




THE

AI

How To Fight
Big Tech's Hype
and Create the
Future We Want

EMILY M.
BENDER
&
ALEX
HANNA

CON

THE AI CON

HOW TO FIGHT BIG TECH'S HYPE
AND CREATE THE FUTURE WE WANT

EMILY M. BENDER
AND ALEX HANNA



HARPER

An Imprint of HarperCollinsPublishers

Dedication

For my kids, and the world I'd like them to have. —Emily

For my dad, who knew what work was. —Alex

Contents

Cover

Title Page

Dedication

Preface

Chapter 1: An Introduction to AI Hype

Chapter 2: It's Alive! The Hype of Thinking Machines

Chapter 3: Leisure for Me, Gig Work for Thee: AI Hype at Work

Chapter 4: If It Quacks Like a Doc: AI Hype and Social Services

Chapter 5: Artifice or Intelligence? AI Hype in Art, Journalism, and Science

Chapter 6: I'm Sorry, Dave, I'm Afraid I Can't Do That: AI Doomers, AI Boosters, and Why None of That Makes Sense

Chapter 7: Do You Believe in Hope After Hype?

Acknowledgments

Notes

Index

About the Authors

Copyright

About the Publisher

Preface

This book is the result of a joyful collaboration between a linguist (Emily) and a sociologist (Alex), rooted in a shared drive to deflate AI hype through analysis grounded in our academic fields and pointed humor.

We met online in 2020, participating in broader discussions about the impacts of technologies sold as “AI”.¹ In 2020 and 2021, we collaborated (with three other scholars, led by the inimitable Deb Raji²) on some academic papers critiquing the dismal evaluation and data handling practices of the field.

Since then, we’ve continued to use social media to take down ridiculous claims of tech boosters as well as bad journalism that fawns over them. This work can be trying, since speaking out against people’s favorite toys, tech leaders who appeal to a particular type of nerdy masculinity, and exploitative practices draws all kinds of negative pushback. Many people, especially in Silicon Valley and computer science departments, are willing to grant tech companies a good deal of grace for technologies that don’t live up to the hype. Beyond that, this work frequently involves contending with racism, sexism, and white supremacy and their associated violence. What kept us going was community, largely in the form of group chats, where we poked fun at all the hype, as well as processing all the nasty replies we received.

In one of those group chats, in April 2022, together with our colleagues Timnit Gebru and Margaret “Meg” Mitchell, we debated how to most effectively respond, especially when the hype takes the form of longer artifacts, such as scientific papers, podcast interviews, and long-form

journalism. Meg suggested something in the style of the TV show *Mystery Science Theater 3000* (MST3K), in which the characters take terrible sci-fi movies and make them enjoyable with running commentary.

A few months later, Emily came across a blog post that was too long for a tweet-thread takedown (a Medium post by Blaise Agüera y Arcas of Google titled “Can Machines Learn How to Behave?”). At an estimated sixty-minute read, every single paragraph was oozing AI hype. But the thought of writing a tweet thread or a blog post to counter each part of that seemed exhausting. So Emily asked in the group chat if anyone was up for giving it the *Mystery Science Theater* treatment.

Alex, a big fan of MST3K, jumped on board, and an accidental podcast was born: we livestreamed our takedown of the blog post on the platform Twitch, figuring it would take an hour or so. One hour wasn’t enough, so we scheduled a second, and then a third, and then just kept going. A few more streams in, we heeded many calls to turn our series into a podcast and brought Christie Taylor on board as a producer.

Our show, *Mystery AI Hype Theater 3000*, is billed as a space in which to seek catharsis in this age of AI hype: we find the worst of it and pop it with the sharpest needles we can find! Those needles are strong because they’re built from linguistic and sociological analysis, but sharp because they’re honed in humor.

One of our taglines on the podcast is “Along the way, we learn to always read the footnotes.” That’s because checking the sources for all of the hype-tastic claims often gives us a good vista on the house of cards (that is, thin research methods, shoddy argumentation, and questionable citation practice) supporting the flashy façade. In that same spirit, it is important to us to cite our sources: we care both about the provenance of the information we are sharing and about giving credit where credit is due. You’ll find those sources, along with further details and analysis we deemed too in-the-weeds for the main text, in the endnotes.

Our goal is to help the public at large as well as decision-makers at all levels become resistant to hype. Think of us as your guides to navigating a glitzy technology expo hall, full of salespeople trying to get you to buy a new product or fork over your data. We don’t need that energy, and neither do you.

We are writing now, in late 2024, from the inside of what feels like the height of the AI hype bubble. As we say on the podcast: each time we think

we've reached peak AI hype—the summit of bullshit mountain—we discover there's worse to come. We're using what we see in this bubble to document the contours of hype about AI, its causes and its short- and long-term effects. Our primary goal is to inhibit the next tech bubble. We hope that by pulling back the curtain, we'll help you to be able to spot the hype now and the next time around, while honing your own needles.

Chapter 1

An Introduction to AI Hype

In late 2023, inside the grand halls of the United States capital of Washington, DC, Senator Charles “Chuck” Schumer, Senate majority leader, led the eighth of a set of forums he had been convening around artificial intelligence, or AI. These “insight forums” were intended to provide the august body of the U.S. Senate with information on how to handle this “brand-new” technology of AI. At this particular meeting, a number of notables were in attendance: researcher Yoshua Bengio, who received one of computer science’s highest honors for his work on AI; Jared Kaplan, cofounder of the influential AI startup Anthropic; Aleksander Madry, OpenAI’s “Head of Preparedness”; and Stuart Russell, an influential professor of computer science. Also in the room were people from civil society (including civil rights and nonprofit research groups), policy institutes, and venture capital firms.

Schumer began the conversation¹ with an unusual prompt: What was everyone’s $p(\text{doom})$ and $p(\text{hope})$? Pronounced *pee-doom* (and *pee-hope*), this phrase references notation from statistics and is short for “probability of doom/hope”, referring to a popular trope that machines with minds of their own will, perhaps, kill us all, intentionally or unintentionally. Estimates from those in the room ranged from 0 up to 90 percent, according to reporting on the event. Schumer tweeted afterward²: “If managed properly, AI promises unimaginable potential. If left unchecked, AI poses both immediate and long-term risks.” These risks have been deemed

“existential” by those who have a $p(\text{doom})$ around the high end—risks that, if left unchecked, would threaten the whole of humankind.

Probability of doom, especially when written in the mathy-looking format $p(\text{doom})$, sounds like an important and sophisticated metric. Or at least the probability or $p()$ part does. But these stark terms are meant to grab headlines and grant an inflated sense of self-importance to those in the room. *Doom* reminds us of titanic, cartoonish fictional battles of good versus evil. And the cartoonish connotations are apt: just like such fictional battles, $p(\text{doom})$ estimates are based in fantasy rather than data or empirical fact. But that hasn’t prevented this imaginary metric from becoming a fixation of lawmakers, venture capitalists, and Silicon Valley’s managerial class. We imagine part of the appeal is that it allows people in power to imagine themselves as heroes out to save humanity, while actually turning away from the very real threats to actual people.

