TECH

Social re-engineering, from Myanmar to Germany

Tech We're Using

BY THE NEW YORK TIMES

How do New York Times journalists use technology in their jobs and in their personal lives? Max Fisher, a reporter based in London who co-writes The Interpreter column and newsletter, discussed the tech he's using.

PHOTOGRAPHS BY BEN QUINTON FOR THE NEW YORK TIMES

Max Fisher, who has written about social media's effects, working with his Interpreter co-writer, Amanda Taub, in London, where he is based. He always packs his wireless, noise-canceling Sony H.ear headphones, above right, and two phones — one just for sensitive information. His loaded Kindle also provides a welcome distraction when he is abroad.

You travel all around the world for your Interpreter columns. What is your tech setup while you're on the road?

Mostly, I try to follow my colleague Sheera Frenkel's security advice. Two laptops and two phones, one of each just for sensitive stuff, which get fully wiped and reloaded after each trip. Everything goes in my Patagonia Headway, the greatest travel bag in human history.

Lots of countries block or monitor certain sites, so I use VPNs, which route all traffic through an anonymous server in another country. And I tether my laptop to my phone so I can avoid sketchy Wi-Fi.

I also need to keep myself sane. So I keep my Kindle fully loaded and always pack my noise-canceling, wireless headphones, the Wirecutter-recommended Sony H.ear. Even if I'm stuck in some edge-of-the-world hotel, living off of granola bars, I can always recharge by escaping for an hour into a Grateful Dead show or the second act of "Doctor Zhivago."

In which country did you find the way people use technology the most surprising and why?

I first went to Myanmar in early 2014, when the country was opening up, and there was no such thing as personal technology. Not even brick phones.

When I went back in late 2017, I could hardly believe it was the same country. Everybody had his or her nose in a smartphone, often logged in to Facebook. You'd meet with the same sources at the same roadside cafe, but now they'd drop a stack of iPhones on the table next to the tea.

It was like the purest possible experiment in what the same society looks like with or without modern consumer technology. Most people loved it, but it also helped drive genocidal violence against the Rohingya minority, empower military hard-liners and spin up riots.

People sometimes talk about this showing that Myanmar wasn't "ready" to come online so rapidly. But it looked to me like the same distorting effect of social media I'd seen in any other country. Maybe the change was just more obvious because it happened so rapidly and Myanmar was already pretty messed up.

You've lately been writing a lot about the effects of social media on the world. What have been some of your major takeaways?

We think of any danger as coming from misuse — scammers, hackers, state-sponsored misinformation — but we're starting to understand the risks that come from these platforms working exactly as designed. Facebook, YouTube and others use algorithms to identify and promote content that will keep us engaged, which turns out to amplify some of our worst impulses.

Even after reporting with Amanda Taub on algorithm-driven violence in Germany and Sri Lanka, I didn't quite appreciate this until I turned on Facebook push alerts this summer. Right away, virtually every gadget I owned started blowing up with multiple daily alerts urging me to check in on my ex, even if she hadn't posted anything. I'd stayed away from her page for months specifically to avoid training Facebook to show me her posts. Yet somehow the algorithm had correctly identified this as the thing likeliest to make me click, then followed me across continents to ensure that I did.

It made me think of the old "Terminator" movies, except instead of a killer robot sent to find Sarah Connor, it's a sophisticated set of programs ruthlessly pursuing our attention. And exploiting our most human frailties to do it.

Facebook's terrorizing me into mourning a breakup hardly matters. But, for a lot of users, unhealthy-but-irresistible content can come in more consequential forms. Like a viral rumor or a statement of hate we might otherwise know to avoid.

Where do you think this might all lead us?

I spend a lot of my time asking people this. What is the aggregate effect of routing an ever-growing share of human social relations through engagement-maximizing algorithms?

Maybe the effect is broadly negative. Maybe it's broadly positive. Probably it's mixed. But it is almost certainly profoundly disruptive in ways that we may spend the rest of our lives trying to understand.

Whether they set out to or not, these companies are conducting the largest social re-engineering experiment in human history, and no one has the slightest clue what the consequences are.

In the meantime, I've turned off Facebook push alerts and have reinstated a longstanding practice of avoiding any activity that would train an algorithm in what makes me click. I use sites like YouTube only anonymously and with my browser in incognito mode. (Separately, like my colleague Nellie Bowles, I set my screens to grayscale.)

