best of all, he has this wonderful passion for life that he doesn't expect me to share."



For years, Facebook had heard concerns about its use of private data and its ability to shape people's behavior. The company's troubles came to a head during the Presidential election of 2016, when propagandists used the site to spread misinformation that helped turn society against itself. Some of the culprits were profiteers who gamed Facebook's automated systems with toxic political clickbait known as "fake news." In a prime example, at least a hundred Web sites were traced to Veles, Macedonia, a small city where entrepreneurs, some still in high school, discovered that posting fabrications to pro-Donald Trump Facebook groups unleashed geysers of traffic. Fake-news sources also paid Facebook to "microtarget" ads at users who had proved susceptible in the past.

The other culprits, according to U.S. intelligence, were Russian agents who wanted to sow political chaos and help Trump win. In February, Robert Mueller, the special counsel investigating Russia's role in the election, charged thirteen Russians with an "interference operation" that made use of Facebook, Twitter, and Instagram. The Internet Research Agency, a firm in St. Petersburg working for the Kremlin, drew hundreds of thousands of users to Facebook groups optimized to stoke outrage,

including Secured Borders, Blacktivist, and Defend the 2nd. They used Facebook to organize offline rallies, and bought Facebook ads intended to hurt Hillary Clinton's standing among Democratic voters. (One read "Hillary Clinton Doesn't Deserve the Black Vote.") With fewer than a hundred operatives, the I.R.A. achieved an astonishing impact: Facebook estimates that the content reached as many as a hundred and fifty million users.

At the same time, former Facebook executives, echoing a growing body of research, began to voice misgivings about the company's role in exacerbating isolation, outrage, and addictive behaviors. One of the largest studies, published last year in the *American Journal of Epidemiology*, followed the Facebook use of more than five thousand people over three years and found that higher use correlated with self-reported declines in physical health, mental health, and life satisfaction. At an event in November, 2017, Sean Parker, Facebook's first president, called himself a "conscientious objector" to social media, saying, "God only knows what it's doing to our children's brains." A few days later, Chamath Palihapitiya, the former vice-president of user growth, told an audience at Stanford, "The short-term, dopamine-driven feedback loops that we have created are destroying how society works—no civil discourse, no coöperation, misinformation, mistruth." Palihapitiya, a prominent Silicon Valley figure who worked at Facebook from 2007 to 2011, said, "I feel tremendous guilt. I think we all knew in the back of our minds." Of his children, he added, "They're not allowed to use this shit." (Facebook replied to the remarks in a statement, noting that Palihapitiya had left six years earlier, and adding, "Facebook was a very different company back then.")

In March, Facebook was confronted with an even larger scandal: the *Times* and the British newspaper the *Observer* reported that a researcher had gained access to the personal information of Facebook users and sold it to Cambridge Analytica, a consultancy hired by Trump and other Republicans which advertised using "psychographic" techniques to manipulate voter behavior. In all, the personal data of eighty-seven million people had been harvested. Moreover, Facebook had known of the problem since December of 2015 but had said nothing to users or regulators. The company acknowledged the breach only after the press discovered it.

The Cambridge Analytica revelations touched off the most serious crisis in Facebook's history, and, with it, a public reckoning with the power of Big Tech. Facebook is now under investigation by the F.B.I., the Securities and Exchange Commission, the Department of Justice, and the Federal Trade Commission, as well as by authorities abroad, from London to Brussels to Sydney. Facebook's peers and rivals have expressed conspicuously little sympathy. Elon Musk deleted his Facebook pages and those of his companies, Tesla and SpaceX. Tim Cook, the C.E.O. of Apple, told an interviewer, "We could make a ton of money if we monetized our customer," but "we've elected not to do that." At Facebook's annual shareholder meeting, in May, executives struggled to keep order. An investor who interrupted the agenda to argue against Zuckerberg's renomination as chairman was removed. Outside, an airplane flew a banner that read "YOU BROKE DEMOCRACY." It was paid for by Freedom from Facebook, a coalition of progressive groups that have asked the F.T.C. to break up the company into smaller units.

On July 25th, Facebook's stock price dropped nineteen per cent, cutting its market value by a hundred and nineteen billion dollars, the largest one-day drop in Wall Street history. Nick Bilton, a technology writer at *Vanity Fair*, tweeted that Zuckerberg was losing \$2.7 million per second, "double what the average American makes in an entire lifetime." Facebook's user base had flatlined in the U.S. and Canada, and dropped slightly in Europe, and executives warned that revenue growth would decline further, in part because the scandals had led users to opt out of allowing Facebook to collect some data. Facebook depends on trust, and the events of the past two years had made people wonder whether the company deserved it.

Zuckerberg's friends describe his travails as a by-product of his success. He is often compared to another Harvard dropout, Bill Gates, who has been his mentor in business and philanthropy. Gates told me, "Somebody who is smart, and rich, and ends up not acknowledging problems as quickly as they should will be attacked as arrogant. That comes with the territory." He added, "I wouldn't say that Mark's an arrogant individual." But, to critics, Facebook is guilty of a willful blindness driven by greed, naïveté, and contempt for oversight.

In a series of conversations over the summer, I talked to Zuckerberg about Facebook's problems, and about his underlying views on technology and society. We spoke at his home, at his office, and by phone. I also interviewed four dozen people inside and outside the company about its culture, his performance, and his decision-making. I found Zuckerberg straining, not always coherently, to grasp problems for which he was plainly unprepared. These are not technical puzzles to be cracked in the middle of the night but some of the subtlest aspects of human affairs, including the meaning of truth, the limits of free speech, and the origins of violence.

Zuckerberg is now at the center of a full-fledged debate about the moral character of Silicon Valley and the conscience of its leaders. Leslie Berlin, a historian of technology at Stanford, told me, "For a long time, Silicon Valley enjoyed an unencumbered embrace in America. And now everyone says, Is this a trick? And the question Mark Zuckerberg is dealing with is: Should my company be the arbiter of truth and decency for two billion people? Nobody in the history of technology has dealt with that."

Facebook's headquarters, at 1 Hacker Way, in Menlo Park, overlooking the salt marshes south of San Francisco, has the feel of a small, prosperous dictatorship, akin to Kuwait or Brunei. The campus is a self-contained universe, with the full range of free Silicon Valley perks: dry cleaning, haircuts, music lessons, and food by the acre, including barbecue, biryani, and salad bars. (New arrivals are said to put on the "Facebook fifteen.") Along with stock options and generous benefits, such trappings have roots in the nineteen-seventies, when, Leslie Berlin said, founders aspired to create pleasant workplaces and stave off the rise of labor unions. The campus, which was designed with the help of consultants from Disney, is arranged as an ersatz town that encircles a central plaza, with shops and restaurants and offices along a main street. From the air, the word "HACK" is visible in gigantic letters on the plaza pavement.

On Zuckerberg's campus, he is king. Executives offer fulsome praise. David Marcus, who runs Facebook's blockchain project, told me recently, "When I see him portrayed in certain ways, it really hurts me personally, because it's not the guy he is." Even when colleagues speak more candidly, on the whole they like him. "He's not an asshole," a former senior executive told me. "That's why people work there so long."

Before I visited Zuckerberg for the first time, in June, members of his staff offered the kind of advice usually reserved for approaching a skittish bird: proceed gingerly, build a connection, avoid surprises. The advice, I discovered, wasn't necessary. In person, he is warmer and more direct than his public pronouncements, which resemble a politician's bland plumb, would suggest. The contrast between the public and the private Zuckerberg reminded me of Hillary Clinton. In both cases, friends complain that the popular image is divorced from the casual, funny, generous person they know. Yet neither Zuckerberg nor Clinton has found a way to publicly express a more genuine persona. In Zuckerberg's case, moments of self-reflection are so rare that, last spring, following a CNN interview in which he said that he wanted to build a company that "my girls are going to grow up and be proud of me for," the network framed the clip as a news event, with the title "Zuckerberg in rare emotional moment."

I asked Zuckerberg about his aversion to opening up. "I'm not the most polished person, and I will say something wrong, and you see the cost of that," he said. "I don't want to inflict that pain, or do something that's going to not reflect well on the people around me." In the most recent flap, a few weeks earlier, he had told Kara Swisher, the host of the "[Recode Decode](#)" podcast, that he permits Holocaust deniers on Facebook because he isn't sure if they are "intentionally getting it wrong." After a furor erupted, he issued a statement saying that he finds Holocaust denial "deeply offensive." Zuckerberg told me, "In an alternate world where there weren't the compounding experiences that I had, I probably would have gotten more comfortable being more personal, and out there, and I wouldn't have felt pushback every time I did something. And maybe my persona, or at least how I felt comfortable acting publicly, would shift."

The downside of Zuckerberg's exalted status within his company is that it is difficult for him to get genuine, unexpurgated feedback. He has tried, at times, to puncture his own bubble. In 2013, as a New Year's resolution, he pledged to meet someone new, outside Facebook, every day. In 2017, he travelled to more than thirty states on a "listening tour" that he hoped would better acquaint him with the outside world. David Plouffe, President Obama's former campaign manager, who is now the head of policy and advocacy at the Chan Zuckerberg Initiative, the family's philanthropic investment company, attended some events on the tour. He told me, "When a politician goes to one of those, it's an hour, and they're talking for fifty of those minutes. He would talk for, like, five, and just ask questions."

But the exercise came off as stilted and tone-deaf. Zuckerberg travelled with a professional photographer, who documented him feeding a calf in Wisconsin, ordering barbecue, and working on an assembly line at a Ford plant in Michigan. Online, people joked that the photos made him look like an extraterrestrial exploring the human race for the first time. A former Facebook executive who was involved in the tour told a friend, "No one wanted to tell Mark, and no one did tell Mark, that this really looks just dumb."

Zuckerberg has spent nearly half his life inside a company of his own making, handpicking his lieutenants, and sculpting his environment to suit him. Even Facebook's signature royal blue reflects his tastes. He is red-green color-blind, and he chose blue because he sees it most vividly. [Sheryl Sandberg](#), the chief operating officer, told me, "Sometimes Mark will say, in front of the company, 'Well, I've never worked anywhere else, but Sheryl tells me . . .'" She went on, "He acknowledges he doesn't always have the most experience. He's only had the experience he's had, and being Mark Zuckerberg is pretty extraordinary."

Long before it seemed inevitable or even plausible, Mark Elliot Zuckerberg had an outsized sense of his own potential. It was "a teleological frame of feeling almost chosen," a longtime friend told me. "I think Mark has always seen himself as a man of history, someone who is destined to be great, and I mean that in the broadest sense of the term." Zuckerberg has observed that more than a few giants of history grew up in bourgeois comfort near big cities and then channelled those advantages into transformative power.

In Zuckerberg's case, the setting was Dobbs Ferry, New York, a Westchester County suburb twenty-five miles north of New York City. His mother, Karen Kempner, grew up in Queens; on a blind date, she met a mailman's son, Edward Zuckerberg, of Flatbush, who was studying to be a dentist. They married and had four children. Mark, the only boy, was the second-oldest. His mother, who had become a psychiatrist, eventually gave up her career to take care of the kids and manage the dental office, which was connected to the family home. Of his father, Zuckerberg told me, "He was a dentist, but he was also a huge techie. So he always had not just a system for drilling teeth but, like, the laser system for drilling teeth that was controlled by the computer." Ed

Zuckerberg marketed himself as the Painless Dr. Z, and later drummed up dentistry business with a direct-mail solicitation that declared, "I am literally the Father of Facebook!" (Since 2013, Zuckerberg's parents have lived in California, where Ed practices part time and lectures on using social media to attract patients.)

In the nineteen-eighties and nineties, Ed bought early personal computers—the Atari 800, the I.B.M. XT—and Mark learned to code. At twelve, he set up his first network, ZuckNet, on which messages and files could be shared between the house and his father's dental office. Rabbi David Holtz, of Temple Beth Abraham, in Tarrytown, told me that he watched Zuckerberg with other kids and sensed that he was "beyond a lot of his peers. He was thinking about things that other people were not." When I asked Zuckerberg where his drive came from, he traced it to his grandparents, who had immigrated from Europe in the early twentieth century. "They came over, went through the Great Depression, had very hard lives," he said. "Their dream for their kids was that they would each become doctors, which they did, and my mom just always believed that we should have a bigger impact." His eldest sister, Randi, an early Facebook spokesperson, has gone on to write books and host a radio show; Donna received her Ph.D. in classics from Princeton and edits an online classics journal; Arielle has worked at Google and as a venture capitalist.