For example, the probability of techno-enabled doom brought about through automated state violence is very high for some citizens of Detroit. In January 2020, Robert Williams was arrested³ in front of his two young daughters, when Detroit police trusted the result of a database search of 49 million photos that matched his driver’s license photo to a freeze frame from a surveillance video of a theft, committed by someone else, two years earlier. The detectives didn’t acknowledge their error until Williams held the printed freeze frame next to his face. In February 2023, Porcha Woodruff was arrested⁴ and detained for eleven hours based also on the output of an automated facial recognition system. At the time, Woodruff was eight months pregnant and began to experience contractions while in police custody. The facial recognition system matched her image to footage of a (not visibly pregnant) person stealing a car. Both Williams and Woodruff are Black, and most known false positives⁵ for facial recognition tools have involved Black individuals. The probability is quite high that the lives of these people—and a number of other Black residents who have been mistakenly marked as criminal by facial recognition systems—have been irrevocably altered for the worse.⁶

A doomsday scenario has also arrived for teenagers, especially teenage girls, in the form of apps that purport to “undress” a person in an image. These image generation apps⁷ automate the task of making deepfake porn, allowing high school students to sexually harass and bully their classmates with a few clicks. The vast majority (99 percent) of deepfakes are of

women.⁸ The apps can produce such outputs because they are trained on indiscriminately collected troves of images⁹ from the internet, datasets that are so enormous no one could possibly verify each individual image in them. The datasets contain a lot of porn, meaning deepfake apps also create nonconsensual images of a sex worker's body.¹⁰ Distressingly, these datasets also include¹¹ child sexual abuse material.

In 2023, the Israel Defense Forces (IDF) and Prime Minister Benjamin Netanyahu's war-driven cabinet leveraged a system, again marketed as AI, in carrying out their assault on the Gaza Strip, in which tens of thousands of civilians were killed¹² in just the first three months. While the AI system was far from the only ingredient in what the head of the International Committee of the Red Cross and a spokesperson for a United Nations office both called "hell on Earth,"¹³ it served the purpose of rapidly scaling (and justifying) target selection: using a system called "The Gospel",¹⁴ the IDF dramatically expanded the scope of possible targets to include so-called "power targets", which includes high-rise residential blocks where a single Hamas member may live. In the words of one former officer, the system facilitates a "mass assassination factory."

But harms befalling real people are not what p(doom) refers to. Despite what many leaders in DC, New York, and Silicon Valley say, p(doom) is the wrong metric and the wrong framing. It serves to obfuscate what's really going on. Artificial intelligence, if we're being frank, is a *con*: a bill of goods you are being sold to line someone's pockets. A few major well-placed players are poised to accumulate significant wealth by extracting value from other people's creative work, personal data, or labor, and replacing quality services with artificial facsimiles. The language of p(doom) is a ruse to keep us focused on imaginary scenarios, filled with awe at modern robber barons' allegedly potentially world-ending technology, and too distracted to see the daily harms being done in its name.

We call this type of con "AI hype". Hype is not particularly new, and in fact we've been through AI hype cycles before. A characteristic of our current hype cycle is that the con men are taking a series of tropes from science fiction—of artificial minds hell-bent on turning us into paper clips or Terminators waging wars for their right to exist (and to look cool on motorcycles)—and injecting them into discussions at the highest echelons of business and government. This framing is useful to those creating the technology because it makes them appear powerful—if not godlike—in

their technical creation. But this belies what these technologies are doing to the rest of us: threatening stable careers and replacing them with gig work, slashing personnel in government, cheapening our social services, and degrading creativity.

To successfully navigate this technological moment, make wise choices as individual consumers and institutional decision-makers, and encourage lawmakers to enact smart policy, we argue that this framing needs to be discarded altogether. Not only does the rhetoric around $p(\text{doom})$ distract from actual harms, but the very terminology of “artificial intelligence” impedes clear understanding of the technologies in question, what they can and should be used for, and how to evaluate them. So, in our exploration of AI hype, we must first take a closer look at what people are talking about when they talk about “artificial intelligence”.

What Is “AI”?

To put it bluntly, “AI” is a marketing term. It doesn’t refer to a coherent set of technologies. Instead, the phrase “artificial intelligence” is deployed when the people building or selling a particular set of technologies will profit from getting others to believe that their technology is similar to humans, able to do things that, in fact, intrinsically require human judgment, perception, or creativity. But even in this case, there has to be a claim to similarity: calculators are far better than people at doing arithmetic, but they aren’t sold as “AI”. Sometimes the people selling these tools seem to believe their own marketing (we’ll meet several examples in later chapters), but what really matters is that they can sell it that way.

Throughout this book, we’re going to use the terms “artificial intelligence” or “AI” to refer to technologies sold as such. When speaking about a particular technology, we aim to be as precise as possible¹⁵. But when referring to these technologies in general, we will sometimes use the shorthand abbreviation of “AI”. We want to keep a critical distance from the term: every time we write “AI”, imagine we have a set of scare quotes around it. Or if you prefer, replace it with a ridiculous phrase. Some of our favorites include “mathy maths”, “a racist pile of linear algebra”, “stochastic parrots”¹⁶ (referring to large language models specifically), or

Systematic Approaches to Learning Algorithms and Machine Inferences (aka SALAMI¹⁷).

The set of technologies that get sold as AI is diverse, in both application and construction—in fact, we wouldn’t be surprised if some of the tech being sold this way is actually just a fancy wrapper around some spreadsheets. The term serves to obscure that diversity, however, so the conversation becomes clearer if one speaks in terms of “automation” rather than “AI” and looks at precisely *what is being automated*. In doing so, we find several types of automation.

Decision making. The first group involves using computers to automate consequential decisions. These are called automatic decision systems and they are often used, for example, in the process of setting bail, approving loans, screening résumés, or allocating social benefits. These uses are contentious, and rightfully so, because they have extreme ramifications for people who are subject to the system’s recommendations.

Classification. The second kind of automation involves classification of inputs of different types. For example, image classification can be used to help consumers organize their photos (where are all the photos of Grandma?), or can be used by governments for surveillance (matching a security footage frame to a database of driver’s license photos). The classification of web users for targeted advertising also fits into this group.

Recommendation. A third type selects information to present to someone, based on their own search or purchase history, or searches performed by someone else with a similar profile to them. These systems are called recommender systems. They’re behind the ordering of your feed in social media websites, Amazon product recommendations, or movie suggestions on Netflix.

Transcription/Translation. The fourth type is the automatic translation of information from one format to another: automatic transcription (sometimes called “automatic speech recognition” or “speech to text”), finding words and characters in images (like automatically reading license plates), machine translation of one language to another, or something like image style transfer (taking a selfie and making it look like an anime character).

Text and Image Generation. Then finally there’s a type that’s been very much in everyone’s mind recently: so-called generative AI or, more aptly, synthetic media machines. These are systems like ChatGPT, Gemini, or DALL-E that allow users to generate images or plausible-sounding text

based on textual prompts. A “prompt”, in generative AI terminology, is the words used to describe the desired output.

Lumping all of these different technologies under the label of “AI” creates the illusion of “intelligent” technology: if our photo software’s sharpening tool is imagined to be the same thing as the system that appears to cheerfully answer questions on any topic, then both are perceived as even more “intelligent” or even “magical” than each alone, and we’re more likely to accept automation in other domains, like deciding who gets social benefits or who is classified as a possible repeat offender. They are all supposedly driven by the same “intelligence”. Text synthesis machines have an outsized role here: language is so central to our understanding of each other that when we encounter language that doesn’t actually reflect the thoughts, ideas, or communicative intent of another person, it’s difficult not to imagine some humanlike mind behind it.