It's not that I fear some devastating privacy breach or misuse of my data. Rather, these platforms are incredibly sophisticated at learning our habits and keeping us engaged in ways that are not necessarily healthy for us or our communities.

Outside work, is there a gadget or software or some other tech tool that you or your family loves using? Why?

I do a lot of cycling, so my iPhone is stuffed with various weather and transit apps. Google Maps has mostly replaced the GPS gadgets and old-fashioned bicycle maps. The only app of those worth recommending, London Air, tracks London's air quality.

To give social media some credit, some of my favorite serialized entertainment of any kind is the Twitter feed of Nicole Cliffe, a writer for various publications, which could exist and feel so personal only on a platform like Twitter. My sister and I regularly send each other tweets of hers, like a recent story about her mother's quest to reclaim stolen marijuana plants. They're funny and well written, as well as unfailingly kind and warm.

A robot mind answers your mail

ON TECHNOLOGY
FROM THE MAGAZINE

BY JOHN HERRMAN

In 1996, Microsoft unleashed Clippit, better known as Clippy, on users of Microsoft Office. The legendarily irritating mascot-helper spent the following years hovering around the edges of documents, blinking dumbly under his lascivious eyebrows and blurting out, "It looks like you're writing a letter," until the company sidelined him in 2001, officially recognized as a mistake.

Clippy's problems were manifold. He announced his presence, via a personified avatar, to tell us something that we already knew (or that should have been obvious in the first place) and then proudly offered us little in the way of actual help. He sat and watched us and learned nothing, and repeated himself. He said too much and did too little.

Nevertheless, over 20 years later, the spawn of Clippy are hiding everywhere, guessing what we're trying to do and offering to help. But Clippy's successors are doing their best to avoid his mistakes. Most of the time they are faceless, and if they speak, they do so in a disembodied but humanlike voice. They tend to wait to be asked for help, rather than telling us what they think they know unprompted. And when they do offer help, they tend to be more subtle, more accurate or both. They have perhaps more in common with Clippy's unassuming partners, like Spelling and Grammar Check or AutoCorrect, which spoke through red underlines or small actions carried out on reasonable assumptions (who would intentionally type "teh"?). These tools have followed us and our clumsy fingers to our new smartphones, where they have become both more assertive and more useful, correcting us and only occasionally requiring us to correct them back, and learning all the while.

What does the tech industry want to assist us with now? Email. If you use Gmail, you've probably interacted with either Smart Reply or Smart Compose, whether or not you know them by name. Google introduced Smart Reply in 2015, and Smart Compose began rolling out this year.

Both, in execution, are self-explanatory. Smart Reply suggests canned responses to inbound emails, based on the company's best guess at what most emailers might be about to type. The suggestions are short, peppy and often adequate, at least as a start. Sometimes their tone prompts unhappy realizations about what Gmail sees in us. The frequency with which they use exclamation marks emphasizes just how peculiar the language of professional email communication has become ("Sounds great!" "Very cool!" "Love it!"). Smart Compose, in contrast, offers word and phrase suggestions, based on similar judgments, as the user types in real time. You write "Take a look," and ghostly text might appear to its right: "and let me know what you think." Its assumptions are more personalized, and they feel that way because it is constantly, visibly, guessing what you're thinking.

Smart Compose and Smart Reply are, at their core, artificial-intelligence technologies: They are programmed to perform tasks, but also to adapt. To start, Smart Reply was trained on publicly available bodies of email text. (Among the most widely used for such projects is the cache of some 500,000 emails collected during the discovery phase of the Enron investigation.) "What makes machine learning different from regular programming is you look at corpuses of data to make guesses about things," says Paul Lambert, a product manager for Gmail. "You create a model."

Once that model was trained to deal with some of the more obvious idiosyncrasies of email communications — corporate disclaimers and phrases like "Sent from Outlook" — Google began training it on anonymized text from actual Gmail users. Phrases that appear frequently enough come under consideration for inclusion in Smart Reply. This, too, requires cleanup. Early testers reported seeing "I love you" as a suggested response to work emails.

PATRICK SEMANSKY/ASSOCIATED PRESS

Google is helping relieve the knowledge workers of the world of the drudgery of email by showing how inhuman it was.