When Zuckerberg was a junior in high school, he transferred to Phillips Exeter Academy, where he spent most of his time coding, fencing, and studying Latin. Ancient Rome became a lifelong fascination, first because of the language ("It's very much like coding or math, and so I appreciated that") and then because of the history. Zuckerberg told me, "You have all these good and bad and complex figures. I think Augustus is one of the most fascinating. Basically, through a really harsh approach, he established two hundred years of world peace." For non-classics majors: Augustus Caesar, born in 63 B.C., staked his claim to power at the age of eighteen and turned Rome from a republic into an empire by conquering Egypt, northern Spain, and large parts of central Europe. He also eliminated political opponents, banished his daughter for promiscuity, and was suspected of arranging the execution of his grandson.

"What are the trade-offs in that?" Zuckerberg said, growing animated. "On the one hand, world peace is a long-term goal that people talk about today. Two hundred years feels unattainable." On the other hand, he said, "that didn't come for free, and he had to do certain things." In 2012, Zuckerberg and Chan spent their honeymoon in Rome. He later said, "My wife was making fun of me, saying she thought there were three people on the honeymoon: me, her, and Augustus. All the photos were different sculptures of Augustus." The couple named their second daughter August.

In 2002, Zuckerberg went to Harvard, where he embraced the hacker mystique, which celebrates brilliance in pursuit of disruption. “The ‘fuck you’ to those in power was very strong,” the longtime friend said. In 2004, as a sophomore, he embarked on the project whose origin story is now well known: the founding of Thefacebook.com with four fellow-students (“the” was dropped the following year); the legal battles over ownership, including a suit filed by twin brothers, Cameron and Tyler Winklevoss, accusing Zuckerberg of stealing their idea; the disclosure of embarrassing messages in which Zuckerberg mocked users for giving him so much data (“they ‘trust me.’ dumb fucks,” he wrote); his regrets about those remarks, and his efforts, in the years afterward, to convince the world that he has left that mind-set behind.

During Zuckerberg’s sophomore year, in line for the bathroom at a party, he met Priscilla Chan, who was a freshman. Her parents, who traced their roots to China, had grown up in Vietnam and arrived in the U.S. as refugees after the war, settling in Quincy, Massachusetts, where they washed dishes in a Chinese restaurant. Priscilla was the eldest of three daughters, and the first member of her family to go to college. “I suddenly go to Harvard, where there’s this world where people had real and meaningful intellectual pursuits,” she said. “Then I met Mark, who so exemplified that.” She was struck by how little Zuckerberg’s background had in common with her own. “Fifty per cent of people go to college from the high school I went to. You could learn how to be a carpenter or a mechanic,” she said. “I was just, like, ‘This person speaks a whole new language and lives in a framework that I’ve never seen before.’” She added, “Maybe there was some judgment on my part: ‘You don’t understand me because you went to Phillips Exeter,’” but, she said, “I had to realize early on that I was not going to change who Mark was.” After Harvard, Chan taught in a primary school and eventually became a pediatrician. In 2017, she stopped seeing patients to be the day-to-day head of the Chan Zuckerberg Initiative. When I asked Chan about how Zuckerberg had responded at home to the criticism of the past two years, she talked to me about *Sitzfleisch*, the German term for sitting and working for long periods of time. “He’d actually sit so long that he froze up his muscles and injured his hip,” she said.

After his sophomore year, Zuckerberg moved to Palo Alto and never left. Even by the standards of Silicon Valley, Facebook’s first office had a youthful feel. Zuckerberg carried two sets of business cards. One said “I’m CEO . . . bitch!” Visitors encountered a graffiti mural of a scantily clad woman riding a Rottweiler. In Adam Fisher’s “Valley of Genius,” an oral history of Silicon Valley, an early employee named Ezra Callahan muses, “‘How much was the direction of the internet influenced by the perspective of nineteen-, twenty-, twenty-one-year-old well-off white boys?’ That’s a real question that sociologists will be studying forever.”

Facebook was fortunate to launch when it did: Silicon Valley was recovering from the dot-com bust and was entering a period of near-messianic ambitions. The Internet was no longer so new that users were scarce, but still new enough that it was largely unregulated; first movers could amass vast followings and consolidate power, and the coming rise of inexpensive smartphones would bring millions of new people online. Most important, Facebook capitalized on a resource that most people hardly knew existed: the willingness of users to subsidize the company by handing over colossal amounts of personal information, for free.

In Facebook, Zuckerberg had found the instrument to achieve his conception of greatness. His onetime speechwriter Katherine Losse, in her memoir, “The Boy Kings,” explained that the “engineering ideology of Facebook” was clear: “Scaling and growth are everything, individuals and their experiences are secondary to what is necessary to maximize the system.” Over time, Facebook devoted ever-greater focus to what is known in Silicon Valley as “growth hacking,” the constant pursuit of scale. Whenever the company talked about “connecting people,” that was, in effect, code for user growth.

Then, in 2007, growth plateaued at around fifty million users and wouldn't budge. Other social networks had maxed out at around that level, and Facebook employees wondered if they had hit a hidden limit. Zuckerberg created a special Growth Team, which had broad latitude to find ways of boosting the numbers. Among other fixes, they discovered that, by offering the site in more languages, they could open huge markets. Alex Schultz, a founding member of the Growth Team, said that he and his colleagues were fanatical in their pursuit of expansion. "You will fight for that inch, you will die for that inch," he told me. Facebook left no opportunity untapped. In 2011, the company asked the Federal Election Commission for an exemption to rules requiring the source of funding for political ads to be disclosed. In filings, a Facebook lawyer argued that the agency "should not stand in the way of innovation."

Sandy Parakilas, who joined Facebook in 2011, as an operations manager, paraphrased the message of his orientation session as "We believe in the religion of growth." He said, "The Growth Team was the coolest. Other teams would even try to call subgroups within their teams the 'Growth X' or the 'Growth Y' to try to get people excited."

To gain greater reach, Facebook had made the fateful decision to become a "platform" for outside developers, much as Windows had been in the realm of desktop computers, a generation before. The company had opened its trove of data to programmers who wanted to build Facebook games, personality tests, and other apps. After a few months at Facebook, Parakilas was put in charge of a team responsible for making sure that outsiders were not misusing the data, and he was unnerved by what he found. Some games were siphoning off users' messages and photographs. In one case, he said, a developer was harvesting user information, including that of children, to create unauthorized profiles on its own Web site. Facebook had given away data before it had a system to check for abuse. Parakilas suggested that there be an audit to uncover the scale of the problem. But, according to Parakilas, an executive rejected the idea, telling him, "Do you really want to see what you'll find?"

Parakilas told me, "It was very difficult to get the kind of resources that you needed to do a good job of insuring real compliance. Meanwhile, you looked at the Growth Team and they had engineers coming out of their ears. All the smartest minds are focussed on doing whatever they can possibly do to get those growth numbers up."

New hires learned that a crucial measure of the company's performance was how many people had logged in to Facebook on six of the previous seven days, a measurement known as L6/7. "You could say it's how many people love this service so much they use it six out of seven days," Parakilas, who left the company in 2012, said. "But, if your job is to get that number up, at some point you run out of good, purely positive ways. You start thinking about 'Well, what are the dark patterns that I can use to get people to log back in?'"

Facebook engineers became a new breed of behaviorists, tweaking levers of vanity and passion and susceptibility. The real-world effects were striking. In 2012, when Chan was in medical school, she and Zuckerberg discussed a critical shortage of organs for transplant, inspiring Zuckerberg to add a small, powerful nudge on Facebook: if people indicated that they were organ donors, it triggered a notification to friends, and, in turn, a cascade of social pressure. Researchers later found that, on the first day the feature appeared, it increased official organ-donor enrollment more than twentyfold nationwide.

Sean Parker later described the company's expertise as "exploiting a vulnerability in human psychology." The goal: "How do we consume as much of your time and conscious attention as possible?" Facebook engineers discovered that people find it nearly impossible not to log in after receiving an e-mail saying that someone has uploaded a picture of them. Facebook also discovered its power to affect people's political behavior. Researchers found that, during the 2010 midterm elections, Facebook was able to prod

users to vote simply by feeding them pictures of friends who had already voted, and by giving them the option to click on an “I Voted” button. The technique boosted turnout by three hundred and forty thousand people—more than four times the number of votes separating Trump and Clinton in key states in the 2016 race. It became a running joke among employees that Facebook could tilt an election just by choosing where to deploy its “I Voted” button.

These powers of social engineering could be put to dubious purposes. In 2012, Facebook data scientists used nearly seven hundred thousand people as guinea pigs, feeding them happy or sad posts to test whether emotion is contagious on social media. (They concluded that it is.) When the findings were published, in the *Proceedings of the National Academy of Sciences*, they caused an uproar among users, many of whom were horrified that their emotions may have been surreptitiously manipulated. In an apology, one of the scientists wrote, “In hindsight, the research benefits of the paper may not have justified all of this anxiety.”

Facebook was, in the words of Tristan Harris, a former design ethicist at Google, becoming a pioneer in “persuasive technology.” He explained, “A hammer, in your hand, is non-persuasive—it doesn’t have its own ways of manipulating the person that holds it. But Facebook and Snapchat, in their design features, are persuading a teen-ager to wake up and see photo after photo after photo of their friends having fun without them, even if it makes them feel worse.” In 2015, Harris delivered a talk at Facebook about his concern that social media was contributing to alienation. “I said, ‘You guys are in the best position in the world to deal with loneliness and see it as a thing that you are amplifying and a thing that you can help make go the other way,’” he told me. “They didn’t do anything about it.” He added, “My points were in their blind spot.”

As Facebook grew, Zuckerberg and his executives adopted a core belief: even if people criticized your decisions, they would eventually come around. In one of the first demonstrations of that idea, in 2006, Facebook introduced the News Feed, a feature that suddenly alerted friends whenever a user changed profile pictures, joined groups, or altered a relationship status. (Until then, users had to visit a friend’s page to see updates.) Users revolted. There was a street protest at the headquarters, and hundreds of thousands of people joined a Facebook group opposing the change. Zuckerberg posted a tepid apology (“Calm down. Breathe. We hear you.”), and people got used to the feed.

“A lot of the early experience for me was just having people really not believe that what we were going to do was going to work,” Zuckerberg told me. “If you think about the early narratives, it was, like, ‘Well, this was just a college thing.’ Or ‘It’s not gonna be a big deal.’ Or ‘O.K., other people are using it, but it’s kind of a fad. There’s Friendster and there’s MySpace, and there will be something after,’ or whatever.” He added, “I feel like it really tests you emotionally to have constant doubt, and the assertion that you don’t know what you are doing.”

In 2006, Zuckerberg made his most unpopular decision at the fledgling company. Yahoo was offering a billion dollars to buy Facebook and, as Matt Cohler, a top aide at the time, recalls, "Our growth had stalled out." Cohler and many others implored Zuckerberg to take the offer, but he refused. "I think nearly all of his leadership team lost faith in him and in the business," Cohler said. Zuckerberg told me that most of his leadership "left within eighteen months. Some of them I had to fire because it was just too dysfunctional. It just completely blew up. But the thing that I learned from that is, if you stick with your values and with what you believe you want to be doing in the world, you can get through. Sometimes it will take some time, and you have to rebuild, but that's a pretty powerful lesson."

On several occasions, Zuckerberg stumbled when it came to issues of privacy. In 2007, Facebook started giving advertisers a chance to buy into a program called Beacon, which would announce to a user's friends what that user was browsing for, or buying, online. Users could opt out, but many had no idea that the feature existed until it revealed upcoming holiday gifts, or, in some cases, exposed extramarital affairs. Zuckerberg apologized ("We simply did a bad job with this release, and I apologize for it," he wrote), and Beacon was withdrawn.

Despite the apology, Zuckerberg was convinced that he was ahead of his users, not at odds with them. In 2010, he said that privacy was no longer a "social norm." That year, the company found itself in trouble again after it revised its privacy controls to make most information public by default. The Federal Trade Commission cited Facebook for "engaging in unfair and deceptive practices" with regard to the privacy of user data. The company signed a consent decree pledging to establish a "comprehensive privacy program" and to evaluate it every other year for twenty years. In a post, Zuckerberg offered a qualified apology: "I think that a small number of high profile mistakes . . . have often overshadowed much of the good work we've done."