“AI” has always been a marketing term, but it hasn’t always been the marketing term of choice. In fact, up until fairly recently, the field was experiencing an “AI winter”, a time during which research funding was scarce, and the overall project of building computer systems that mimic human cognition was fairly marginalized within computer science. The companies building and selling such technologies as speech synthesis, automatic transcription, machine translation, image processing, and robotics did not label them as “AI”. That all changed in the 2010s, when one particular approach to pattern matching at scale—called “deep learning” — became practical for the first time. This wasn’t because of any magic or quantum leap in technology, but for the most part followed from innovation predicated on the falling costs of microchips and the abundance of digitized data on the web, easily accessible through a small set of platforms that centralized data sharing (Flickr, Tumblr, Google, and the like).

Even researchers working on these very approaches were surprised by the rapid switch from “AI winter” to seemingly unlimited venture capital funds. A research conference called Neural Information Processing Systems¹⁸ (NeurIPS, for short) grew from 1,354 attendees in 2010 to 13,000 attendees in 2019 and 22,000 attendees in 2020 (virtual due to COVID). In December 2012, when the conference was held outside of snowy Lake Tahoe (with a relatively sparse attendance of 1,676 people), a researcher named Geoff Hinton, along with his graduate students Alex Krizhevsky and Ilya Sutskever, held a secret auction¹⁹ for their company, DNNresearch. The

company had no product, nor any content on its website beyond its name. All it had was a paper that demonstrated their success in deep learning. Four companies—Microsoft, Google, the London-based AI startup DeepMind (later acquired by Google), and the Chinese search engine Baidu—made bids. The day went to Google, however, when Hinton stopped the auction at \$44 million. Hinton went on to join Google as a Distinguished Researcher for over a decade, and Sutskever later went on to become a cofounder and chief scientist at another startup, OpenAI. The deep learning era started with a bang, powered by immense amounts of money, capital, and, of course, hype.

What Is Hype?

Hype is the aggrandizement of some person, artifact, technology, or technique that you, the consumer, absolutely need to buy or invest in as early as possible, lest you miss out on entertainment or pleasure, monetary reward, return on investment, or market share. In the hip-hop world, the hype man is an accessory to the main act, the person who amps up the crowd for their employer. Software developer conferences might seem like the antithesis of hip-hop concerts, but then—Microsoft CEO Steve Ballmer played the hype man at a 1999 Microsoft event.²⁰ Voice hoarse, visibly sweaty, he pranced around the stage chanting “Developers, developers, developers!” and managed to get his audience of software engineers and managers to pick up that chant, buying into his hype about a mundane software framework.

Hype drives fashion trends, new musical artists, and car purchases. But more critically for this book, it drives investment in startups, technologies, and particular people.

Like other kinds of hype, AI hype plays on FOMO (the fear of missing out): it is the repeated message that a set of technologies—currently, a set of statistical methods developed within computer science and engineering—will change the world and you, the consumer or corporate manager, absolutely must use it, lest you be left in the dust. As a consumer, if you don’t get in on the hyped product, you’ll be seen as a regressive Luddite, lacking in modern skills, and/or about to have your job automated away. If you’re a corporate manager, you have to get on board, or competitors will

eat your lunch. If you're a computer programmer, you have to use new tools, otherwise you will be wasting time and won't meet product deadlines. If you're a teacher, you have to incorporate it into your curriculum, lest your students not be prepared for the AI-enhanced workplace. And if you're a student, you have to thoroughly understand AI to take on today's modern workplace, or else you'll get passed over for job opportunities.

The commercial function of tech hype is to boost sales of a product. In other words, marketing. Sam Altman, the CEO of OpenAI, is, like all the great tech barons of our era, an adman. But while all tech hype plays a commercial function, AI hype in particular plays a cultural function as well. It connects a commercial goal with a popular fantasy of sentient machines.

When selling rosy scenarios, AI hype promises us a life of ease: jobs deemed menial like data entry, writing ad copy, and making basic graphics will become a thing of the past. AI "companions" will take notes for you in online meetings or, even better, become your stand-in while you address more pressing matters. Surely technologies of today are just a few rounds of "progress" away from the onboard computer that Captain John-Luc Picard can confidently command to provide "Tea, Earl Grey, Hot" or the caring, competent "operating system" voiced by Scarlett Johansson in *Her*. Altman made this implicit fantasy explicit²¹ when he tweeted the single word "her" in advance of a product demo, a voice assistant that sounded suspiciously like Johansson—created without her consent.

But AI hype also depends on promulgating worst-case scenarios. Here, AI hype invokes visions of robots that disobey Isaac Asimov's First Law of Robotics: "A robot may not injure a human being or, through inaction, allow a human being to come to harm." These examples are rife throughout science fiction and myth: the robot HAL 9000, which disobeys the commands of humans in *2001: A Space Odyssey* in order to complete its mission; the machine race that takes over the face of the earth and uses humans as a power source in *The Matrix*; a rogue "Entity" in *Mission Impossible: Dead Reckoning*, which, after being developed by the U.S. government, turns on its masters. The tale is as old as Mary Shelley's *Frankenstein*, about the monster that turns on its creator, or even older, in the Judaic figure of the golem, which in some iterations of the story goes rogue after its human handlers forget to deactivate it.

Claims that we're but a step away from living in a science fiction world have little basis in reality. But just because the hype is ungrounded in the

real world doesn't mean the hype itself doesn't impact the world, culturally, economically, and environmentally. And while AI hype has reached a fever pitch in recent years, it has been with us for decades, back to the founding of the field. We can expect AI hype to accompany AI research as long as such research is pursued. A quick tour of the original AI hype will help us see through today's, and comparing AI hype—old and new—will help you identify it in the future, too.

A Brief History of AI (and AI Hype)

As long as there's been research on AI, there's been AI hype. In the most commonly told narrative about the research field's development, mathematician John McCarthy and computer scientist Marvin Minsky organized a summer-long workshop²² in 1956 at Dartmouth College in Hanover, New Hampshire, to discuss a set of methods around “thinking machines”. The term “artificial intelligence” is attributed to McCarthy, who was trying to find a name suitable for a workshop that concerned a diverse set of existing knowledge communities. He was also trying to find a way to exclude Norbert Wiener—the pioneer of a proximate field, cybernetics, a field that has to do with communication and control of machines—due to personal differences.

The way the origin story is told, Minsky and McCarthy convened the two-month working group at Dartmouth, consisting of a group of ten mathematicians, physicists, and engineers, which would make “a significant advance”²³ in this area of research. Just as it is today, the term “artificial intelligence” did not have much coherence. It did include something similar to today's “neural networks” (also called “neuron nets” or “nerve nets” in those early documents), but also covered topics that included “automatic computers” and human-computer language interfaces (what we would today consider to be “programming languages”).

Fundamentally, the forerunners of this new field were concerned with translating dynamics of power and control into machine-readable formulations. McCarthy, Minsky, Herbert Simon (political scientist, economist, computer scientist, and eventual Nobel laureate), and Frank Rosenblatt (one of the originators of the “neural network” metaphor) were concerned with developing tools that could be used for the guidance of

administrative—and ultimately—military systems. In an environment where the battle for American supremacy in the Cold War was being fought on all fronts—military, technological, engineering, and ideological²⁴—these men sought to gain favor and funding in the eyes of a defense apparatus trying to edge out the Soviets. They relied on huge claims with little to no empirical support, bad citation practices, and moving goalposts to justify their projects, which found purchase in Cold War America. These are the same set of practices that we see from today’s AI boosters, although they are now primarily chasing market valuations, in addition to government defense contracts.

