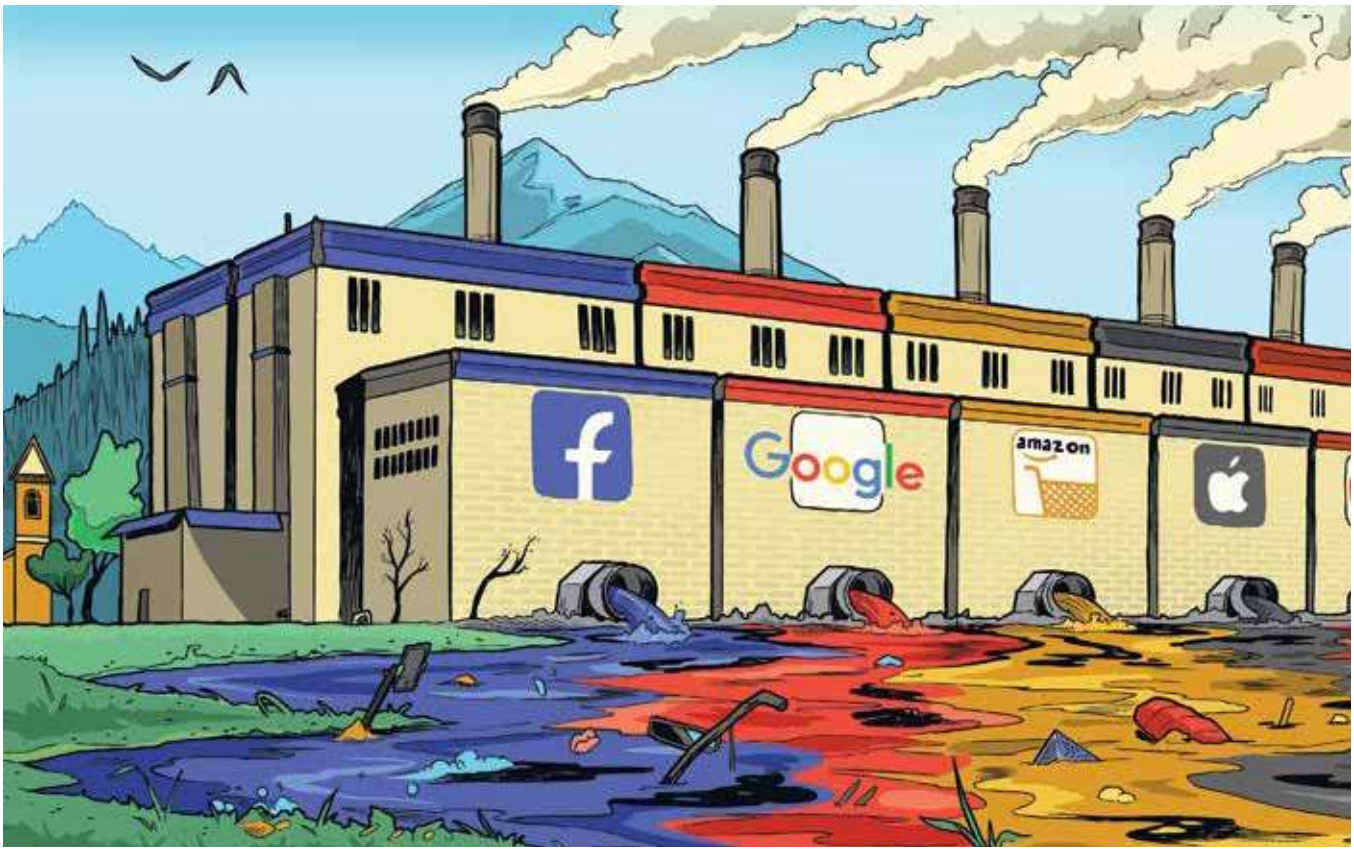

Nick Sirotich

The World Is Choking on Digital Pollution

Society figured out how to manage the waste produced by the Industrial Revolution. We must do the same thing with the internet today.

By Judy Estrin and Sam Gill

Tens of thousands of Londoners died of cholera from the 1830s to the 1860s. The causes were simple: mass quantities of human waste and industrial contaminants were pouring into the Thames, the central waterway of a city at the center of a rapidly industrializing world. The river gave off an odor so rank that Queen Victoria once had to cancel a leisurely boat ride. By the summer of 1858, Parliament couldn't hold hearings due to the overwhelming stench coming through the windows.