Facebook had adopted a buccaneering motto, "Move fast and break things," which celebrated the idea that it was better to be flawed and first than careful and perfect. Andrew Bosworth, a former Harvard teaching assistant who is now one of Zuckerberg's longest-serving lieutenants and a member of his inner circle, explained, "A failure can be a form of success. It's not the form you want, but it can be a useful thing to how you learn." In Zuckerberg's view, skeptics were often just fogies and scolds. "There's always someone who wants to slow you down," he said in a commencement address at Harvard last year. "In our society, we often don't do big things because we're so afraid of making mistakes that we ignore all the things wrong today if we do nothing. The reality is, anything we do will have issues in the future. But that can't keep us from starting."

Zuckerberg's disregard for criticism entered a more emphatic phase in 2010, with the release of the movie "The Social Network," an account of Facebook's early years, written by Aaron Sorkin and directed by David Fincher. Some of the film was fictionalized. It presented Zuckerberg's motivation largely as a desire to meet girls, even though, in real life, he was dating Priscilla Chan for most of the time period covered in the movie. But other elements cut close to the truth, including the depiction of his juvenile bravado and the early feuds over ownership. Zuckerberg and Facebook had chosen not to be involved in the production, and the portrayal was unflattering. Zuckerberg, played by Jesse Eisenberg, is cocksure and cold, and the real Zuckerberg found the depiction hurtful. "First impressions matter a lot, and for a lot of people that was their introduction to me," he told me. "My reaction to this, to all these things, is primarily that I perceive it through the employees." His concern was less about how people would think of him, he said, than about "how is our company, how are our employees—these people I work with and care so much about—how are they going to process this?"

Before the movie came out, Facebook executives debated how to respond. Zuckerberg settled on a stance of effortful good cheer, renting a movie theatre to screen it for the staff. Eight years later, Facebook executives still mention what they call, resentfully, “the movie.” Sandberg, who is the company’s second most important public figure, and one of Zuckerberg’s most ardent defenders, told me, “From its facts to its essence to its portrayal, I think that was a very unfair picture. I still think it forms the basis of a lot of what people believe about Mark.”

While the movie contributed to the fortress mentality on campus, Zuckerberg made a series of decisions that solidified his confidence in his instincts. In 2012, he paid a billion dollars for Instagram, the photo-sharing service, which at the time had only thirteen employees. Outside the industry, the startup appeared wildly overpriced, but it proved to be one of the best investments in the history of the Internet. (Today, Instagram is valued at more than a hundred times what Zuckerberg paid for it, and, even more important, it is popular with young people, a cohort that shows declining interest in Facebook.) That spring, Facebook went public on the Nasdaq, at a valuation of a hundred and four billion dollars. There were technical glitches on the day of the listing, and many people doubted that the company could earn enough money to justify the valuation. The share price promptly sank. The *Wall Street Journal* called the I.P.O. a “fiasco,” and shareholders sued Facebook and Zuckerberg. “We got a ton of criticism,” he recalled. “Our market cap got cut in half. But what I felt was, we were at a sufficient skill and complexity that it was going to take a couple years to work through the problem, but I had strong conviction that we were doing the right thing.” (Even with its recent plunge, the value of Facebook stock has more than quadrupled in the years since.)

Zuckerberg was happy to make sharp turns to achieve his aims. In 2011, when users started moving from desktop computers to phones, Facebook swerved toward mobile technology. Zuckerberg told employees that he would kick them out of his office if their ideas did not account for the transition. “Within a month, you literally can’t meet with Mark if you’re not bringing him a mobile product,” Bosworth recalled.

In 2014, as problems accumulated, Facebook changed its motto, “Move fast and break things,” to the decidedly less glamorous “Move fast with stable infrastructure.” Still, internally, much of the original spirit endured, and the push for haste began to take a toll in the offline world. In early 2016, Zuckerberg directed employees to accelerate the release of Facebook Live, a video-streaming service, and expanded its team of engineers from twelve to more than a hundred. When the product emerged, two months later, so did unforeseen issues: the service let users flag videos as inappropriate, but it didn’t give them a way to indicate where in a broadcast the problem appeared. As a result, Facebook Live videos of people committing suicide, or engaged in criminal activity, started circulating before reviewers had time to race through, find the issues, and take the videos down. A few months after the service launched, a Chicago man named Antonio Perkins was fatally shot on Facebook Live and the video was viewed hundreds of thousands of times.

The incident might have served as a warning to slow down, but, instead, the next day, Bosworth sent around a remarkable internal memo justifying some of Facebook’s “ugly” physical and social effects as the trade-offs necessary for growth: “Maybe it costs a life by exposing someone to bullies. Maybe someone dies in a terrorist attack coordinated on our tools. And still we connect people. The ugly truth is that we believe in connecting people so deeply that anything that allows us to connect more people more often is *de facto* good.”

This spring, after the memo leaked to BuzzFeed, Bosworth said that he had been playing devil’s advocate, and Zuckerberg issued a statement: “Boz is a talented leader who says many provocative things. This was one that most people at Facebook including

myself disagreed with strongly. We've never believed the ends justify the means."

Zuckerberg was also experimenting with philanthropy. In 2010, shortly before the release of "The Social Network," he made a high-profile gift. Appearing onstage at "The Oprah Winfrey Show," along with Chris Christie, the governor of New Jersey, and Cory Booker, the mayor of Newark, he announced a hundred-million-dollar donation to help Newark's struggling public-school system. The project quickly encountered opposition from local groups that saw it as out of touch, and, eight years later, it's generally considered a failure. In May, Ras Baraka, Newark's mayor, said of the donation, "You can't just cobble up a bunch of money and drop it in the middle of the street and say, 'This is going to fix everything.'"

For all the criticism, the project has produced some measurable improvements. A Harvard study found greater gains in English than the state average, and a study by MarGrady Research, an education-policy group, found that high-school graduation rates and over-all student enrollment in Newark have risen since the donation. Zuckerberg emphasizes those results, even as he acknowledges flaws in his approach. "Your earning potential is dramatically higher if you graduate from high school versus not. That part of it, I think, is the part that worked and it was effective," he said. "There were a bunch of other things that we tried that either were much harder than we thought or just didn't work." Strategies that helped him in business turned out to hurt him in education reform. "I think in a lot of philanthropy and government-related work, if you try five things and a few of them fail, then the ones that fail are going to get a lot of the attention," he said.

In 2015, Zuckerberg and Chan pledged to spend ninety-nine per cent of their Facebook fortune "to advance human potential and promote equality for all children in the next generation." They created the Chan Zuckerberg Initiative, a limited-liability company that gives to charity, invests in for-profit companies, and engages in political advocacy. David Plouffe said that the lessons of the Newark investment shaped the initiative's perspective. "I think the lesson was, you have to do this in full partnership with the community, not just the leaders," he said. "You need to have enthusiastic buy-in from superintendents, and teachers, and parents."

In contrast to a traditional foundation, an L.L.C. can lobby and give money to politicians, without as strict a legal requirement to disclose activities. In other words, rather than trying to win over politicians and citizens in places like Newark, Zuckerberg and Chan could help elect politicians who agree with them, and rally the public directly by running ads and supporting advocacy groups. (A spokesperson for C.Z.I. said that it has given no money to candidates; it has supported ballot initiatives through a 501(c)(4) social-welfare organization.) "The whole point of the L.L.C. structure is to allow a coördinated attack," Rob Reich, a co-director of Stanford's Center on Philanthropy and Civil Society, told me. The structure has gained popularity in Silicon Valley but has been criticized for allowing wealthy individuals to orchestrate large-scale social agendas behind closed doors. Reich said, "There should be much greater transparency, so that it's not dark. That's not a criticism of Mark Zuckerberg. It's a criticism of the law."

In 2016, Zuckerberg announced, onstage and in a Facebook post, his intention to "help cure all disease in our children's lifetime." That was partly bluster: C.Z.I. is working on a slightly more realistic agenda, to "cure, prevent or manage all diseases." The theatrics irritated some in the philanthropy world who thought that Zuckerberg's presentation minimized the challenges, but, in general, scientists have applauded the ambition. When I asked Zuckerberg about the reception of the project, he said, "It's funny, when I talk to people here in the Valley, you get a couple of reactions. A bunch of people have the reaction of 'Oh, that's obviously going to happen on its own—why don't you just spend your time doing something else?' And then a bunch of people have the reaction of 'Oh, that seems almost impossible—why are you setting your sights so high?'"

Characteristically, Zuckerberg favors the optimistic scenario. "On average, every year for the last eighty years or so, I think, life expectancy has gone up by about a quarter of a year. And, if you believe that technological and scientific progress is not going to slow, there is a potential upside to speeding that up," he said. "We're going to get to a point where the life expectancy implied by extrapolating that out will mean that we'll basically have been able to manage or cure all of the major things that people suffer from and die from today. Based on the data that we already see, it seems like there's a reasonable shot."

I asked Bill Gates, whose private foundation is the largest in the U.S., about Zuckerberg's objectives. "There are aspirations and then there are plans," he said. "And plans vary in terms of their degree of realism and concreteness." He added that Zuckerberg's long-range goal is "very safe, because you will not be around to write the article saying that he overcommitted."

As Facebook expanded, so did its blind spots. The company's financial future relies partly on growth in developing countries, but the platform has been a powerful catalyst of violence in fragile parts of the globe. In India, the largest market for Facebook's WhatsApp service, hoaxes have triggered riots, lynchings, and fatal beatings. Local officials resorted to shutting down the Internet sixty-five times last year. In Libya, people took to Facebook to trade weapons, and armed groups relayed the locations of targets for artillery strikes. In Sri Lanka, after a Buddhist mob attacked Muslims this spring over a false rumor, a Presidential adviser told the *Times*, "The germs are ours, but Facebook is the wind."

Nowhere has the damage been starker than in Myanmar, where the Rohingya Muslim minority has been subject to brutal killings, gang rapes, and torture. In 2012, around one per cent of the country's population had access to the Internet. Three years later, that figure had reached twenty-five per cent. Phones often came preloaded with the Facebook app, and Buddhist extremists seeking to inflame ethnic tensions with the Rohingya mastered the art of misinformation. Wirathu, a monk with a large Facebook following, sparked a deadly riot against Muslims in 2014 when he shared a fake report of a rape and warned of a "Jihad against us." Others gamed Facebook's rules against hate speech by fanning paranoia about demographic change. Although Muslims make up no more than five per cent of the country, a popular graphic appearing on Facebook cautioned that "when Muslims become the most powerful" they will offer "Islam or the sword."

Beginning in 2013, a series of experts on Myanmar met with Facebook officials to warn them that it was fuelling attacks on the Rohingya. David Madden, an entrepreneur based in Myanmar, delivered a presentation to officials at the Menlo Park headquarters, pointing out that the company was playing a role akin to that of the radio broadcasts that spread hatred during the Rwandan genocide. In 2016, C4ADS, a Washington-based nonprofit, published a detailed analysis of Facebook usage in Myanmar, and described a "campaign of hate speech that actively dehumanizes Muslims." Facebook officials said that they were

hiring more Burmese-language reviewers to take down dangerous content, but the company repeatedly declined to say how many had actually been hired. By last March, the situation had become dire: almost a million Rohingya had fled the country, and more than a hundred thousand were confined to internal camps. The United Nations investigator in charge of examining the crisis, which the U.N. has deemed a genocide, said, "I'm afraid that Facebook has now turned into a beast, and not what it was originally intended." Afterward, when pressed, Zuckerberg repeated the claim that Facebook was "hiring dozens" of additional Burmese-language content reviewers.

More than three months later, I asked Jes Kaliebe Petersen, the C.E.O. of Phandeeeyar, a tech hub in Myanmar, if there had been any progress. "We haven't seen any tangible change from Facebook," he told me. "We don't know how much content is being reported. We don't know how many people at Facebook speak Burmese. The situation is getting worse and worse here."

I saw Zuckerberg the following morning, and asked him what was taking so long. He replied, "I think, fundamentally, we've been slow at the same thing in a number of areas, because it's actually the same problem. But, yeah, I think the situation in Myanmar is terrible." It was a frustrating and evasive reply. I asked him to specify the problem. He said, "Across the board, the solution to this is we need to move from what is fundamentally a reactive model to a model where we are using technical systems to flag things to a much larger number of people who speak all the native languages around the world and who can just capture much more of the content."