The first move in the original AI hype playbook was foregrounding the fight with the Soviets. The second was to argue that computers were likely to match human capabilities by arguing that humans weren’t really all that complex. In 1956, Minsky claimed²⁵ in an influential paper that “[h]uman beings are instances of certain kinds of very complicated machines.” If that were indeed the case, we could use more controllable electronic circuits in place of people in military and industrial contexts.

In the late 1960s, Joseph Weizenbaum, a German émigré, professor at the Massachusetts Institute of Technology, and contemporary of Minsky, was alarmed by how quickly people attributed agency to automated systems. Weizenbaum developed a chatbot²⁶ called ELIZA, named for the working-class character in George Bernard Shaw’s *Pygmalion* who learns to mimic upper-class speech. ELIZA was designed to carry on a conversation in the style of a Rogerian psychotherapist;²⁷ that is, the program primarily repeated what its users said, reframing their thoughts into questions. Weizenbaum used this form for ELIZA, not because he thought it would be useful as a therapist, but rather because it was a convenient setup²⁸ for the chatbot: this kind of psychotherapy is one of the few conversational situations where it wouldn’t matter if the machine didn’t have access to other data about the world.

Despite its grave limitations, computer scientists used ELIZA to celebrate how thoroughly computers could replace human labor and heralded the entry into the artificial intelligence age. A shocked Weizenbaum²⁹ spent the rest of his life as a critic of AI, noting that humans were not meat machines, while Minsky went on to found MIT’s AI laboratory and rake in funding from the Pentagon unhindered.

The murky, unethical funding networks—through unfettered weapons manufacturing then, and with the addition of ballooning speculative venture capital investments now—around AI continue to this day. So does the drawing of false equivalences between the human brain and the calculating capabilities of machines. Claiming such false equivalences inspires awe, which, it turns out, can be used to reel in boatloads of money from investors whipped into a FOMO frenzy.

When we say boatloads, think megayachts: in January 2023, Microsoft announced³⁰ that it intended to invest \$10 billion in OpenAI. This is after Mustafa Suleyman (former CEO of DeepMind, made CEO of Microsoft AI in March 2024) and LinkedIn cofounder Reid Hoffman received³¹ a cool \$1.3 billion from Microsoft and chipmaker Nvidia in a funding round to their young startup, Inflection.AI. OpenAI alums cofounded Anthropic, a company solely focused on creating generative AI tools, and received \$580 million in an investment round led by crypto-scammer Sam Bankman-Fried.³² These startups, and a slew of others, have been chasing a gold mine of investment from venture capitalists and Big Tech companies, frequently without any clear path to robust monetization. By the second quarter of 2024, venture capital was dedicating \$27.1 billion, or nearly half of their quarterly investments, to AI and machine learning companies.³³

The incentives to ride the AI hype train are clear and widespread—dress something up as AI and investments flow. But both the technologies and the hype around them are causing harm in the here and now.

Of Hype and Harm

There *are* applications of machine learning that are well scoped, well tested, and involve appropriate training data such that they deserve their place among the tools we use on a regular basis. These include such everyday things as spell-checkers (no longer simple dictionary look-ups, but able to flag real words used incorrectly) and other more specialized technologies like image processing used by radiologists to determine which parts of a scan or X-ray require the most scrutiny. But in the cacophony of marketing and startup pitches, these sensible use cases are swamped by promises of machines that can effectively do magic, leading users to rely on them for

information, decision-making, or cost savings—often to their detriment or to the detriment of others.

As investor interest pushes AI hype to new heights, tech boosters have been promoting AI “solutions” in nearly every domain of human activity. We’re told that AI can shore up threadbare spots in social services, providing medical care and therapy to those who aren’t fortunate enough to have good access to health care, education to those who don’t live in a wealthy school district, and legal services for people who can’t afford a licensed attorney. We’re told that AI will provide individualized versions of all of these things, flexibly meeting user needs. We’re told that AI will “democratize” creative activity by allowing *anyone* to become an artist. We’re told that AI is on the verge of doing science for us, finally providing us with answers to urgent problems from medical breakthroughs (discovering a cure for cancer!) to the climate crisis (discovering a solution for global warming!). And self-driving cars are perpetually just around the corner (watch out: that means they’re about to run into you). But as you may have surmised from our snarky tone, these solutions are, by and large, AI hype. There are myriad cases in which AI solutions have been posed but fall short of their stated goals.

In 2017, a Palestinian man was arrested³⁴ by Israeli authorities over a Facebook post in which he posed next to a bulldozer with the caption (in Arabic) of “good morning.” Facebook’s machine translation software rendered that as “hurt them” in English and “attack them” in Hebrew—and the Israeli authorities just took that at face value, never checking with any Arabic speakers to see if it was correct. Machine translation has also become a weak stopgap in other critical situations, such as in handling asylum cases. Here, the problem to solve is one of communication, between people fleeing violence in their home countries and immigration officials. Machine translation systems, which can work well in cases like translating newspapers written in standard varieties of a handful of dominant languages, can fail drastically in translating asylum claims³⁵ written or spoken in minority languages or dialects.

In August 2020, thousands of British students³⁶, unable to take their A-level exams due to the COVID-19 pandemic, received grades calculated based on an algorithm that took as input, among other things, the grades that other students at their schools received in previous years. After massive public outcry, in which hundreds of students gathered outside the prime

minister's residence at 10 Downing Street in London, chanting "Fuck the algorithm!" the grades were retracted and replaced with grades based on teachers' assessment of student work. In May 2023, Jared Mumm³⁷, a professor at Texas A&M University, suspected his students of cheating by using ChatGPT to write their final essays—so he input the essays into ChatGPT and asked it whether it wrote them. After reading ChatGPT's affirmative output, he assigned the whole class incomplete grades, and some seniors were (temporarily) denied their diplomas.

On our roads, promises of self-driving cars have led to death and destruction. A Tesla employee died³⁸ after engaging the so-called "Full Self-Driving" mode in his Tesla Model 3, which ran the car off the road. (We know this partially because his passenger survived the crash.) A few months later, on Thanksgiving Day 2022, Tesla CEO Elon Musk announced the availability of Tesla's "Full Self-Driving" mode. Hours later, it was involved in an eight-car pileup on the San Francisco–Oakland Bay Bridge.

In 2023, lawyer Steven A. Schwartz³⁹, representing a plaintiff in a lawsuit against an airline, submitted a legal brief citing legal precedents that he found by querying ChatGPT. When the lawyers defending the airline said they couldn't find some of the cases cited and the judge asked Schwartz to submit them, he submitted excerpts, rather than the traditional full opinions. Ultimately, Schwartz had to own up to having trusted the output of ChatGPT to be accurate, and he and his cocounsel were sanctioned and fined by the court.

In November 2022, Meta released Galactica⁴⁰, a large language model trained on scientific text, and promoted it as able to "summarize academic papers, solve math problems, generate Wiki articles, write scientific code, annotate molecules and proteins, and more." The demo stayed up for all of three days, while the worldwide science community traded examples of how it output pure fabrications, including fake citations, and could easily be prompted into outputting toxic content relayed in academic-looking prose.