The problem was finally solved by a talented engineer and surveyor named Joseph Bazalgette, who designed and oversaw the construction of an industrial-scale, fully integrated sewer system. Once it was complete, London never suffered a major cholera outbreak again.

London's problem was not a new one for humanity. Natural and industrial waste is a fact of life. We start excreting in the womb and, despite all the inconveniences, keep at it for the rest of our lives. And, since at least the Promethean moment when we began to control fire, we've been contributing to human-generated emissions through advances intended to make our lives easier and more productive, often with little regard for the costs.

As industrialization led to increased urbanization, the by-products of combined human activity grew to such levels that their effects could not be ignored. The metaphorical heart of the world's industrial capital, the Thames was also the confluence of the effects of a changing society. "Near the bridges the feculence rolled up in clouds so dense that they were visible at the surface, even

in water of this kind," noted Michael Faraday, a British scientist now famous for his contributions to electromagnetism.

Relief came from bringing together the threads needed to tackle this type of problem—studying the phenomenon, assigning responsibility, and committing to solutions big enough to match the scope of what was being faced. It started with the recognition that direct and indirect human waste was itself an industrial-scale problem. By the 1870s, governmental authorities were starting to give a more specific meaning to an older word: they started calling the various types of waste "pollution."

A problem without a name cannot command attention, understanding, or resources—three essential ingredients of change. Recognizing that at some threshold industrial waste ceases to be an individual problem and becomes a social problem—a problem we can name—has been crucial to our ability to manage it. From the Clean Air Act to the Paris Accords, we have debated the environmental costs of progress with participants from all corners of society: the companies that produce energy or industrial products; the scientists who study our environment and our behaviors; the officials we elect to represent us; and groups of concerned citizens who want to take a stand. The outcome of this debate is not predetermined. Sometimes, we take steps to restrain industrial externalities. Other times, we unleash them in the name of some other good.

By the 1870s, governmental authorities were giving a more specific meaning to an old word: they called industrial waste "pollution." Now, we are confronting new and alarming by-products of progress, and the stakes may be just as high.

Now, we are confronting new and alarming by-products of progress, and the stakes for our planet may be just as high as they were during the Industrial Revolution. If the steam engine and blast furnace heralded our movement into the industrial age, computers and smartphones now signal our entry into the next age, one defined not by physical production but by the ease of services provided through the commercial internet. In this new age, names like Zuckerberg, Bezos, Brin, and Page are our new Carnegies, Rockefellers, and Fords.

As always, progress has not been without a price. Like the factories of 200 years ago, digital advances have given rise to a pollution that is reducing the quality of our lives and the strength of our democracy. We manage what we choose to measure. It is time to name and measure not only the progress the information revolution has brought, but also the harm that has come with it. Until we do, we will never know which costs are worth bearing.

We seem to be caught in an almost daily reckoning with the role of the internet in our society. This past March, Facebook lost $134 billion in market value over a matter of weeks after a scandal involving the misuse of user data by the political consulting firm Cambridge Analytica. In August, several social media companies banned InfoWars, the conspiracy-mongering platform of right-wing commentator Alex Jones. Many applauded this decision, while others cried of a left-wing conspiracy afoot in the C-suites of largely California-based technology companies.

Perhaps the most enduring political news story over the past two years has been whether Donald Trump and his campaign colluded with Russian efforts to influence the 2016 U.S. presidential election—efforts that almost exclusively targeted vulnerabilities in digital information services. Twitter, a website that started as a way to let friends know what you were up to, might now be used to help determine intent in a presidential obstruction of justice investigation.

And that's just in the realm of American politics. Facebook banned senior Myanmar military officials from the social network after a United Nations report accusing the regime of genocide against the Muslim Rohingya minority cited the platform's role in fanning the flames of violence. The spread of hoaxes and false kidnapping allegations on Facebook and messaging application WhatsApp (which is owned by Facebook) was linked to ethnic violence, including lynchings, in India and Sri Lanka.