I told him that people in Myanmar are incredulous that a company with Facebook's resources has failed to heed their complaints. "We're taking this seriously," he said. "You can't just snap your fingers and solve these problems. It takes time to hire the people and train them, and to build the systems that can flag stuff for them." He promised that Facebook would have "a hundred or more Burmese-speaking people by the end of the year," and added, "I hate that we're in this position where we are not moving as quickly as we would like." A few weeks after our conversation, Facebook announced that it was banning Myanmar's Army chief and several other military officials.

Over the years, Zuckerberg had come to see his ability to reject complaints as a virtue. But, by 2016, that stance had primed the company for a crisis. Tristan Harris, the design ethicist, said, "When you're running anything like Facebook, you get criticized all the time, and you just stop paying attention to criticism if a lot of it is not well founded. You learn to treat it as naïve and uninformed." He went on, "The problem is it also puts you out of touch with genuine criticism from people who actually understand the issues."

The 2016 election was supposed to be good for Facebook. That January, Sheryl Sandberg told investors that the election would be "a big deal in terms of ad spend," comparable to the Super Bowl and the World Cup. According to Borrell Associates, a research and consulting firm, candidates and other political groups were on track to spend \$1.4 billion online in the election, up ninefold from four years earlier.

Facebook offered to "embed" employees, for free, in Presidential campaign offices to help them use the platform effectively. Clinton's campaign said no. Trump's said yes, and Facebook employees helped his campaign craft messages. Although Trump's language was openly hostile to ethnic minorities, inside Facebook his behavior felt, to some executives, like just part of the distant cesspool of Washington. Americans always seemed to be choosing between a hated Republican and a hated Democrat, and Trump's descriptions of Mexicans as rapists was simply an extension of that.

During the campaign, Trump used Facebook to raise two hundred and eighty million dollars. Just days before the election, his team paid for a voter-suppression effort on the platform. According to *Bloomberg Businessweek*, it targeted three Democratic constituencies—“idealistic white liberals, young women, and African Americans”—sending them videos precisely tailored to discourage them from turning out for Clinton. Theresa Hong, the Trump campaign’s digital-content director, later told an interviewer, “Without Facebook we wouldn’t have won.”

After the election, Facebook executives fretted that the company would be blamed for the spread of fake news. Zuckerberg’s staff gave him statistics showing that the vast majority of election information on the platform was legitimate. At a tech conference a few days later, Zuckerberg was defensive. “The idea that fake news on Facebook—of which, you know, it’s a very small amount of the content—influenced the election in any way, I think, is a pretty crazy idea,” he said. To some at Facebook, Zuckerberg’s defensiveness was alarming. A former executive told *Wired*, “We had to really flip him on that. We realized that if we didn’t, the company was going to start heading down this pariah path.”

When I asked Zuckerberg about his “pretty crazy” comment, he said that he was wrong to have been “glib.” He told me, “Nobody wants any amount of fake news. It is an issue on an ongoing basis, and we need to take that seriously.” But he still bristles at the implication that Facebook may have distorted voter behavior. “I find the notion that people would only vote some way because they were tricked to be almost viscerally offensive,” he said. “Because it goes against the whole notion that you should trust people and that individuals are smart and can understand their own experience and can make their own assessments about what direction they want their community to go in.”

Shortly after the election, Mark Warner, the ranking Democrat on the Senate Intelligence Committee, contacted Facebook to discuss Russian interference. “The initial reaction was completely dismissive,” he told me. But, by the spring, he sensed that the company was realizing that it had a serious problem. “They were seeing an enormous amount of Russian activity in the French elections,” Warner said. “It was getting better, but I still don’t think they were putting nearly enough resources behind this.” Warner, who made a fortune in the telecom business, added, “Most of the companies in the Valley think that policymakers, one, don’t get it, and, two, that ultimately, if they just stonewall us, then we’ll go away.”

Facebook moved fitfully to acknowledge the role it had played in the election. In September of 2017, after Robert Mueller obtained a search warrant, Facebook agreed to give his office an inventory of ads linked to Russia and the details of who had paid for them. In October, Facebook disclosed that Russian operatives had published about eighty thousand posts, reaching a hundred and twenty-six million Americans.

In March, after the Cambridge Analytica news broke, Zuckerberg and Facebook were paralyzed. For five days, Zuckerberg said nothing. His personal Facebook profile offered no statements or analysis. Its most recent post was a photo of him and Chan baking hamantaschen for Purim.

“I feel like we’ve let people down and that feels terrible,” he told me later. “But it goes back to this notion that we shouldn’t be making the same mistake multiple times.” He insists that fake news is less common than people imagine: “The average person might perceive, from how much we and others talk about it, that there is more than ten times as much misinformation or hoax content on Facebook than the academic measures that we’ve seen so far suggest.” He is still not convinced that the spread of misinformation had an impact on the election. “I actually don’t consider that a closed thing,” he said. “I still think that’s the kind of thing that needs to be studied.”

In conversation, Zuckerberg is, unsurprisingly, highly analytical. When he encounters a theory that doesn't accord with his own, he finds a seam of disagreement—a fact, a methodology, a premise—and hammers at it. It's an effective technique for winning arguments, but one that makes it difficult to introduce new information. Over time, some former colleagues say, his deputies have begun to filter out bad news from presentations before it reaches him. A former Facebook official told me, "They only want to hear good news. They don't want people who are disagreeing with them. There is a culture of 'You go along to get along.'"

I once asked Zuckerberg what he reads to get the news. "I probably mostly read aggregators," he said. "I definitely follow Techmeme"—a roundup of headlines about his industry—"and the media and political equivalents of that, just for awareness." He went on, "There's really no newspaper that I pick up and read front to back. Well, that might be true of most people these days—most people don't read the physical paper—but there aren't many news Web sites where I go to browse."

A couple of days later, he called me and asked to revisit the subject. "I felt like my answers were kind of vague, because I didn't necessarily feel like it was appropriate for me to get into which specific organizations or reporters I read and follow," he said. "I guess what I tried to convey, although I'm not sure if this came across clearly, is that the job of uncovering new facts and doing it in a trusted way is just an absolutely critical function for society."

Zuckerberg and Sandberg have attributed their mistakes to excessive optimism, a blindness to the darker applications of their service. But that explanation ignores their fixation on growth, and their unwillingness to heed warnings. Zuckerberg resisted calls to reorganize the company around a new understanding of privacy, or to reconsider the depth of data it collects for advertisers.

James P. Steyer, the founder and C.E.O. of Common Sense Media, an organization that promotes safety in technology and media for children, visited Facebook's headquarters in the spring of 2018 to discuss his concerns about a product called Messenger Kids, which allows children under thirteen—the minimum age to use the primary Facebook app—to make video calls and send messages to contacts that a parent approves. He met with Sandberg and Elliot Schrage, at the time the head of policy and communications. "I respect their business success, and like Sheryl personally, and I was hoping they might finally consider taking steps to better protect kids. Instead, they said that the best thing for young kids was to spend more time on Messenger Kids," Steyer told me. "They still seemed to be in denial. Would *you* 'move fast and break things' when it comes to children? To our democracy? No, because you can damage them forever."

To some people in the company, the executives seemed concentrated not on solving the problems or on preventing the next ones but on containing the damage. Tavis McGinn, a former Google pollster, started working at Facebook in the spring of 2017, doing

polls with a narrow focus: measuring the public perception of Zuckerberg and Sandberg. During the next six months, McGinn conducted eight surveys and four focus groups in three countries, collecting the kinds of measurements favored by politicians and advertisers. Facebook polled reactions to the company's new stated mission to "bring the world closer together," as well as to items on Zuckerberg's social-media feed, including his writings, photographs, and even his casual banter during a back-yard barbecue broadcast on Facebook Live.

In September, McGinn resigned. In an interview, he told the Web site the Verge that he had become discouraged. "I was not going to be able to change the way that the company does business," he said. "I couldn't change the values. I couldn't change the culture." He concluded that measuring the "true social outcomes" of Facebook was of limited interest to senior staffers. "I think research can be very powerful, if people are willing to listen," he said. "But I decided after six months that it was a waste of my time to be there. I didn't feel great about the product. I didn't feel proud to tell people I worked at Facebook. I didn't feel I was helping the world." (McGinn, who has signed a nondisclosure agreement with Facebook, declined to comment for this article.)

In March, Zuckerberg agreed to testify before Congress for the first time about Facebook's handling of user data. The hearing was scheduled for April. As the date approached, the hearing acquired the overtones of a trial.

In barely two years, the mood in Washington had shifted. Internet companies and entrepreneurs, formerly valorized as the vanguard of American ingenuity and the astronauts of our time, were being compared to Standard Oil and other monopolists of the Gilded Age. This spring, the *Wall Street Journal* published [an article](#) that began, "Imagine a not-too-distant future in which trustbusters force Facebook to sell off Instagram and WhatsApp." It was accompanied by a sepia-toned illustration in which portraits of Zuckerberg, Tim Cook, and other tech C.E.O.s had been grafted onto overstuffed torsos meant to evoke the robber barons. In 1915, Louis Brandeis, the reformer and future Supreme Court Justice, testified before a congressional committee about the dangers of corporations large enough that they could achieve a level of near-sovereignty "so powerful that the ordinary social and industrial forces existing are insufficient to cope with it." He called this the "curse of bigness." Tim Wu, a Columbia law-school professor and the author of a forthcoming book inspired by Brandeis's phrase, told me, "Today, no sector exemplifies more clearly the threat of bigness to democracy than Big Tech." He added, "When a concentrated private power has such control over what we see and hear, it has a power that rivals or exceeds that of elected government."

Shortly before Zuckerberg was due to testify, a team from the Washington law firm of WilmerHale flew to Menlo Park to run him through mock hearings and to coach him on the requisite gestures of humility. Even before the recent scandals, Bill Gates had advised Zuckerberg to be alert to the opinions of lawmakers, a lesson that Gates had learned in 1998, when Microsoft faced accusations of monopolistic behavior. Gates testified to Congress, defiantly, that "the computer-software industry is not broken, and there is no need to fix it." Within months, the Department of Justice sued Microsoft for violating federal antitrust law, leading to three years of legal agony before a settlement was reached. Gates told me that he regretted "taunting" regulators, saying, "Not something I would choose to repeat." He encouraged Zuckerberg to be attentive to D.C. "I said, 'Get an office there—now.' And Mark did, and he owes me," Gates said. Last year, Facebook spent \$11.5 million on lobbying in Washington, ranking it between the American Bankers Association and General Dynamics among top spenders.

On April 10th, when Zuckerberg arrived at the [Senate hearing](#), he wore a sombre blue suit, and took a seat before more than forty senators. In front of him, his notes outlined likely questions and answers, including the prospect that a senator might ask him to

step down from the company. His answer, in shorthand, would be: “Founded Facebook. My decisions. I made mistakes. Big challenge, but we’ve solved problems before, going to solve this one. Already taking action.”

As it turned out, nobody asked him to resign—or much of anything difficult. Despite scattered moments of pressure, the overwhelming impression left by the event was how poorly some senators grasped the issues. In the most revealing moment, Orrin Hatch, the eighty-four-year-old Republican from Utah, demanded to know how Facebook makes money if “users don’t pay for your service.” Zuckerberg replied, “Senator, we run ads,” allowing a small smile.

To observers inclined to distrust Zuckerberg, he was evasive to the point of amnesiac—he said, more than forty times, that he would need to follow up—but when the hearing concluded, after five hours, he had emerged unscathed, and Wall Street, watching closely, rewarded him by boosting the value of Facebook’s stock by twenty billion dollars. A few days later, on the internal Facebook message board, an employee wrote that he planned to buy T-shirts reading “Senator, we run ads.”