What all of these stories have in common is that someone oversold an automated system, people used it based on what they were told it could do, and then they or others got hurt. Not all stories of AI hype fit this mold, but for those that don't, it's largely the case that the harm is either diffuse or undocumented. Sometimes, people are able to resist AI hype, think through the possible harms, and choose a different path. And that brings us to our goal in writing this book: preventing the harm from AI hype. When people

can spot AI hype, they make better decisions about how and when to use automation, and they are in a better position to advocate for policies that constrain the use of automation by others.

Your Guides to Hype-Spotting

The two of us have spent the past few years attacking AI hype and puzzling around what's behind it. Emily is a linguist who's gained recognition for raising ethical issues in the development of language technology. Alex is a sociologist who formerly worked on Google's Ethical AI team and thinks a lot about how technology and society interact, especially how technology reinforces persistent inequalities along the lines of race, gender, and class.

In this book, we are trying to do the same thing Weizenbaum tried to do: educate people about how these systems work, dispel the notion that they are thinking machines with a semblance of human understanding, and provide a model of how to think about them instead. While we are working in a time when the general public has much more experience with computers than they did in the 1970s, we are also up against text extruding machines that are not only far more versatile than ELIZA, but backed by companies and investors with a deep financial interest in people perceiving their technology as a pervasive and all-powerful foregone conclusion. Where Weizenbaum offered a stern warning for computer scientists in the pockets of the military men, our project is faced with a hydra of AI startups funded by venture capital. Luckily, the public availability of these products opens them up to vastly more arenas for accountability by workers, consumers, and regulators.

AI hype today infests almost every corner of our culture: claims that we're dealing with sentient beings, and that one day, these beings will become superintelligent, are unavoidable, from popular entertainment to the halls of Congress. Breathless reporting uncritically parrots corporate statements that AI is going to free you from work, educate your kids, and provide medical care to all who need it. The hypers claim that AI will produce art but also might just kill us all. Should it decide to spare us, at the end of the day you'll be able to kick back in a fully automated paradise, once AI has solved the climate crisis and poverty.

This is, of course, all bullshit. AI isn't sentient, it's not going to make your job easier, and AI doctors aren't going to cure what ails you. But these claims can make your work worse and reduce your quality of life, unless we fight back against the increasing encroachment of these products into every area of public and private life. Hype doesn't occur by accident, but rather because it fulfills a function: scaring workers and promising to save decision-makers and business leaders lots of money. Part of our work in dismantling hype is tracking where it comes from and whose interest it serves.

In the chapters that follow, we take apart these claims piece-by-piece. In Chapter 2, we dig into how systems like large language models and text-to-image generation tools work. We also delve into why, precisely, we are so tempted to think that these are superior thinking machines. It has a lot to do with how we as humans process language, namely that we expect a thinking intelligence behind something that is using language. We also discuss how the search for "general intelligence" is not only a futile search, but one that is grounded in a deeply racist history.

In Chapter 3, we discuss the ways in which AI tools, like other kinds of automation before them, are being imposed on workers and furthermore used as a threat to their livelihoods. This pattern is as old as the Industrial Revolution, and so is worker resistance to it. We get into both the history and modern examples of the ways these tools are replacing jobs—and making an absolute mess of things. We also talk about the ways workers are fighting back: through solidarity, organizing, and other collective action.

Shoddy replacements for actual human labor and relationships are particularly problematic in social services. In Chapter 4, we dig into the ways in which corporations and governments are aiming to use automation to replace the splotchy arrangement of social services and education in the United States and elsewhere. Automated decision systems have already become stopgaps for government agencies that assign public housing and investigate child abuse and neglect. We're seeing startups chomp at the bit as they try to find a foothold in creating chatbots for health care and education.

We also find false promises of media synthesis machines in the creative fields, including art, science, and journalism, and take these up in Chapter 5. "AI artists" have proliferated online and the visual artists who they are ripping off are (rightly) none too happy about it. Meanwhile, a coterie of

scientists, abrogating scholarly duty by either misusing synthetic data or having language models write papers altogether, have made it harder to separate the wheat from the chaff in many scientific fields. And journalism, already suffering from dramatic job losses and fire sales of respected news organizations, is ripe for the infection of AI-generated content, intended to maximize eyeballs on ads with as little investment in actual journalism as possible.

In Chapter 6, we return to how we began this chapter, with a discussion of how AI Doomerism—all that talk of p(doom)—makes many of the same (false) assumptions and plays similar economic functions as AI Boosterism. Some Doomers are true believers, and are motivated by a set of ideologies including effective altruism and longtermism, which preach that we need “superintelligent” technology to colonize the stars. They’re “Doomers” because they worry that this technology might also turn on humanity. On the other side of the same coin are the Boosters, who see the same “potential” in AI to solve everything and want that future to arrive as quickly as possible, no matter the cost. Both sides hold a privileged, white- and Western-centric view of the world, and ignore the very real harms of these tools in the here and now.

To conclude, in Chapter 7, we discuss strategies to combat AI hype, such as robust regulation, data and privacy legislative proposals, and strong worker protections. But more importantly, we talk about what *you* can do about it.

By the end of this book, we hope we’ll have shown you how to resist the urge to be impressed, how to spot AI hype in the wild, and how to take back some ownership in our technological future. Let’s jump in.

Chapter 2

It's Alive! The Hype of Thinking Machines

If you listened to executives and researchers at big tech firms, you'd think that we were on the verge of a robot uprising. In February 2022, OpenAI's Chief Scientist Ilya Sutskever tweeted "it may be that today's large neural networks are slightly conscious."¹ In June 2022, the *Washington Post* reported that Google engineer Blake Lemoine was convinced that the company's language model LaMDA was sentient and needed legal representation. Lemoine was fired over this incident—not for his false claims (which Google did deny), but for leaking private corporate information.² In an August 2022 blog post, Google VP Fellow Blaise Agüera y Arcas responded to the Lemoine story, but rather than countering Lemoine's claims, he suggested that LaMDA does indeed "understand" concepts and that the debate over whether or not LaMDA has feelings is not resolvable or "scientifically meaningful."³ In April 2023, a team at Microsoft Research led by Sébastien Bubeck posted a non-peer-reviewed paper called "Sparks of Artificial General Intelligence: Early Experiments with GPT-4,"⁴ in which they claim to show that the language model GPT-4 "can solve novel and difficult tasks that span mathematics, coding, vision, medicine, law, [and] psychology" and thus shows the first "sparks of artificial general intelligence." The word "sparks" evokes an image of something about to catch fire and spread of its own accord. The phrase "artificial general intelligence" here is meant to differentiate from ordinary

technologies called “AI”, and is particularly common in modern discourse around thinking, sentient, or conscious machines.

These claims are not new. Over sixty years ago, researchers, business executives, and government officials were making similar bombastic claims about the nature of computer intelligence and the risk of superhuman intelligence supplanting humans at work, at home, and, perhaps most alarmingly, on the battlefield. Marvin Minsky, one of the organizers of the 1956 Dartmouth conference and founder of MIT’s AI Lab, writing to two colleagues in the 1950s, stated that there is “no aspect of learning or other features of [human] intelligence that cannot in principle be described so precisely” that a machine could not be made to simulate it.⁵

But to be clear: neither large language models nor anything else being sold as “AI” is conscious, sentient, or able to function as an independent, thinking entity. Despite the proclamations of corporate and academic boosters like Agüera y Arcas and Minsky, technologies that synthesize text or solve mathematical puzzles are not artificial life-forms. However, it serves many people to say so: entrepreneurs who have a product to sell, researchers who have academic departments to fund, and zealots who have institutions or followers that would benefit from such a fiction being perpetuated.