Concerns about the potential addictiveness of on-demand, mobile technology have grown acute. A group of institutional investors pressured Apple to do something about the problem, pointing to studies showing technology's negative impact on students' ability to focus, as well as links between technology use and mental health issues. The Chinese government announced plans to control use of video games by children due to a rise in levels of nearsightedness. Former Facebook executive Chamath Palihapitiya described the mechanisms the company used to hold users' attention as "short-term, dopamine-driven feedback loops we've created [that] are destroying how society works," telling an audience at the Stanford Graduate School of Business that his own children "aren't allowed to use that shit."

The feculence has become so dense that it is visible—and this is only what has floated to the top.

For all the good the internet has produced, we are now grappling with effects of *digital pollution* that have become so potentially large that they implicate our collective well-being. We have moved beyond the point at which our anxieties about online services stem from individuals seeking to do harm—committing crimes, stashing child pornography, recruiting terrorists. We are now face-to-face with a system that is embedded in every structure of our lives and institutions, and that is itself shaping our society in ways that deeply impact our basic values.

We are right to be concerned. Increased anxiety and fear, polarization, fragmentation of a shared context, and loss of trust are some of the most apparent impacts of digital pollution. Potential degradation of intellectual and emotional capacities, such as critical thinking, personal authority, and emotional well-being, are harder to detect. We don't fully understand the cause and effect of digital toxins. The amplification of the most odious beliefs in social media posts, the dissemination of inaccurate information in an instant, the anonymization of our public discourse, and the vulnerabilities that enable foreign governments to interfere in our elections are just some of the many phenomena that have accumulated to the point that we now have real angst about the future of democratic society.

In one sense, the new technology giants largely shaping our online world aren't doing anything new. Amazon sells goods directly to consumers and uses consumer data to drive value and sales; Sears Roebuck delivered goods to homes, and Target was once vilified for using data on customer behavior to sell maternity products to women who had yet to announce their pregnancies. Google and Facebook grab your attention with information you want or need, and in exchange put advertisements in front of you; newspapers started the same practice in the nineteenth century and have continued to do it into the twenty-first—even if, thanks, in part, to Google and Facebook, it's not longer as lucrative.

But there are fundamental and far-reaching differences. The instantaneity and connectivity of the internet allow new digital pollution to flow in unprecedented ways. This can be understood through three ideas: scope, scale, and complexity.

The *scope* of our digital world is wider and deeper than we tend to recognize.

It is *wider* because it touches every aspect of human experience, reducing them all to a single small screen that anticipates what we want or "should" want. After the widespread adoption of social media and smartphones, the internet evolved from a tool that helped us do certain things to the primary surface for our very existence. Data flows into our smart TV, our smart fridge, and the location and voice assistants in our phones, cars, and gadgets, and comes back out in the form of services, reminders, and notifications that shape what we do and how we behave.

It is *deeper* because the influence of these digital services goes all the way down, penetrating our mind and body, our core chemical and biological selves. Evidence is mounting that the 150 times a day we check our phones could be profoundly influencing our behaviors and trading on our psychological reward systems in ways more pervasive than any past medium. James Williams, a ten-year Google employee who worked on advertising and then left to pursue a career in academia, has been sounding the alarm for years. "When, exactly, does a 'nudge' become a 'push'?" he asked five years ago. "When we call these types of technology 'persuasive,' we're implying that they shouldn't cross the line into being coercive or manipulative. But it's hard to say where that line is."

Madison Avenue had polls and focus groups. But they could not have imagined what artificial intelligence systems now do. Predictive systems curate and filter. They interpret our innermost selves and micro-target content we will like in order to advance the agendas of marketers, politicians, and bad actors. And with every click (or just time spent looking at something), these tools get immediate feedback and more insights, including the Holy Grail in advertising: determining cause and effect between ads and human behavior. The ability to gather data, target, test, and endlessly loop is every marketer's dream—brought to life in Silicon Valley office parks. And the more we depend on technology, the more it changes us.

The scope of the internet's influence on us comes with a problem of *scale*. The instantaneity with which the internet connects most of the globe, combined with the kind of open and participatory structure that the "founders" of the internet sought and valorized, has created a flow of information and interaction that we may not be able to manage or control in a safe way.

After the widespread adoption of social media and smartphones, the internet evolved from a tool that helped us do certain things to the primary surface for our very existence. And the more we depend on technology, the more it changes us.