When I asked Zuckerberg whether policymakers might try to break up Facebook, he replied, adamantly, that such a move would be a mistake. The field is “extremely competitive,” he told me. “I think sometimes people get into this mode of ‘Well, there’s not, like, an exact replacement for Facebook.’ Well, actually, that makes it *more* competitive, because what we really are is a system of different things: we compete with Twitter as a broadcast medium; we compete with Snapchat as a broadcast medium; we do messaging, and iMessage is default-installed on every iPhone.” He acknowledged the deeper concern. “There’s this other question, which is just, laws aside, how do we feel about these tech companies being big?” he said. But he argued that efforts to “curtail” the growth of Facebook or other Silicon Valley heavyweights would cede the field to China. “I think that anything that we’re doing to constrain them will, first, have an impact on how successful we can be in other places,” he said. “I wouldn’t worry in the near term about Chinese companies or anyone else winning in the U.S., for the most part. But there are all these places where there are day-to-day more competitive situations—in Southeast Asia, across Europe, Latin America, lots of different places.”

The rough consensus in Washington is that regulators are unlikely to try to break up Facebook. The F.T.C. will almost certainly fine the company for violations, and may consider blocking it from buying big potential competitors, but, as a former F.T.C. commissioner told me, “in the United States you’re allowed to have a monopoly position, as long as you achieve it and maintain it without doing illegal things.”

Facebook is encountering tougher treatment in Europe, where antitrust laws are stronger and the history of fascism makes people especially wary of intrusions on privacy. One of the most formidable critics of Silicon Valley is the European Union’s top antitrust regulator, Margrethe Vestager. Last year, after an investigation of Google’s search engine, Vestager accused the company of giving an “illegal advantage” to its shopping service and fined it \$2.7 billion, at that time the largest fine ever imposed by the E.U. in an antitrust case. In July, she added another five-billion-dollar fine for the company’s practice of requiring device makers to preinstall Google apps.

In Brussels, Vestager is a high-profile presence—nearly six feet tall, with short black-and-silver hair. She grew up in rural Denmark, the eldest child of two Lutheran pastors, and, when I spoke to her recently, she talked about her enforcement powers in philosophical terms. “What we’re dealing with, when people start doing something illegal, is exactly as old as Adam and Eve,” she said. “Human decisions very often are guided by greed, by fear of being pushed out of the marketplace, or of losing something

that's important to you. And then, if you throw power into that cocktail of greed and fear, you have something that you can recognize throughout time."

Vestager told me that her office has no open cases involving Facebook, but she expressed concern that the company was taking advantage of users, beginning with terms of service that she calls "unbalanced." She paraphrased those terms as "It's your data, but you give us a royalty-free global license to do, basically, whatever we want." Imagine, she said, if a brick-and-mortar business asked to copy all your photographs for its unlimited, unspecified uses. "Your children, from the very first day until the confirmation, the rehearsal dinner for the wedding, the wedding itself, the first child being baptized. You would never accept that," she said. "But this is what you accept without a blink of an eye when it's digital."

In Vestager's view, a healthy market should produce competitors to Facebook that position themselves as ethical alternatives, collecting less data and seeking a smaller share of user attention. "We need social media that will allow us to have a nonaddictive, advertising-free space," she said. "You're more than welcome to be successful and to dramatically outgrow your competitors if customers like your product. But, if you grow to be dominant, you have a special responsibility not to misuse your dominant position to make it very difficult for others to compete against you and to attract potential customers. Of course, we keep an eye on it. If we get worried, we will start looking."

As the pressure on Facebook has intensified, the company has been moving to fix its vulnerabilities. In December, after Sean Parker and Chamath Palihapitiya spoke publicly about the damaging psychological effects of social media, Facebook acknowledged evidence that heavy use can exacerbate anxiety and loneliness. After years of perfecting addictive features, such as "auto-play" videos, it announced a new direction: it would promote the quality, rather than the quantity, of time spent on the site. The company modified its algorithm to emphasize updates from friends and family, the kind of content most likely to promote "active engagement." In a post, Zuckerberg wrote, "We can help make sure that Facebook is time well spent."

The company also grappled with the possibility that it would once again become a vehicle for election-season propaganda. In 2018, hundreds of millions of people would be voting in elections around the world, including in the U.S. midterms. After years of lobbying against requirements to disclose the sources of funding for political ads, the company announced that users would now be able to look up who paid for a political ad, whom the ad targeted, and which other ads the funders had run.

Samidh Chakrabarti, the product manager in charge of Facebook's "election integrity" work, told me that the revelations about Russia's Internet Research Agency were deeply alarming. "This wasn't the kind of product that any of us thought that we were

working on,” he said. With the midterms approaching, the company had discovered that Russia’s model for exploiting Facebook had inspired a generation of new actors similarly focussed on skewing political debate. “There are lots of copycats,” Chakrabarti said.

Zuckerberg used to rave about the virtues of “frictionless sharing,” but these days Facebook is working on “imposing friction” to slow the spread of disinformation. In January, the company hired Nathaniel Gleicher, the former director for cybersecurity policy on President Obama’s National Security Council, to blunt “information operations.” In July, it removed thirty-two accounts running disinformation campaigns that were traced to Russia. A few weeks later, it removed more than six hundred and fifty accounts, groups, and pages with links to Russia or Iran. Depending on your point of view, the removals were a sign either of progress or of the growing scale of the problem. Regardless, they highlighted the astonishing degree to which the security of elections around the world now rests in the hands of Gleicher, Chakrabarti, and other employees at Facebook.

As hard as it is to curb election propaganda, Zuckerberg’s most intractable problem may lie elsewhere—in the struggle over which opinions can appear on Facebook, which cannot, and who gets to decide. As an engineer, Zuckerberg never wanted to wade into the realm of content. Initially, Facebook tried blocking certain kinds of material, such as posts featuring nudity, but it was forced to create long lists of exceptions, including images of breast-feeding, “acts of protest,” and works of art. Once Facebook became a venue for political debate, the problem exploded. In April, in a call with investment analysts, Zuckerberg said glumly that it was proving “easier to build an A.I. system to detect a nipple than what is hate speech.”

The cult of growth leads to the curse of bigness: every day, a billion things were being posted to Facebook. At any given moment, a Facebook “content moderator” was deciding whether a post in, say, Sri Lanka met the standard of hate speech or whether a dispute over Korean politics had crossed the line into bullying. Zuckerberg sought to avoid banning users, preferring to be a “platform for all ideas.” But he needed to prevent Facebook from becoming a swamp of hoaxes and abuse. His solution was to ban “hate speech” and impose lesser punishments for “misinformation,” a broad category that ranged from crude deceptions to simple mistakes. Facebook tried to develop rules about how the punishments would be applied, but each idiosyncratic scenario prompted more rules, and over time they became byzantine. According to Facebook training slides published by the *Guardian* last year, moderators were told that it was permissible to say “You are such a Jew” but not permissible to say “Irish are the best, but really French sucks,” because the latter was defining another people as “inferiors.” Users could not write “Migrants are scum,” because it is dehumanizing, but they could write “Keep the horny migrant teen-agers away from our daughters.” The distinctions were explained to trainees in arcane formulas such as “Not Protected + Quasi protected = not protected.”

In July, the issue landed, inescapably, in Zuckerberg’s lap. For years, Facebook had provided a platform to the conspiracy theorist Alex Jones, whose delusions include that the parents of children killed in the Sandy Hook school massacre are paid actors with an anti-gun agenda. Facebook was loath to ban Jones. When people complained that his rants violated rules against harassment and fake news, Facebook experimented with punishments. At first, it “reduced” him, tweaking the algorithm so that his messages would be shown to fewer people, while feeding his fans articles that fact-checked his assertions.

Then, in late July, Leonard Pozner and Veronique De La Rosa, the parents of Noah Pozner, a child killed at Sandy Hook, published an open letter addressed “Dear Mr Zuckerberg,” in which they described “living in hiding” because of death threats from conspiracy theorists, after “an almost inconceivable battle with Facebook to provide us with the most basic of protections.” In

their view, Zuckerberg had “deemed that the attacks on us are immaterial, that providing assistance in removing threats is too cumbersome, and that our lives are less important than providing a safe haven for hate.”

Facebook relented, somewhat. On July 27th, it took down four of Jones’s videos and suspended him for a month. But public pressure did not let up. On August 5th, the dam broke after Apple, saying that the company “does not tolerate hate speech,” stopped distributing five podcasts associated with Jones. Facebook shut down four of Jones’s pages for “repeatedly” violating rules against hate speech and bullying. I asked Zuckerberg why Facebook had wavered in its handling of the situation. He was prickly about the suggestion: “I don’t believe that it is the right thing to ban a person for saying something that is factually incorrect.”

Jones seemed a lot more than factually incorrect, I said.

“O.K., but I think the facts here are pretty clear,” he said, homing in. “The *initial* questions were around misinformation.” He added, “We don’t take it down and ban people unless it’s directly inciting violence.” He told me that, after Jones was reduced, more complaints about him flooded in, alerting Facebook to older posts, and that the company was debating what to do when Apple announced its ban. Zuckerberg said, “When they moved, it was, like, O.K., we shouldn’t just be sitting on this content and these enforcement decisions. We should move on what we know violates the policy. We need to make a decision now.”

It will hardly be the last quandary of this sort. Facebook’s free-speech dilemmas have no simple answers—you don’t have to be a fan of Alex Jones to be unnerved by the company’s extraordinary power to silence a voice when it chooses, or, for that matter, to amplify others, to pull the levers of what we see, hear, and experience. Zuckerberg is hoping to erect a scalable system, an orderly decision tree that accounts for every eventuality and exception, but the boundaries of speech are a bedeviling problem that defies mechanistic fixes. The Supreme Court, defining obscenity, landed on “I know it when I see it.” For now, Facebook is making do with a Rube Goldberg machine of policies and improvisations, and opportunists are relishing it. Senator Ted Cruz, Republican of Texas, seized on the ban of Jones as a fascist assault on conservatives. In a moment that was rich even by Cruz’s standards, he quoted Martin Niemöller’s famous lines about the Holocaust, saying, “As the poem goes, you know, ‘First they came for Alex Jones.’”

Last week, testifying before the Senate Intelligence Committee, Sandberg said, “We now have over twenty thousand people and we are able to review reports in fifty languages, twenty-four hours a day.” (In the hallway before the session, as if to underscore the complexities ahead, Alex Jones had made a scene, heckling Senator Marco Rubio for not doing more to get him back on Facebook.) In recent years, Sandberg has gained fame for her work outside the company, including her books, “Lean In,” a best-seller on women’s empowerment, and “Option B,” which she wrote after the sudden death of her husband, Dave Goldberg. But her responsibility for the Facebook fallout is likely to grow, and her reputation as C.O.O. hinges on the implementation of changes. Many people in Silicon Valley believe that Sandberg and Facebook’s board of directors must do more to prevent the company from making another major mistake. “I know a couple of guys who are color-blind,” a prominent executive told me, “and their wives lay their clothes out for them in the morning if they don’t want to go out every day looking like Bozo the Clown. Sheryl and the board are expected to lay the clothes out for Mark.” He went on, “If you have blind spots, then you rely on the people around you to tell you where they are.”

In one of our conversations, I asked Zuckerberg whether he finds it insulting when people speculate that he lacks emotions. “Insulting?” he asked, and then paused for several seconds to consider. “I don’t find it insulting. I don’t think it’s *accurate*. I

mean, I definitely care a lot. There's a difference between letting emotions drive impulsive decisions and caring." He went on, "Ultimately, I think the reason that we built this successful thing is because we just solve problem after problem after problem, and typically you don't do that by making impulsive, emotional decisions."

The caricature of Zuckerberg is that of an automaton with little regard for the human dimensions of his work. The truth is something else: he decided long ago that no historical change is painless. Like Augustus, he is at peace with his trade-offs. Between speech and truth, he chose speech. Between speed and perfection, he chose speed. Between scale and safety, he chose scale. His life thus far has convinced him that he can solve "problem after problem after problem," no matter the howling from the public it may cause.

At a certain point, the habits of mind that served Zuckerberg well on his ascent will start to work against him. To avoid further crises, he will have to embrace the fact that he's now a protector of the peace, not a disrupter of it. Facebook's colossal power of persuasion has delivered fortune but also peril. Like it or not, Zuckerberg is a gatekeeper. The era when Facebook could learn by doing, and fix the mistakes later, is over. The costs are too high, and idealism is not a defense against negligence.