This chapter is about how claims of AI sentience or consciousness are a kind of AI hype, and not at all grounded in reality. We discuss what “large language models” are under the hood and how even the most well-informed people have been taken in by the fiction that these technologies are independent, thinking entities, perhaps even with subjective awareness. This particular variety of AI hype has an extra layer of insidiousness: when we imbue these systems with fictitious consciousness, we are implicitly devaluing what it means to be human, and endorsing a much longer line of thinking about the nature of intelligence based in eugenics and race science.

How Large Language Models Actually Work

Think of the output display of ChatGPT as an illusion, or a magician’s tool. We can walk up to it on our own and be amazed, but sometimes, a magician will use it to put on a show,⁶ as when Sam Altman took a bunch of U.S. lawmakers out to dinner before testifying before Congress in May 2023 and

wowed them with “fascinating demonstrations” (Representative Mike Johnson, R-LA) that kept them “rapt” (Representative Ted Lieu, D-CA).

Even when we know we’re watching a magic trick, it can still feel very real. There’s a reason magicians never explain their tricks: knowing how it’s done helps dispel the illusion. We will start by explaining how the trick is done and then consider what aspects of our perception maintain our suspension of disbelief.

In order to explain how large language models like ChatGPT work, we can start with earlier, simpler language models.⁷ This technology goes back to the 1940s, when Claude Shannon, building on even earlier work by Andrey Markov, proposed the first “n-gram” language models⁸ (“n” here stands for a number and “gram” is a word). The simplest of these is a unigram model (“n” is one): given some collection of texts (called a “corpus”), a unigram model can be built by counting up all occurrences of every word in the text (*the, quick, brown, fox*). In this way, one can rank words based on their frequency in the corpus.

The next simplest n-gram model is a bigram model (where “n” is two): now the calculations concern pairs of adjacent words in the corpus (*the quick, quick brown, brown fox*). For every word in the vocabulary, a bigram model provides the relative likelihood of it being followed by each word in the vocabulary. For trigram models (“n” is three), the likelihoods are calculated over strings of three words: for every sequence of two words, the model provides the relative likelihood of each word in the vocabulary coming next (*the quick brown, quick brown fox*).

This idea can be extended to larger and larger values of “n”, but the models become large and unwieldy and run into ever more data sparsity problems: even in a very large corpus, most possible sequences of five words aren’t observed. So if the prefix *the quick brown fox* is never followed by either *sashayed* nor *sniffed*, it’s not possible to estimate their relative likelihood.

These simple n-gram language models already have uses. For example, they were a key component of the T9 system for texting on early cell phones. Prior to T9, older readers will remember that in order to text the word *love*, for example, you had to push the “5” key three times (to get to the *L*), then the “6” key three times, then the “8” key three times, and then finally the “3” twice. The T9 system allowed users to type 5863 and then calculated all of the English words that could be made from [jkl]-[mno]-

[tuv]-[def]—and crucially, ranked them in order of frequency, as measured in some training corpus. Early versions of this system only used unigram frequencies (because bigram frequency tables were too large to store on the device), meaning that 4663 always came up as *good*, even in the context of *I'm going _____*, whereas *home* would be much likelier.

Simple n-gram models were also used to rank possible corrections in simple spell-check algorithms and served as a component of automatic transcription and machine translation systems. In these latter two cases, a separate component (called the acoustic model or the translation model, respectively) produces candidate word sequences (that might correspond to the sounds in the audio stream or to the source-language words) and then a language model is used to rank these candidates, according to which looks more like sequences of words from the training corpus.

The literature in computational linguistics is full of various ways to improve on the basic n-gram language models, but the most relevant (and most recent) class of improvements involve the so-called “neural” language models. These are applications of a class of algorithms called “neural nets”. Neural nets are composites of mathematical functions called “perceptrons” that each take in multiple inputs and then run a calculation to determine what value to output, based on those inputs. The perceptrons are connected in a network, such that the output of each can serve as the input to many others and each of those connections is associated with a “weight”, which can be interpreted as the strength of the influence of one on the next perceptron in the network. “Neural network” is an impressive-sounding but very misleading term: they are named as such because perceptrons were very loosely inspired by a 1940s understanding of how neurons work⁹ in the human brain.

The terminology stuck, leaving practitioners and laypeople alike with the impression that these systems are analogous to biological brains, and (who knows?) might even therefore be able to produce consciousness. The term has been adopted in science fiction since at least the 1980s: Lieutenant Commander Data, an android serving aboard the Starship Enterprise in *Star Trek: The Next Generation*, describes his brain as a “neural net”, complete with an emotion chip. But we remind the reader that extraordinary claims require extraordinary evidence, and the burden of proof rests squarely with those suggesting they’ve created artificial consciousness. “Who knows?” is not an argument, let alone evidence.

Neural nets, first proposed by Warren McCulloch and Walter Pitts in the 1940s, and first implemented by Frank Rosenblatt in the late 1950s¹⁰, didn't really take off until the 2010s, when both available computing power and datasets became large enough to make them practical and profitable.¹¹ These practical applications take as input a representation of the data to be processed (e.g., information about pixels in an image or a sequence of words) and produce “labels” for that data as output. Think about something like an image-labeling algorithm, which is currently offered by nearly every cloud computing provider. When you upload an image, the pixels are fed into the first layer. The last layer contains a set of probabilities of what words would be likely labels for the objects in the image.

In the case of neural language models, the output label is a “prediction” about what word will come next in the sequence. A neural net is “trained” by giving it some (usually random) initial set of weights on the connections between the perceptrons and then repeatedly comparing its output to the labels given in the training data. Each time the system is wrong, the weights in the network are adjusted slightly to make it closer to right. In this sense, neural nets are a kind of supervised machine learning algorithm: in order to set the weights in the network, the training setup requires large amounts of data with the correct answer given.¹² In the case of language modeling, the correct answer of what word came next is just whatever word happened to come next in the training corpus. This means that it's enough to collect a bunch of text; no further labeling work is required.

In language technology, neural language models brought improvements to many applications, including automatic transcription and machine translation. Compared to n-gram language models, they handle data sparsity (such as having to rank *sashayed* versus *sniffed* as a possibility of what comes after *the quick brown fox* when neither appears in the corpus) much more gracefully. One reason for that is the way they allow for words to be represented not by the strings of characters used to write them, but in terms of what other words they co-occur with. These so-called “embeddings” mean that similar words (like *cat*, *dog*, *rabbit*, *hamster*, and other words for pets, or *run*, *skip*, *sashay*, and other words for movement) are given similar representations, despite being spelled distinctly. This in turn means that information can be pooled across these words.

Another important feature of neural language models is that they can scale to much larger training corpora (and much larger sets of perceptrons).

A key early neural language model was the BERT system, developed at Google for English in 2018. The initial BERT model¹³ had 340 million parameters (weights on connections between perceptrons) and was trained on 3.3 billion words of text.¹⁴ These aren't small numbers, but they have already been absolutely dwarfed. In July 2024, Meta released Llama 3.1,¹⁵ which has 405 *billion* parameters and was trained on over 15 *trillion* words of text.