A key driver of this scale is how easy and cheap it is to create and upload content, or to market services or ideas. Internet-enabled services strive to drain all friction out of every transaction. Anyone can now rent their apartment, sell their junk, post an article or idea—or just amplify a sentiment by hitting "like." The lowering of barriers has, in turn, incentivized how we behave on the internet—in both good and bad ways. The low cost of production has allowed more free expression than ever before, sparked new means of providing valued services, and made it easier to forge virtuous connections across the globe. It also makes it easier to troll or pass along false information to thousands of others. It has made us vulnerable to manipulation by people or governments with malevolent intent.

The sheer volume of connections and content is overwhelming. Facebook has more than two billion active users each month. Google executes three and a half billion searches per day. YouTube streams over one billion hours of video per day. These numbers challenge basic human comprehension. As

one Facebook official said in prepared testimony to Congress this year, "People share billions of pictures, stories, and videos on Facebook daily. Being at the forefront of such a high volume of content means that we are also at the forefront of new and challenging legal and policy questions."

Translation: *We're not sure what to do either*. And, instead of confronting the ethical questions at stake, the corporate response is often to define incremental policies based on what technology can do. Rather than considering actual human needs, people and society evolve toward what digital technology will support.

The third challenge is that the scope and scale of these effects relies on increasingly *complex* algorithmic and artificial intelligence systems, limiting our ability to exercise any human management. When Henry Ford's assembly line didn't work, a floor manager could investigate the problem and identify the source of human or mechanical error. Once these systems became automated, the machines could be subjected to testing and diagnostics and taken apart if something went wrong. After digitization, we still had a good sense of what computer code would produce and could analyze the code line by line to find errors or other vulnerabilities.

Commercial forces are taking basic questions out of our hands. It is treated as inevitable that there must be billons of posts, billions of pictures, billions of videos. The focus is on business: more users, more engagement, and greater activity.

Large-scale machine-learning systems cannot be audited in this way. They use information to learn how to do things. Like a human brain, they change as they learn. When they go wrong, artificial intelligence systems cannot be seen from a God's-eye view that tells us what happened. Nor can we predict exactly what they will do under unknown circumstances. Because they evolve based on the data they take in, they have the potential to behave in unexpected ways.

Taken together, these three kinds of change—the scope of intertwining digital and non-digital experience, the scale and frequency leading to unprecedented global reach, and the complexity of the machines—have resulted in impacts at least as profound as the transition from agricultural to industrial society, over a much shorter period of time. And the very elements that have made the internet an incredible force for good also come together to create new problems. The shift is so fundamental that we do not really understand the impacts with any clarity or consensus. What do we call hate speech when it is multiplied by tens of thousands of human and nonhuman users for concentrated effect? What do we call redlining when it is being employed implicitly by a machine assigning thousands of credit ratings per second in ways the machine's creator can't quite track? What do we call the deterioration of our intellectual or emotional capacities that results from checking our phones too often?

We need a common understanding, not just of the benefits of technology, but also of its costs—to our society and ourselves.

Human society now faces a critical choice: Will we treat the effects of digital technology and digital experience as something to be managed collectively? Right now, the answer being provided by those with the greatest concentration of power is no.

The major internet companies treat many of these decisions as theirs, even as CEOs insist that they make no meaningful decisions at all. Jack Dorsey warned against allowing Twitter to become a forum "constructed by our [Twitter employees'] personal views." Mark Zuckerberg, in reference to various conspiracy theories, including Holocaust denialism, stated, "I don't believe that our platform should take that down because I think there are things that different people get wrong. I don't think that they're intentionally getting it wrong."

These are just the explicit controversies, and the common refrain of "We are just a platform for our users" is a decision by default. There can be no illusions here: corporate executives are making critical societal choices. Every major internet company has some form of "community standards" about acceptable practices and content; these standards are expressions of their own values. The problem is that, given their pervasive role, these companies' values come to govern all of our lives without our input or consent.

Commercial forces are taking basic questions out of our hands. We go along through our acceptance of a kind of technological determinism: the technology simply marches forward toward less friction, greater ubiquity, more convenience. This is evident, for example, when leaders in tech talk about the volume of content. It is treated as inevitable that there must be billons of posts, billions of pictures, billions of videos. It is evident, too, when these same leaders talk to institutional investors in quarterly earnings calls. The focus is on business: more users, more engagement, and greater activity. Stagnant growth is punished in the stock price.