Armed with this catalog of phrases — currently more than 20,000, according to the company — the model can then start incorporating more contextual clues: What was the subject of the email? Is the email asking a question? Is it expressing a happy sentiment, or is it offering condolences? Phrases are scored based on their utility — how much typing they save, basically — as well as the A.I.'s confidence in the prediction. Both features then take into account how people use them.

If, for example, it suggests a certain completion, and enough users take it, that one will be more likely to appear in the future. If a canned reply is never used, this is a signal that it should be purged; if it is frequently used, it will show up more often. This could, in theory, create feedback loops: common phrases becoming more common as they're offered back to users, winning a sort of election for the best way to say "O.K." with polite verbosity, and even training users, A.I.-like, to use them elsewhere. Such a dynamic would take root only where a behavior is already substantially automated — typed, at work, more as a learned performance rather than as an expression of will, or even an idea. Smart Compose is, in other words, good at isolating the ways we've already been programmed — by work, by social convention, by communication tools — and taking them off our hands.

Using these features is a bit like minding a machine that is trying to learn how to do what you do for a living. And even if it's the part of the job you wish you didn't have to do, it still prompts uncomfortable thoughts of replacement — or, if not replacement, then something close to it.

It is not remotely implausible that in the near future, a tremendous amount of communication could be conducted in tandem with an A.I. But constant sweeping changes in office communication — from speaking and writing to phones and printing to emailing and instant messaging — do not tell a tidy tale of increased efficiency or decreased workload, even as they represent progress. Already, an undefined but undeniable portion of workplace email amounts to human self-automation: an uncanny form of communication where clichés aren't shunned so much as recognized for their usefulness; where a tone of polite enthusiasm is taken to its exclamatory extreme to mash any ambivalence you may have about, say, "circling back later." One can visualize in the near future hundred-email chains between colleagues unfurling from a single human starting point, composed of nothing but routinized replies. Depending on what your current inbox looks like, this might not require much imagination at all. A study conducted in 2016 by researchers at Carleton University's Sprott School of Business in Canada tried to understand the role email had come to play in the modern office. They surveyed "highly educated baby boomer or Gen X" subjects who were mostly "managers or professionals" working in office jobs and found that they spend on average a full third of their workweeks "processing" email. Whatever their titles, they are — like many office workers — to a large extent professional emailers. Even if their roles are otherwise highly specialized, in this significant way they are not. They are their own assistants.

In 1930, John Maynard Keynes wrote that, thanks to new efficiencies, workers of the future could expect "three-hour shifts or a 15-hour week." He guessed that this would happen within a century. Automation and the abundance it produced has indeed led to countless economic changes, but it did not negate or replace the entire order.

Asked for evidence of the success of this newest tool, Google says that Smart Compose is already "saving people a billion characters of typing each week." This statistic supports one half of what Keynes might have predicted at the dawn of automated communication — the abundance and the glut — but is tellingly silent on the other half, the same half he couldn't quite see the first time. Self-automation can free us only to the extent that it actually belongs to us. We can be sure of only one thing that will result from automating email: It will create more of it.

John Herrman is a technology reporter for The Times.

Can a computer write a novel?

BERKELEY, CALIF.

An author's software take his text snippets and runs with the ideas

BY DAVID STREITFELD

Robin Sloan has a collaborator on his new novel: a computer.

The idea that a novelist is someone struggling alone in a room, equipped with nothing more than determination and inspiration, could soon be obsolete. Mr. Sloan is writing his book with the help of home-brewed software that finishes his sentences with the push of a tab key.

It's probably too early to add "novelist" to the long list of jobs that artificial intelligence will eliminate. But if you watch Mr. Sloan at work, it is quickly clear that programming is on the verge of redefining creativity.

Mr. Sloan, who won acclaim for his debut, "Mr. Penumbra's 24-Hour Bookstore," composes by writing snippets of text, which he sends to himself as messages and then works over into longer passages. His new novel, which is still untitled, is set in a near-future California where nature is resurgent. The other day, the writer made this note: "The bison are back. Herds 50 miles long."

In his cluttered man-cave of an office in an industrial park here, he is now expanding this slender notion. He writes: *The bison are gathered around the canyon. . . .* What comes next? He hits tab. The computer makes a noise like "pock," analyzes the last few sentences, and adds the phrase "by the bare sky."

Mr. Sloan likes it. "That's kind of fantastic," he said. "Would I have written 'bare sky' by myself? Maybe, maybe not."