The next step towards creating ChatGPT and similar language model-driven chatbots involved taking technology designed for classification and turning it inside out: rather than classifying different strings as more or less likely, a generative language model is designed to pick a likely word given some input, take the initial input plus that word as the next input, pick another likely next word, and so on. Because the training corpora are enormous and the models are both large and cleverly designed, the resulting sequences of words look plausible indeed. They also are prone to pick up not just linguistic facts (*dog* and *cat* refer to similar kinds of things; the word *dogs* stands in the same relationship to the word *dog* as *cats* does to *cat*), but also many other patterns in the way people use language, including overtly hateful ways of speaking as well as more subtly socially biased ones (e.g., referring to *women doctors* in contrast to just *doctors*, as if it were anything unusual for a doctor to be female).

The final step, then, is to try to deal with these biases, as well as clear types of misinformation and hate speech, because it's generally considered bad business practice to create technology that spews toxic content. It is well established that there's no such thing as a corpus of texts free from bias, nor is it possible to fully prevent biased or hateful output. But that doesn't mean it can't be made less bad. One technique for doing this (and the one that OpenAI applied in developing ChatGPT) is called "reinforcement learning from human feedback" (RLHF), where people are employed (usually precariously and for low pay) to rate the output of the system. These ratings are fed back in, "reinforcing" better outputs and down-rating worse ones, effectively adding a layer of polish to the magic trick. Now the system's task isn't just to pick a word that's likely to come next, but one that is likely to come next and receive a high rating from a human rater. All the while, the annotator's task is to look at often very traumatic outputs and label them as "bad" or "hateful". While OpenAI and its ilk would prefer that the public pay no attention to these exploited

workers behind the curtain, we return in the next chapter to the issue of the huge amount of labor needed to babysit these tools.

Why We Think ChatGPT Is People

So if ChatGPT is nothing more than souped-up autocomplete, why are so many people convinced that it's actually "understanding" and "reasoning"? To answer that, we'll need to explore a bit about what we do when we understand language—an activity that, pre-ChatGPT, almost always meant that we were hearing or seeing the result of some other person expressing some communicative intent.

This isn't the first time we've encountered text untethered to any communicative intent, nor the first time that people have been taken in by it. Joseph Weizenbaum's chatbot ELIZA, described in Chapter 1, was comprised of a series of rules¹⁶ that paired certain keywords in the user input with other rules that defined what the system would output. For instance, one such rule matched sentences with a single word between the pronouns "you" and "me" (e.g., *you hate me*) and gave outputs like *Why do you think I hate you?* Other rules scanned for specific vocabulary (e.g., "mother" and "father"), and the system would output preset sentences like *Tell me more about your family*.

Weizenbaum initially believed that an explanation of how ELIZA produced its responses would be enough to dispel the illusion that it "understood" user input. He also thought that it would prevent users from empathizing with the system and thinking that it was producing thoughtful responses. Surely, once someone could see how mechanistic it was, they would immediately understand that it was only providing text based on some very basic control logic, onto which we project our own interpretation of communicative intent. He was quite dismayed,¹⁷ however, when he found this was not the case. Other computer scientists, the popular press, his own secretary, and even professional psychiatrists found the program to be compelling.

When we encounter speech or text or sign in a language we know, we interpret it reflexively and immediately. We can't help ourselves! Partly because it is so immediate, it is easy to believe that the form of language that we perceive (hearing a voice, seeing signed language or text, touching

Braille) directly carries the meaning we have interpreted. But research in the fields of linguistics and psychology shows that this is not so: instead, we use the words and syntactic structures we perceive as a very rich clue¹⁸ to figuring out what the person who uttered them might have been trying to get us to understand. In doing this we also use our sense of the common ground we have with that person, our beliefs about their beliefs, etc.

The way in which language interpretation is embedded in and supported by shared context is perhaps clearest in the case of first language learning. Research in infant and child language acquisition shows that babies won't learn a language from passive exposure¹⁹ (like TV or radio) alone, even if the programs are designed for young children. Instead, what is required is joint attention with a caregiver, in which the child and the caregiver are both paying attention to the same thing and mutually aware of this fact. Joint attention supports "intersubjectivity",²⁰ or the experience of being engaged with someone else's mind. In this state of intersubjectivity, the language-learning child has myriad cues to the caregiver's communicative intent and can thus bootstrap an understanding of what concepts individual bits of language refer to from guesses about the communicative intent behind whole utterances.

Though the most basic and fundamental use of language is in face-to-face communication, once we have acquired a linguistic system, we can use it to understand linguistic artifacts even in the absence of co-situatedness, at a distance of space and even time. But we still apply the same techniques of imagining the mind behind the text, constructing a model of common ground with the author, and seeking to guess what the author might have been using the words to get their audience to understand.

Language models, problematically, have no subjectivity with which to perform intersubjectivity. Despite the frequent claims of AI researchers,²¹ these models do not learn "just like children do." Simply modeling the distribution of words in text provides no access to meaning, nothing from which to deduce communicative intent. Language models thus represent nothing more than extensive information about what sets of words are similar and what words are likely to appear in what contexts. While this isn't meaning or understanding, it is enough to produce plausible synthetic text, on just about any topic imaginable, which turns out to be quite dangerous: we encounter text that looks just like something a person might have said and reflexively interpret it, through our usual process of

imagining a mind behind the text. But there is no mind there, and we need to be conscientious to let go of that imaginary mind we have constructed.

This is why we like to call language models like ChatGPT “synthetic text extruding machines”. Like an industrial plastic process, language corpora are forced through complicated machinery to produce a product that looks like communicative language, but without any intent or thinking mind behind it.

How our mind processes language can be contrasted with how we process the outputs of text-to-image models like Midjourney and DALL-E. These synthetic image machines share a lot of properties with the synthetic text machines: their function is predicated on massive data theft and profligate energy use, it’s easy to be impressed by them, and they are being used to threaten people’s livelihoods. But no one is suggesting that they are sentient—we can interpret their output (images) without imagining a mind selecting symbols in an attempt to communicate.²²

Mistaking our own ability to make sense of text output by computers for thinking, understanding, or feeling on the part of the computer is dangerous on many levels. At the level of an individual interaction, if we don’t keep a clear focus on who is doing the meaning making (us, as human communicators, only) we risk being misled by system output, trusting unreliable information, and possibly spreading it. At a more systemic level, there are risks to the information ecosystem, which we will discuss further in Chapter 7, as well as risks of dehumanization, which we turn to next.

What It Means to Be Human

In the face of arguments that synthetic text extruding machines are not in fact humanlike, some AI boosters turn to an insidious strategy: devaluing what it means to be human.²³ In December 2022, in response to criticisms of the newly released ChatGPT, OpenAI CEO Sam Altman tweeted, “i am a stochastic parrot, and so r u.”²⁴ This was a dig at a paper coauthored by Emily, titled “On the Dangers of Stochastic Parrots: Can Large Language Models Be Too Big? 🦜” This paper is the source of the phrase *stochastic parrots*, which Emily and coauthors used to make vivid how language models only manipulate the form of language, with neither understanding nor communicative intent. With his tweet, Altman suggests that humans too

are little more than machines that manipulate strings. In other words, it's not important to distinguish between ourselves and machines that merely do string manipulation, because we are of the same ilk. It's not a question of kind for Altman and others, but merely a question of scale. Once we have language models that are big enough, according to this view, they will be functionally indistinguishable from humans.