In some sense, the "Mark Zuckerberg production"—as he called Facebook in its early years—has only just begun. Zuckerberg is not yet thirty-five, and the ambition with which he built his empire could well be directed toward shoring up his company, his country, and his name. The question is not whether Zuckerberg has the power to fix Facebook but whether he has the will; whether he will kick people out of his office—with the gusto that he once mustered for the pivot to mobile—if they don't bring him ideas for preventing violence in Myanmar, or protecting privacy, or mitigating the toxicity of social media. He succeeded, long ago, in making Facebook great. The challenge before him now is to make it good. ♦

Published in the print edition of the September 17, 2018, issue, with the headline "Ghost in the Machine."



Evan Osnos is a staff writer at The New Yorker. His most recent book is "Wildland: The Making of America's Fury."

More: [Mark Zuckerberg](#) [Facebook](#) [Technology](#) [Silicon Valley](#) [Social Media](#) [Tech Companies](#) [Chief Executive Officers \(C.E.O.s\)](#)

[2016 Election](#)

THIS WEEK'S ISSUE

Never miss a big *New Yorker* story again. Sign up for This Week's Issue and get an e-mail every week with the stories you have to read.




Illustration by Shawna X; Animation by Robby Abaya

DECADE IN REVIEW

THE AGE OF INSTAGRAM FACE

How social media, FaceTune, and plastic surgery created a single, cyborgian look.

By Jia Tolentino

December 12, 2019

This past summer, I booked a plane ticket to Los Angeles with the hope of investigating what seems likely to be one of the oddest legacies of our rapidly expiring decade: the gradual

emergence, among professionally beautiful women, of a single, cyborgian face. It's a young face, of course, with poreless skin and plump, high cheekbones. It has catlike eyes and long, cartoonish lashes; it has a small, neat nose and full, lush lips. It looks at you coyly but blankly, as if its owner has taken half a Klonopin and is considering asking you for a private-jet ride to Coachella. The face is distinctly white but ambiguously ethnic—it suggests a *National Geographic* composite illustrating what Americans will look like in 2050, if every American of the future were to be a direct descendant of Kim Kardashian West, Bella Hadid, Emily Ratajkowski, and Kendall Jenner (who looks exactly like Emily Ratajkowski). “It’s like a sexy ... baby ... tiger,” Cara Craig, a high-end New York colorist, observed to me recently. The celebrity makeup artist Colby Smith told me, “It’s Instagram Face, duh. It’s like an unrealistic sculpture. Volume on volume. A face that looks like it’s made out of clay.”

Instagram, which launched as the decade was just beginning, in October, 2010, has its own aesthetic language: the ideal image is always the one that instantly pops on a phone screen. The aesthetic is also marked by a familiar human aspiration, previously best documented in wedding photography, toward a generic sameness. Accounts such as Insta Repeat illustrate the platform’s monotony by posting grids of indistinguishable photos posted by different users—a person in a yellow raincoat standing at the base of a waterfall, or a hand holding up a bright fall leaf. Some things just perform well.

The human body is an unusual sort of Instagram subject: it can be adjusted, with the right kind of effort, to perform better and better over time. Art directors at magazines have long edited photos of celebrities to better match unrealistic beauty standards; now you can do that to pictures of yourself with just a few taps on your phone. Snapchat, which launched in 2011 and was originally known as a purveyor of disappearing messages, has maintained its user base in large part by providing photo filters, some of which allow you to become intimately familiar with what your face would look like if it were ten-per-cent more conventionally attractive—if it were thinner, or had smoother skin, larger eyes, fuller lips. Instagram has added an array of flattering selfie filters to its Stories feature.

FaceTune, which was released in 2013 and promises to help you “wow your friends with every selfie,” enables even more precision. A number of Instagram accounts are dedicated to identifying the tweaks that celebrities make to their features with photo-editing apps. Celeb Face, which has more than a million followers, posts photos from the accounts of celebrities, adding arrows to spotlight signs of careless FaceTuning. Follow Celeb Face for a month, and this constant perfecting process begins to

seem both mundane and pathological. You get the feeling that these women, or their assistants, alter photos out of a simple defensive reflex, as if FaceTuning your jawline were the Instagram equivalent of checking your eyeliner in the bathroom of the bar.

“I think ninety-five per cent of the most-followed people on Instagram use FaceTune, easily,” Smith told me. “And I would say that ninety-five per cent of these people have also had some sort of cosmetic procedure. You can see things getting trendy—like, everyone’s getting brow lifts via Botox now. Kylie Jenner didn’t used to have that sort of space around her eyelids, but now she does.”

Twenty years ago, plastic surgery was a fairly dramatic intervention: expensive, invasive, permanent, and, often, risky. But, in 2002, the Food and Drug Administration approved Botox for use in preventing wrinkles; a few years later, it approved hyaluronic-acid fillers, such as Juvéderm and Restylane, which at first filled in fine lines and wrinkles and now can be used to restructure jawlines, noses, and cheeks. These procedures last for six months to a year and aren’t nearly as expensive as surgery. (The average price per syringe of filler is six hundred and eighty-three dollars.) You can go get Botox and then head right back to the office.

A class of celebrity plastic surgeons has emerged on Instagram, posting time-lapse videos of injection procedures and before-and-after photos, which receive hundreds of thousands of views and likes. According to the American Society of Plastic Surgeons, Americans received more than seven million neurotoxin injections in 2018, and more than two and a half million filler injections. That year, Americans spent \$16.5 billion on cosmetic surgery; ninety-two per cent of these procedures were performed on women. Thanks to injectables, cosmetic procedures are no longer just for people who want huge changes, or who are deep in battle with the aging process—they’re for millennials, or even, in rarefied cases, members of Gen Z. Kylie Jenner, who was born in 1997, spoke on her reality-TV show “Life of Kylie” about wanting to get lip fillers after a boy commented on her small lips when she was fifteen.

Ideals of female beauty that can only be met through painful processes of physical manipulation have always been with us, from tiny feet in imperial China to wasp waists in nineteenth-century Europe. But contemporary systems of continual visual self-broadcasting—reality TV, social media—have created new disciplines of continual visual self-improvement. Social media has supercharged the propensity to regard one’s personal identity as a potential source of profit—and, especially for young

women, to regard one's body this way, too. In October, Instagram announced that it would be removing "all effects associated with plastic surgery" from its filter arsenal, but this appears to mean all effects *explicitly* associated with plastic surgery, such as the ones called "Plastica" and "Fix Me." Filters that give you Instagram Face will remain. For those born with assets—natural assets, capital assets, or both—it can seem sensible, even automatic, to think of your body the way that a McKinsey consultant would think about a corporation: identify underperforming sectors and remake them, discard whatever doesn't increase profits and reorient the business toward whatever does.

Smith first started noticing the encroachment of Instagram Face about five years ago, "when the lip fillers started," he said. "I'd do someone's makeup and notice that there were no wrinkles in the lips at all. Every lipstick would go on so smooth." It has made his job easier, he noted, archly. "My job used to be to make people look like that, but now people come to me already looking like that, because they're surgically enhanced. It's great. We used to have to contour you to give you those cheeks, but now you just went out and got them."

There was something strange, I said, about the racial aspect of Instagram Face—it was as if the algorithmic tendency to flatten everything into a composite of greatest hits had resulted in a beauty ideal that favored white women capable of manufacturing a look of rootless exoticism. "Absolutely," Smith said. "We're talking an overly tan skin tone, a South Asian influence with the brows and eye shape, an African-American influence with the lips, a Caucasian influence with the nose, a cheek structure that is predominantly Native American and Middle Eastern." Did Smith think that Instagram Face was actually making people look better? He did. "People are absolutely getting prettier," he said. "The world is so visual right now, and it's only getting more visual, and people want to upgrade the way they relate to it."

This was an optimistic way of looking at the situation. I told Smith that I couldn't shake the feeling that technology is rewriting our bodies to correspond to its own interests—rearranging our faces according to whatever increases engagement and likes. "Don't you think it's scary to imagine people doing this forever?" I asked.

"Well, yeah, it's *obviously* terrifying," he said.

B palm trees and department stores of Wilshire and the palm trees and boutique eateries of Santa Monica, there's a doctor, or several, on every block. On a Wednesday afternoon, I parked my rental car in a tiny underground lot, emerged next to a Sprinkles Cupcakes and a bougie psychic's office, and walked to a consultation appointment I had made with one of the best-known celebrity plastic surgeons, whose before-and-after Instagram videos frequently attract half a million views.

I'd booked the consultation because I was curious about the actual experience of a would-be millennial patient—a fact I had to keep mentioning to my boyfriend, who seemed moderately worried that I would come back looking like a human cat. A few weeks before, I had downloaded Snapchat for the first time and tried out the filters, which were in fact very flattering: they gave me radiant skin, doe lashes, a face shaped like a heart. It wasn't lost on me that when I put on a lot of makeup I am essentially trying to create a version of this face. And it wasn't hard for me to understand why millennial women who were born within spitting distance of Instagram Face would want to keep drawing closer to it. In a world where women are rewarded for youth and beauty in a way that they are rewarded for nothing else—and where a strain of mainstream feminism teaches women that self-objectification is progressive, because it's profitable—cosmetic work might seem like one of the few guaranteed high-yield projects that a woman could undertake.

The plastic surgeon's office was gorgeous and peaceful, a silvery oasis. A receptionist, humming along to "I Want to Know What Love Is," handed me intake forms, which asked about stress factors and mental health, among other things. I signed an arbitration agreement. A medical assistant took photos of my face from five different angles. A medical consultant with lush hair and a deeply warm, caring aura came into the room. Careful not to lie, and lightly alarmed by the fact that I didn't need to, I told her that I'd never gotten fillers or Botox but that I was interested in looking better, and that I wanted to know what experts would advise. She was complimentary, and told me that I shouldn't get too much done. After a while, she suggested that maybe I would want to pay attention to my chin as I aged, and maybe my cheeks, too—maybe I'd want to lift them a little bit.

Then the celebrity doctor came in, giving off the intensity of a surgeon and the focus of a glassblower. I said to him, too, that I was just interested in looking better, and wanted to know what an expert would recommend. I showed him one of my filtered Snapchat photos. He glanced at it, nodded, and said, "Let me show you what we could do." He took a photo of my face on his phone

and projected it onto a TV screen on the wall. “I like to use FaceTune,” he said, tapping and dragging.

Within a few seconds, my face was shaped to match the Snapchat photo. He took another picture of me, in profile, and FaceTuned the chin again. I had a heart-shaped face, and visible cheekbones. All of this was achievable, he said, with chin filler, cheek filler, and perhaps an ultrasound procedure that would dissolve the fat in the lower half of my cheeks—or we could use Botox to paralyze and shrink my masseter muscles.

I asked the doctor what he told people who came to see him wanting to look like his best-known patients. “People come in with pictures of my most famous clients all the time,” he said. “I say, ‘I can’t turn you into them. I can’t, if you’re Asian, give you a Caucasian face, or I could, but it wouldn’t be right—it wouldn’t look right.’ But if they show me a specific feature they want then I can work with that. I can say, ‘If you want a sharp jaw like that, we can do that.’ But, also, these things are not always right for all people. For you, if you came in asking for a sharp jaw, I would say no—it would make you look masculine.”

“Does it seem like more people my age are coming in for this sort of work?” I asked.

“I think that ten years ago it was seen as anti-cerebral to do this,” he said. “But now it’s empowering to do something that gives you an edge. Which is why young people are coming in. They come in to enhance something, rather than coming in to fix something.”

“And it’s subtle,” I said.

“Even with my most famous clients, it’s very subtle,” the doctor said. “If you look at photos taken five years apart, you can tell the difference. But, day to day, month to month, you can’t.”

I felt that I was being listened to very carefully. I thanked him, sincerely, and then a medical assistant came in to show me the recommendations and prices: injectables in my cheeks (\$5,500 to \$6,900), injectables in my chin (same price), an ultrasound “lipofreeze” to fix the asymmetry in my jawline (\$8,900 to \$18,900), or Botox in the TMJ region (\$2,500). I walked out of the clinic into the Beverly Hills sunshine, laughing a little, imagining what it’d be like to have a spare thirty thousand dollars on

hand. I texted photos of my FaceTuned jaw to my friends and then touched my actual jaw, a suddenly optional assemblage of flesh and bone.