He moves on: *The bison have been traveling for two years back and forth. . . .* Tab, pock. The computer suggests *between the main range of the city.*

"That wasn't what I was thinking at all, but it's interesting," the writer said. "The lovely language just pops out and I go, 'Yes.'"

His software is not labeled anything as grand as artificial intelligence. It's machine learning, facilitating and extending his own words, his own imagination. At one level, it merely helps him do what fledgling writers have always done — immerse themselves in the works of those they want to emulate. Hunter Thompson, for instance, strived to write in the style of F. Scott Fitzgerald, so he retyped "The Great Gatsby" several times as a shortcut to that objective.

Writers are readers, after all. "I have read some uncounted number of books and words over the years that all went into my brain and stewed together in unknown and unpredictable ways, and then certain things come out," Mr. Sloan said. "The output can't be anything but a function of the input."

But the input can be pushed in certain directions.

The program came up with a sentence that impressed him: "The slow-sweeping tug moved across the emerald harbor."

A quarter-century ago, an electronic surveillance consultant named Scott French used a supercharged Mac to imitate Jacqueline Susann's sex-drenched tales. His approach was different from Mr. Sloan's. Mr. French wrote thousands of computer-coded rules suggesting how certain character types derived from Ms. Susann's works might plausibly interact.

It took Mr. French and his Mac eight years to finish the tale — he estimated he could have done it by himself in one.

"Just This Once" was commercially published, a significant achievement in itself, although it did not join Ms. Susann's "Valley of the Dolls" on the best-seller list.

A tinkerer and experimenter, Mr. Sloan started down the road of computer-assisted creation driven by little more than "basic, nerdy curiosity." Many others have been experimenting with fiction that pushes in the direction of A.I.

PETER PRATO FOR THE NEW YORK TIMES

Robin Sloan, in his office in the Murray Street Media Lab in Berkeley, Calif. He is using a computer program he created to help write a novel set in the near future.

Botnik Studios used a predictive text program to generate four pages of rather wild Harry Potter fan fiction, which featured lines like this: "He saw Harry and immediately began to eat Hermione's family." On a more serious level, the Alibaba Group, the Chinese e-commerce company, said in January that its software for the first time outperformed humans on a global reading comprehension test. If the machines can read, then they can write.

Mr. Sloan wanted to see for himself. He acquired from the Internet Archive a database of texts: issues of Galaxy and If, two popular science fiction magazines in the 1950s and '60s.

After trial and error, the program came up with a sentence that impressed him: "The slow-sweeping tug moved across the emerald harbor."

"It was a line that made you say, 'Tell me more,'" Mr. Sloan said.

Those original magazines were too limiting, however, full of clichés and stereotypes. So Mr. Sloan augmented the pool with what he calls "The California Corpus," which includes the digital text of novels by John Steinbeck, Dashiell Hammett, Joan Didion, Philip K. Dick and others; Johnny Cash's poems; Silicon Valley oral histories; old Wired articles; the California Department of Fish and Wildlife's Fish Bulletin; and more. "It's growing and changing all the time," he said.

Unlike Mr. French a quarter-century ago, Mr. Sloan probably will not use his computer collaborator as a selling point for the finished book.

He's restricting the A.I. writing in the novel to an A.I. computer that is a significant character, which means the majority of the story will be his own inspiration. But while he has no urge to commercialize the software, he is intrigued by the possibilities. Megasellers like John Grisham and Stephen King could relatively easily market programs that used their many published works to assist fans in producing authorized imitations.

As for the more distant prospects, another San Francisco Bay Area science fiction writer long ago anticipated a time when novelists would turn over the composing to computerized "wordmills."

In Fritz Leiber's "The Silver Eggheads," published in 1961, the human "novelists" spend their time polishing the machines and their reputations. When they try to rebel and crush the wordmills, they find they have forgotten how to write.

Mr. Sloan has finished his paragraph:

"The bison were lined up fifty miles long, not in the cool sunlight, gathered around the canyon by the bare sky. They had been traveling for two years, back and forth between the main range of the city. They ring the outermost suburbs, grunting and muttering, and are briefly an annoyance, before returning to the beginning again, a loop that had been destroyed and was now reconstituted."

"I like it, but it's still primitive," the writer said. "What's coming next is going to make this look like crystal radio kits from a century ago."