Commercial pressures have impacted how the companies providing services on the internet have evolved. Nicole Wong, a former lawyer for Google (and later a White House official) recently reflected during a podcast interview on how Google's search priorities changed over time. In the early days, she said, it was about getting people all the right information quickly. "And then in the mid-2000s, when social networks and behavioral ad-

vertising came into play, there was this change in the principles," she continued. After the rise of social media, Google became more focused on "personalization, engagement . . . what keeps you here, which today we now know very clearly: It's the most outrageous thing you can find."

The drive for profits and market dominance is instilled in artificial intelligence systems that aren't wired to ask why. But we aren't machines; we can ask why. We must confront how these technologies work, and evaluate the consequences and costs for us and other parts of our society. We can question whether the companies' "solutions"—like increased staffing and technology for content moderation—are good enough, or if they are the digital equivalent of "clean coal." As the services become less and less separable from the rest of our lives, their effects become ever more pressing social problems. Once London's industrial effluvia began making tens of thousands fall ill, it became a problem that society shared in common and in which all had a stake. How much digital pollution will we endure before we take action?

We tend to think of pollution as something that needs to be eradicated. It's not. By almost every measure, our ability to tolerate some pollution has improved society. Population, wealth, infant mortality, life span, and morbidity have all dramatically trended in the right direction since the industrial revolution. Pollution is a by-product of systems that are intended to produce a collective benefit. That is why the study of industrial pollution itself is not a judgment on what actions are overall good or bad. Rather, it is a mechanism for understanding effects that are large enough to influence us at a level that dictates we respond collectively.

We must now stake a collective claim in controlling digital pollution. What we face is not the good or bad decision of any one individual or even one company. It is not just about making economic decisions. It is about dispassionately analyzing the economic, cultural, and health impacts on society and then passionately debating the values that should guide our choices—as companies, as individual employees, as consumers, as citizens, and through our leaders and elected representatives.

Hate speech and trolling, the proliferation of misinformation, digital addiction—these are not the unstoppable consequences of technology. A society can decide at what level it will tolerate such problems in exchange for the benefits, and what it is willing to give up in corporate profits or convenience to prevent social harm.

We have a model for this urgent discussion. Industrial pollution is studied and understood through descriptive sciences that name and measure the harm. Atmospheric and environmental scientists research how industrial by-products change the air and water. Ecologists measure the impact of industrial processes on plant and animal species. Environmental economists create models that help us understand the trade-offs between a rule limiting vehicle emissions and economic growth.

We require a similar understanding of digital phenomena—their breadth, their impact, and the mechanisms that influence them. What are the various digital pollutants, and at what level are they dangerous? As with environmental sciences, we must take an interdisciplinary approach, drawing not just from engineering and design, law, economics, and political science but also from fields with a deep understanding of our humanity, including sociology, anthropology, psychology, and philosophy.

To be fair, digital pollution is more complicated than industrial pollution. Industrial pollution is the by-product of a value-producing process, not the product itself. On the internet, value and harm are often one and the same. It is the convenience of instantaneous communication that forces us to constantly check our phones out of worry that we might miss a message or notification. It is the way the internet allows more expression that amplifies hate speech, harassment, and misinformation than at any point in human history. And it is the helpful personalization of services that demands the constant collecting and digesting of personal information. The complex task of identifying where we might sacrifice some individual value to prevent collective harm will be crucial to curbing digital pollution. Science and data inform our decisions, but our collective priorities should ultimately determine what we do and how we do it.

Digital pollution is more complicated than industrial pollution. Industrial pollution is the by-product of a value-producing process, not the product itself. On the internet, value and harm are often one and the same.

The question we face in the digital age is not how to have it all, but how to maintain valuable activity at a societal price on which we can agree. Just as we have made laws about tolerable levels of waste and pollution, we can make rules, establish norms, and set expectations for technology.

Perhaps the online world will be less instantaneous, convenient, and entertaining. There could be fewer cheap services. We might begin to add friction to some transactions rather than relentlessly subtracting it. But these constraints would not destroy innovation. They would channel it, driving creativity in more socially desirable directions. Properly managing the waste of millions of Londoners took a lot more work than dumping it in the Thames. It was worth it. WM

Judy Estrin is an Internet pioneer, business executive, technology entrepreneur, the CEO of JLabs, and the author of *Closing the Innovation Gap.* **Sam Gill** is a vice president at the James S. and John L. Knight Foundation.