The plastic surgeon Jason Diamond was a recurring star of the reality show “Dr. 90210” and has a number of famous clients, including the twenty-nine-year-old “Vanderpump Rules” star Lala Kent, who has posted photos taken in Diamond’s office on Instagram, and who told *People*, “I’ve had every part of my face injected.” Another client is Kim Kardashian West, whom Colby Smith described to me as “patient zero” for Instagram Face. (“Ultimately, the goal is always to look like Kim,” he said.) Kardashian West, who has inspired countless cosmetically altered doppelgängers, insists that she hasn’t had major plastic surgery; according to her, it’s all just Botox, fillers, and makeup. But she also hasn’t tried to hide how her appearance has changed. In 2015, she published a coffee-table book of selfies, called “Selfish,” which begins when she is beautiful the way a human is beautiful and ends when she’s beautiful in the manner of a computer animation.

I scheduled an interview with Diamond, whose practice occupies the penthouse of a building in Beverly Hills. On the desk in his office was a thank-you note from Chrissy Teigen. (It sat atop two of her cookbooks.) As with the doctor I’d seen the day before, Diamond, who has pool-blue eyes and wore black scrubs and square-framed glasses, looked nothing like the tabloid caricature of a plastic surgeon. He was youthful in a way that was only slightly surreal.

Diamond had trained with an old guard of top L.A. plastic surgeons, he told me—people who thought it was taboo to advertise. When, in 2004, he had the opportunity to appear on “Dr. 90210,” he decided to do it, against the advice of his wife and his nurses, because, he said, “I knew that I would be able to show results that the world had never seen.” In 2016, a famous client persuaded him to set up an Instagram account. He now has just under a quarter million followers. The employees at his practice who run the account like that Instagram allows patients to see him as a father of two and as a friend, not only as a doctor.

Diamond had long had a Web site, but in the past his celebrity patients didn’t volunteer to offer testimonials there. “And, of course, we never asked,” he said. “But now—it’s amazing. Maybe thirty per cent of the celebrities I take care of will just ask and offer to shout us out on social media. All of a

sudden, it's popular knowledge that all these people are coming here. For some reason, Instagram made it more acceptable." Cosmetic work had come to seem more like fitness, he suggested. "I think it's become much more mainstream to think about taking care of your face and your body as part of your general well-being. It's kind of understood now: it's O.K. to try to look your best."

There was a sort of cleansing, crystalline honesty to this high-end intersection of superficiality and pragmatism, I was slowly realizing. I hadn't needed to bother posing as a patient—these doctors spent all day making sure that people no longer felt they had anything to hide.

I asked Diamond if he had thoughts about Instagram Face. "You know, there's this look—this Bella Hadid, Kim Kardashian, Kylie Jenner thing that seems to be spreading," I said. Diamond said that he practiced all over the world, and that there were different regional preferences, and that no one template worked for every face. "But there are constants," he said. "Symmetry, proportion, harmony. We are always trying to create balance in the face. And when you look at Kim, Megan Fox, Lucy Liu, Halle Berry, you'll find elements in common: the high contoured cheekbones, the strong projected chin, the flat platform underneath the chin that makes a ninety-degree angle."

"What do you make of the fact that it's much more possible now for people to look at these celebrity faces and think, somewhat correctly, that they could look like that, too?" I asked.

"We could spend two whole days discussing that question," Diamond said. "I'd say that thirty per cent of people come in bringing a photo of Kim, or someone like Kim—there's a handful of people, but she's at the very top of the list, and understandably so. It's one of the biggest challenges I have, educating the person about whether it's reasonable to try to move along that path toward Kim's face, or toward whoever. Twenty years of practice, thousands and thousands of procedures, go into each individual answer—when I can do it, when I can't do it, and when we can do something but shouldn't, for any number of reasons." I told Diamond that I was afraid that if I ever tried injectables, I'd never stop. "It is true that the vast majority of our patients absolutely love their results, and they come back," he said.

We talked about the word "addiction." I said that I dyed my hair and wore makeup most days, and that I knew I would continue to dye my hair and spend money on makeup, and that I didn't consider this an addiction but a choice. (I thought about a line from the book *Perfect Me*, by the

philosopher Heather Widdows: “Choice cannot make an unjust or exploitative practice or act somehow, magically, just or non-exploitative.”) I asked Diamond if his patients felt more like themselves after getting work done.

“I can answer that in part because I do these things, too,” he said, gesturing to his face. “You know when you get a really good haircut, and you feel like the best version of yourself? This is that feeling, but exponential.”

On the way to Diamond’s office, I had passed a café that looked familiar: pale marble-topped tables, blond-wood floors, a row of Prussian-green snake plants, pendant lamps, geometrically patterned tiles. The writer Kyle Chayka has coined the term “AirSpace” for this style of blandly appealing interior design, marked by an “anesthetized aesthetic” and influenced by the “connective emotional grid of social media platforms”—these virtual spaces where hundreds of millions of people learn to “see and feel and want the same things.” WeWork, the collapsing co-working giant—which, like Instagram, was founded in 2010—once convinced investors of a forty-seven-billion-dollar vision in which people would follow their idiosyncratic dreams while enmeshed in a global network of near-indistinguishable office spaces featuring reclaimed wood, neon signs, and ficus trees. Direct-to-consumer brands fill podcast ad breaks with promises of the one true electric toothbrush and meals that arrive in the mail, selling us on the relief of forgoing choice altogether. The general idea seems to be that humans are so busy pursuing complicated forms of self-actualization that we’d like much of our life to be assembled for us, as if from a kit.

I went to see another Beverly Hills plastic surgeon, one who had more than three hundred thousand Instagram followers. I told the doctor that I was a journalist, and that I was there for a consultation. He studied my face from a few angles, felt my jaw, and suggested exactly what the first doctor had recommended. The prices were lower this time—if I had wanted to put the whole thing on my credit card, I could have.

I took the elevator down to the street with three very pretty women who all appeared to be in their early twenties. As I drove back to my hotel, I felt sad and subdued and self-conscious. I had thought that I was researching this subject at a logical distance: that I could inhabit the point of view of an ideal millennial client, someone who wanted to enhance rather than fix herself, who was ambitious

and pragmatic. But I left with a very specific feeling, a kind of bottomless need that I associated with early adolescence, and which I had not experienced in a long time.

I had worn makeup at sixteen to my college interviews; I'd worn makeup at my gymnastic meets when I was ten. In the photos I have of myself at ballet recitals when I was six or seven, I'm wearing mascara and blush and lipstick, and I'm so happy. What did it mean, I wondered, that I have spent so much of my life attempting to perform well in circumstances where an unaltered female face is aberrant? How had I been changed by an era in which ordinary humans receive daily metrics that appear to quantify how our personalities and our physical selves are performing on the market? What was the logical end of this escalating back-and-forth between digital and physical improvement?

On Instagram, I checked up on the accounts of the plastic surgeons I had visited, watching comments roll in: "this is what I need! I need to come see you ASAP!", "want want want," "what is the youngest you could perform this procedure?" I looked at the Instagram account of a singer born in 1999, who had become famous as a teen-ager and had since given herself an entirely new face. I met up with a bunch of female friends for dinner in L.A. that night, two of whom had already adopted injectables as part of their cosmetic routine. They looked beautiful. The sun went down, and the hills of L.A. started to glitter. I had the sense that I was living in some inexorable future. For some days afterward, I noticed that I was avoiding looking too closely at my face.



Jia Tolentino is a staff writer at The New Yorker. She is the author of the essay collection "Trick Mirror."

More: [Instagram](#) [Plastic Surgery](#) [Social Media](#) [California](#)

ANNALS OF ARTIFICIAL INTELLIGENCE

THERE IS NO A.I.

*There are ways of controlling the new technology—but first we have to
stop mythologizing it.*

By Jaron Lanier

April 20, 2023




Illustration by Nicholas Konrad / The New Yorker

As a computer scientist, I don't like the term "A.I." In fact, I think it's misleading—maybe even a little dangerous. Everybody's already using the term, and it might seem a little late in the day to be arguing about it. But we're at the beginning of a new technological era—and the easiest way to mismanage a technology is to misunderstand it.

The term “artificial intelligence” has a long history—it was coined in the nineteen-fifties, in the early days of computers. More recently, computer scientists have grown up on movies like “The Terminator” and “The Matrix,” and on characters like Commander Data, from “Star Trek: The Next Generation.” These cultural touchstones have become an almost religious mythology in tech culture. It’s only natural that computer scientists long to create A.I. and realize a long-held dream.

What’s striking, though, is that many of the people who are pursuing the A.I. dream also worry that it might mean doomsday for mankind. It is widely stated, even by scientists at the very center of today’s efforts, that what A.I. researchers are doing could result in the annihilation of our species, or at least in great harm to humanity, and soon. In a recent poll, half of A.I. scientists agreed that there was at least a ten-per-cent chance that the human race would be destroyed by A.I. Even my colleague and friend Sam Altman, who runs OpenAI, has made similar comments. Step into any Silicon Valley coffee shop and you can hear the same debate unfold: one person says that the new code is just code and that people are in charge, but another argues that anyone with this opinion just doesn’t get how profound the new tech is. The arguments aren’t entirely rational: when I ask my most fearful scientist friends to spell out how an A.I. apocalypse might happen, they often seize up from the paralysis that overtakes someone trying to conceive of infinity. They say things like “Accelerating progress will fly past us and we will not be able to conceive of what is happening.”

I don’t agree with this way of talking. Many of my friends and colleagues are deeply impressed by their experiences with the latest big models, like GPT-4, and are practically holding vigils to await the appearance of a deeper intelligence. My position is not that they are wrong but that we can’t be sure; we retain the option of classifying the software in different ways.

The most pragmatic position is to think of A.I. as a tool, not a creature. My attitude doesn’t eliminate the possibility of peril: however we think about it, we

can still design and operate our new tech badly, in ways that can hurt us or even lead to our extinction. Mythologizing the technology only makes it more likely that we'll fail to operate it well—and this kind of thinking limits our imaginations, tying them to yesterday's dreams. We can work better under the assumption that there is no such thing as A.I. The sooner we understand this, the sooner we'll start managing our new technology intelligently.

If the new tech isn't true artificial intelligence, then what is it? In my view, the most accurate way to understand what we are building today is as an innovative form of social collaboration.

A program like OpenAI's GPT-4, which can write sentences to order, is something like a version of Wikipedia that includes much more data, mashed together using statistics. Programs that create images to order are something like a version of online image search, but with a system for combining the pictures. In both cases, it's people who have written the text and furnished the images. The new programs mash up work done by human minds. What's innovative is that the mashup process has become guided and constrained, so that the results are usable and often striking. This is a significant achievement and worth celebrating—but it can be thought of as illuminating previously hidden concordances between human creations, rather than as the invention of a new mind.

As far as I can tell, my view flatters the technology. After all, what is civilization but social collaboration? Seeing A.I. as a way of working together, rather than as a technology for creating independent, intelligent beings, may make it less mysterious—less like HAL 9000 or Commander Data. But that's good, because mystery only makes mismanagement more likely.

It's easy to attribute intelligence to the new systems; they have a flexibility and unpredictability that we don't usually associate with computer technology. But this flexibility arises from simple mathematics. A large language model like GPT-4

contains a cumulative record of how particular words coincide in the vast amounts of text that the program has processed. This gargantuan tabulation causes the system to intrinsically approximate many grammar patterns, along with aspects of what might be called authorial style. When you enter a query consisting of certain words in a certain order, your entry is correlated with what's in the model; the results can come out a little differently each time, because of the complexity of correlating billions of entries.

The non-repeating nature of this process can make it feel lively. And there's a sense in which it can make the new systems more human-centered. When you synthesize a new image with an A.I. tool, you may get a bunch of similar options and then have to choose from them; if you're a student who uses an L.L.M. to cheat on an essay assignment, you might read options generated by the model and select one. A little human choice is demanded by a technology that is non-repeating.

Many of the uses of A.I. that I like rest on advantages we gain when computers get less rigid. Digital stuff as we have known it has a brittle quality that forces people to conform to it, rather than assess it. We've all endured the agony of watching some poor soul at a doctor's office struggle to do the expected thing on a front-desk screen. The face contorts; humanity is undermined. The need to conform to digital designs has created an ambient expectation of human subservience. A positive spin on A.I. is that it might spell the end of this torture, if we use it well. We can now imagine a Web site that reformulates itself on the fly for someone who is color-blind, say, or a site that tailors itself to someone's particular cognitive abilities and styles. A humanist like me wants people to have more control, rather than be overly influenced or guided by technology. Flexibility may give us back some agency.

Still, despite these possible upsides, it's more than reasonable to worry that the new technology will push us around in ways we don't like or understand.

Recently, some friends of mine circulated a petition asking for a pause on the most ambitious A.I. development. The idea was that we'd work on policy during the pause. The petition was signed by some in our community but not others. I found the notion too hazy—what level of progress would mean that the pause could end? Every week, I receive new but always vague mission statements from organizations seeking to initiate processes to set A.I. policy.

These efforts are well intentioned, but they seem hopeless to me. For years, I worked on the E.U.'s privacy policies, and I came to realize that we don't know what privacy is. It's a term we use every day, and it can make sense in context, but we can't nail it down well enough to generalize. The closest we have come to a definition of privacy is probably "the right to be left alone," but that seems quaint in an age when we are constantly dependent on digital services. In the context of A.I., "the right to not be manipulated by computation" seems almost correct, but doesn't quite say everything we'd like it to.

A.I.-policy conversations are dominated by terms like "alignment" (is what an A.I. "wants" aligned with what humans want?), "safety" (can we foresee guardrails that will foil a bad A.I.?), and "fairness" (can we forestall all the ways a program might treat certain people with disfavor?). The community has certainly accomplished much good by pursuing these ideas, but that hasn't quelled our fears. We end up motivating people to try to circumvent the vague protections we set up. Even though the protections do help, the whole thing becomes a game—like trying to outwit a sneaky genie. The result is that the A.I.-research community communicates the warning that their creations might still kill all of humanity soon, while proposing ever more urgent, but turgid, deliberative processes.

Recently, I tried an informal experiment, calling colleagues and asking them if there's anything specific on which we can all seem to agree. I've found that there is a foundation of agreement. We all seem to agree that deepfakes—false but real-seeming images, videos, and so on—should be labelled as such by the programs

that create them. Communications coming from artificial people, and automated interactions that are designed to manipulate the thinking or actions of a human being, should be labelled as well. We also agree that these labels should come with actions that can be taken. People should be able to understand what they're seeing, and should have reasonable choices in return.

How can all this be done? There is also near-unanimity, I find, that the black-box nature of our current A.I. tools must end. The systems must be made more transparent. We need to get better at saying what is going on inside them and why. This won't be easy. The problem is that the large-model A.I. systems we are talking about aren't made of explicit ideas. There is no definite representation of what the system "wants," no label for when it is doing a particular thing, like manipulating a person. There is only a giant ocean of jello—a vast mathematical mixing. A writers'-rights group has proposed that real human authors be paid in full when tools like GPT are used in the scriptwriting process; after all, the system is drawing on scripts that real people have made. But when we use A.I. to produce film clips, and potentially whole movies, there won't necessarily be a screenwriting phase. A movie might be produced that appears to have a script, soundtrack, and so on, but it will have been calculated into existence as a whole. Similarly, no sketch precedes the generation of a painting from an illustration A.I. Attempting to open the black box by making a system spit out otherwise unnecessary items like scripts, sketches, or intentions will involve building another black box to interpret the first—an infinite regress.

At the same time, it's not true that the interior of a big model has to be a trackless wilderness. We may not know what an "idea" is from a formal, computational point of view, but there could be tracks made not of ideas but of people. At some point in the past, a real person created an illustration that was input as data into the model, and, in combination with contributions from other people, this was transformed into a fresh image. Big-model A.I. is made of people—and the way to open the black box is to reveal them.

This concept, which I've contributed to developing, is usually called "data dignity." It appeared, long before the rise of big-model "A.I.," as an alternative to the familiar arrangement in which people give their data for free in exchange for free services, such as internet searches or social networking. Data dignity is sometimes known as "data as labor" or "plurality research." The familiar arrangement has turned out to have a dark side: because of "network effects," a few platforms take over, eliminating smaller players, like local newspapers. Worse, since the immediate online experience is supposed to be free, the only remaining business is the hawking of influence. Users experience what seems to be a communitarian paradise, but they are targeted by stealthy and addictive algorithms that make people vain, irritable, and paranoid.

In a world with data dignity, digital stuff would typically be connected with the humans who want to be known for having made it. In some versions of the idea, people could get paid for what they create, even when it is filtered and recombined through big models, and tech hubs would earn fees for facilitating things that people want to do. Some people are horrified by the idea of capitalism online, but this would be a more honest capitalism. The familiar "free" arrangement has been a disaster.

One of the reasons the tech community worries that A.I. could be an existential threat is that it could be used to toy with people, just as the previous wave of digital technologies have been. Given the power and potential reach of these new systems, it's not unreasonable to fear extinction as a possible result. Since that danger is widely recognized, the arrival of big-model A.I. could be an occasion to reformat the tech industry for the better.

Implementing data dignity will require technical research and policy innovation. In that sense, the subject excites me as a scientist. Opening the black box will only make the models more interesting. And it might help us understand more about language, which is the human invention that truly impresses, and the one that we are still exploring after all these hundreds of thousands of years.

Could data dignity address the economic worries that are often expressed about A.I.? The main concern is that workers will be devalued or displaced. Publicly, techies will sometimes say that, in the coming years, people who work with A.I. will be more productive and will find new types of jobs in a more productive economy. (A worker might become a prompt engineer for A.I. programs, for instance—someone who collaborates with or controls an A.I.) And yet, in private, the same people will quite often say, “No, A.I. will overtake this idea of collaboration.” No more remuneration for today’s accountants, radiologists, truck drivers, writers, film directors, or musicians.

A data-dignity approach would trace the most unique and influential contributors when a big model provides a valuable output. For instance, if you ask a model for “an animated movie of my kids in an oil-painting world of talking cats on an adventure,” then certain key oil painters, cat portraitists, voice actors, and writers—or their estates—might be calculated to have been uniquely essential to the creation of the new masterpiece. They would be acknowledged and motivated. They might even get paid.

There is a fledgling data-dignity research community, and here is an example of a debate within it: How detailed an accounting should data dignity attempt? Not everyone agrees. The system wouldn’t necessarily account for the billions of people who have made ambient contributions to big models—those who have added to a model’s simulated competence with grammar, for example. At first, data dignity might attend only to the small number of special contributors who emerge in a given situation. Over time, though, more people might be included, as intermediate rights organizations—unions, guilds, professional groups, and so on—start to play a role. People in the data-dignity community sometimes call these anticipated groups mediators of individual data (MIDS) or data trusts. People need collective-bargaining power to have value in an online world—especially when they might get lost in a giant A.I. model. And when people share responsibility in a group, they self-police, reducing the need, or temptation, for governments and

companies to censor or control from above. Acknowledging the human essence of big models might lead to a blossoming of new positive social institutions.

Data dignity is not just for white-collar roles. Consider what might happen if A.I.-driven tree-trimming robots are introduced. Human tree trimmers might find themselves devalued or even out of work. But the robots could eventually allow for a new type of indirect landscaping artistry. Some former workers, or others, might create inventive approaches—holographic topiary, say, that looks different from different angles—that find their way into the tree-trimming models. With data dignity, the models might create new sources of income, distributed through collective organizations. Tree trimming would become more multifunctional and interesting over time; there would be a community motivated to remain valuable. Each new successful introduction of an A.I. or robotic application could involve the inauguration of a new kind of creative work. In ways large and small, this could help ease the transition to an economy into which models are integrated.

Many people in Silicon Valley see universal basic income as a solution to potential economic problems created by A.I. But U.B.I. amounts to putting everyone on the dole in order to preserve the idea of black-box artificial intelligence. This is a scary idea, I think, in part because bad actors will want to seize the centers of power in a universal welfare system, as in every communist experiment. I doubt that data dignity could ever grow enough to sustain all of society, but I doubt that any social or economic principle will ever be complete. Whenever possible, the goal should be to at least establish a new creative class instead of a new dependent class.

There are also non-altruistic reasons for A.I. companies to embrace data dignity. The models are only as good as their inputs. It's only through a system like data dignity that we can expand the models into new frontiers. Right now, it's much easier to get an L.L.M. to write an essay than it is to ask the program to generate an interactive virtual-reality world, because there are very few virtual worlds in

existence. Why not solve that problem by giving people who add more virtual worlds a chance for prestige and income?

Could data dignity help with any of the human-annihilation scenarios? A big model could make us incompetent, or confuse us so much that our society goes collectively off the rails; a powerful, malevolent person could use A.I. to do us all great harm; and some people also think that the model itself could “jailbreak,” taking control of our machines or weapons and using them against us.

We can find precedents for some of these scenarios not just in science fiction but in more ordinary market and technology failures. An example is the 2019 catastrophe related to Boeing’s 737 MAX jets. The planes included a flight-path-correction feature that in some cases fought the pilots, causing two mass-casualty crashes. The problem was not the technology in isolation but the way that it was integrated into the sales cycle, training sessions, user interface, and documentation. Pilots thought that they were doing the right thing by trying to counteract the system in certain circumstances, but they were doing exactly the wrong thing, and they had no way of knowing. Boeing failed to communicate clearly about how the technology worked, and the resulting confusion led to disaster.

Anything engineered—cars, bridges, buildings—can cause harm to people, and yet we have built a civilization on engineering. It’s by increasing and broadening human awareness, responsibility, and participation that we can make automation safe; conversely, if we treat our inventions as occult objects, we can hardly be good engineers. Seeing A.I. as a form of social collaboration is more actionable: it gives us access to the engine room, which is made of people.

Let’s consider the apocalyptic scenario in which A.I. drives our society off the rails. One way this could happen is through deepfakes. Suppose that an evil person, perhaps working in an opposing government on a war footing, decides to

stoke mass panic by sending all of us convincing videos of our loved ones being tortured or abducted from our homes. (The data necessary to create such videos are, in many cases, easy to obtain through social media or other channels.) Chaos would ensue, even if it soon became clear that the videos were faked. How could we prevent such a scenario? The answer is obvious: digital information must have context. Any collection of bits needs a history. When you lose context, you lose control.

Why don't bits come attached to the stories of their origins? There are many reasons. The original design of the Web didn't keep track of where bits came from, likely to make it easier for the network to grow quickly. (Computers and bandwidth were poor in the beginning.) Why didn't we start remembering where bits came from when it became more feasible to at least approximate digital provenance? It always felt to me that we wanted the Web to be more mysterious than it needed to be. Whatever the reason, the Web was made to remember everything while forgetting its context.

Today, most people take it for granted that the Web, and indeed the Internet it is built on, is, by its nature, anti-contextual and devoid of provenance. We assume that decontextualization is intrinsic to the very idea of a digital network. That was never so, however; the initial proposals for digital-network architecture, put forward by the monumental scientist Vannevar Bush in 1945 and the computer scientist Ted Nelson in 1960, preserved provenance. Now A.I. is revealing the true costs of ignoring this approach. Without provenance, we have no way of controlling our A.I.s, or of making them economically fair. And this risks pushing our society to the brink.

If a chatbot appears to be manipulative, mean, weird, or deceptive, what kind of answer do we want when we ask why? Revealing the indispensable antecedent examples from which the bot learned its behavior would provide an explanation: we'd learn that it drew on a particular work of fan fiction, say, or a soap opera. We could react to that output differently, and adjust the inputs of the model to

improve it. Why shouldn't that type of explanation always be available? There may be cases in which provenance shouldn't be revealed, so as to give priority to privacy—but provenance will usually be more beneficial to individuals and society than an exclusive commitment to privacy would be.

The technical challenges of data dignity are real and must inspire serious scientific ambition. The policy challenges would also be substantial—a sign, perhaps, that they are meaningful and concrete. But we need to change the way we think, and to embrace the hard work of renovation. By persisting with the ideas of the past—among them, a fascination with the possibility of an A.I. that lives independently of the people who contribute to it—we risk using our new technologies in ways that make the world worse. If society, economics, culture, technology, or any other spheres of activity are to serve people, that can only be because we decide that people enjoy a special status to be served.

This is my plea to all my colleagues. Think of people. People are the answer to the problems of bits. ♦

NEW YORKER FAVORITES

- In India, a clean-power plant the size of Manhattan could be a model for the world—or a cautionary tale.
- A heat shield for the most important ice on Earth.
- A case for climate optimism, and pragmatism, from John Podesta.
- The climate solutions we can't live without
- Saving the climate will depend on blue-collar workers. Can we train enough of them before time runs out?
- A major Black novelist made a remarkable début. Why did he disappear?