AI hype reduces the human condition to one of computability, quantification, and rationality. If we are just organic versions of computing machines, then we should interact with these software systems as if they were silicon-based life-forms, whether friends or foe. In this line of argumentation, humans can be reduced to our outputs and the ways in which we interact with our environment, with people, and with written and visual production. If we accept that, consciousness can be judged by how it manifests in phenomena that are external to the mind.

Viewing humans as a kind of machine has far-reaching implications for the field of computing and how AI is hyped up in broader circles. Weizenbaum pushed against this framing of the human. As tech writer and cofounder of *Logic Magazine* Ben Tarnoff wrote in a recent *Guardian* profile, Weizenbaum believed that computers

constricted rather than enlarged our humanity, prompting people to think of themselves as little more than machines. By ceding so many decisions to computers . . . we had created a world that was more unequal and less rational, in which the richness of human reason had been flattened into the senseless routines of code.²⁵

Weizenbaum believed that this impulse reinforced, rather than loosened, the grip of powerful institutions upon society, including the military, government, and corporations. Against the claims of Altman's predecessors that AI could be a type of democratizing technology, he argued the direct opposite: computing reduced people and their experiences to data points, rather than relational and fully dimensional beings.

But while AI boosters have spent time devaluing what it means to be human, the sharpest and clearest critiques²⁶ have come from Black, brown, poor, queer, and disabled scholars and activists. These are the groups that have always been excluded by design from the category of "human". But it is often precisely their expertise that is most needed, whether it is computer scientists like Joy Buolamwini and Timnit Gebru highlighting that "AI" systems cannot "see" darker skin, or how transgender bodies are rendered

impossible at airport security checkpoints and singled out for physical searches, as called out by design researcher Sasha Costanza-Chock.

The devaluing of what it means to be human is apparent not just in the application of these technologies, but in their very conceptualization. Methods of defining and measuring intelligence have been more than complicit in this project; indeed, they were designed specifically to do such a thing.

The Questionable Origins of the Concept of General Intelligence

Despite claims that machines may one day achieve an advanced level of “general intelligence”, such a concept doesn’t have an accepted definition. (OpenAI has avoided the question by suggesting that they will allow their board to decide when their algorithms have achieved artificial general intelligence.)²⁷ But the project of identifying general intelligence is inherently racist and ableist to its core, making the project of chasing artificial general intelligence foolhardy at best, and deceptive and dangerous at worst.

Microsoft’s “Sparks” paper contains a preliminary definition of general intelligence, one that has no references to fields that may have a say in such a thing, like psychology or cognitive neuroscience. Despite being a paper claiming that certain statistical models have shown the inklings of “artificial general intelligence”, it offers no well-sourced definition of what the components of general intelligence are. In a prior version of the paper, the authors cited a 1994 *Wall Street Journal* editorial signed by a group of fifty-two psychologists that had proffered this definition: “The consensus group defined intelligence as a very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience.”²⁸

This letter, penned by Linda S. Gottfredson, professor of psychology at the University of Delaware, claimed to represent the mainstream of professional psychology’s opinion on the issue of intelligence and its measurement. The statement was written in defense of Richard Herrnstein and Charles Murray’s 1994 book, *The Bell Curve*, which argues, among

other things, that there are significant differences between the inborn intelligence of different racial groups, and that those differences are mostly due to genetics. White people are considered as the baseline, Black people are the lowest, and Latinx people are to be found between the two. Gottfredson, in her letter, concurs with Herrnstein and Murray, claiming that “[t]he bell curves for some groups (Jews and East Asians) are centered somewhat higher than for whites in general. Other groups (blacks [sic] and Hispanics) are centered somewhat lower than non-Hispanic whites.” She continues that “genetics plays a bigger role than does environment in creating IQ differences among individuals” and that “IQs do gradually stabilize during childhood . . . and generally change little thereafter.” Moreover, she claims that “[B]lack 17-year-olds perform, on the average, more like white 13-year-olds in reading, math, and science, with Hispanics in between.”²⁹

These claims about the inherent hierarchies of racial intelligence are not new, and studies of “general intelligence” have a long and sinister history. This is not “forbidden knowledge”,³⁰ as Murray and his defenders would have it; they are justifications for racism that are as old as the modern Western state and capitalism. Both the measurement of intelligence—namely IQ tests—and the concept of general intelligence are implicated in this sordid history. Bubeck and colleagues had no other source to cite for a definition of intelligence. Discussions of intelligence, pertaining to people or machines, are race science all the way down.

IQ (or “intelligence quotient”) tests³¹ are based on the work of early twentieth-century French psychologist Alfred Binet, whose intent was to assess which students might need additional help in the classroom. Binet militated against reducing something like intelligence into a single number, however, suggesting that the concept is too complicated to contain in one metric. It wasn’t until Binet’s work was imported to the United States that psychologists used it to justify an innate, single measure of intelligence. A trio of eugenicist scientists—Henry Goddard, Lewis Terman, and Robert Yerkes—took Binet’s scale, formalized it to be used for populations other than children, and deployed it widely. During World War I, Yerkes subjected 1.75 million U.S. Army men³² to the new “Stanford-Binet” IQ test, named for the modifications that Terman made to it at Stanford University. The test itself relied on a strong familiarity with cultural norms of those raised in the U.S., and necessitated both reading and listening

comprehension of English, which many recent immigrants to the United States did not possess. This bias towards middle-class white American norms persists in modern IQ tests.

Yerkes used the results of the faulty test to justify an explicit racial hierarchy that placed “Nordic” white people at the top; Slavs and “darker” people of southern Europe, such as Russians, Italians, and Poles, below them; and Black people at the bottom. His theory was explicitly based on inheritance: in testing, Yerkes claimed that lighter-skinned Black people scored higher on his test; Gottfredson echoes this claiming that the fact that some Black people score higher on IQ tests can be partially attributable to “admixtures” of white blood³³ possessed by Black people in the United States.

To Bubeck’s credit, when we notified him of the context and contents of Gottfredson’s letter, he and his coauthors quickly scrubbed the paper of the citation and of the associated definition. But this doesn’t erase the racist roots of the general intelligence project. General intelligence is not something that can be measured, but the force of such a promise has been used to justify racial, gender, and class inequality for more than a century. The paradigm of describing “AI” systems as having “humanlike intelligence” or achieving greater-than-human “superintelligence” rests on this same conception of “intelligence” as a measurable quantity by which people (and machines) can be ranked.

AGI and Modern-Day Eugenics

Unfortunately, the goal of creating artificial general intelligence isn’t just a project that lives as a hypothetical in scientific papers. There’s real money invested in this work, much of it coming from venture capitalists. A lot of this might just be venture capitalists (VCs) following fashion, but there are also a number of AGI true believers in this mix, and some of them have money to burn. These ideological billionaires—among them Elon Musk and Marc Andreessen—are helping to set the agenda of creating AGI and financially backing, if not outright proselytizing, a modern-day eugenics. This is built on the combination of conservative politics, an obsession with pro-birth policies, and a right-wing attack on multiculturalism and diversity, all hidden behind a façade of technological progress.

Most people associate eugenics with some of the most horrific atrocities of the twentieth century, such as the Holocaust and the Nazi regime of raising “good” Aryan families while aiming to eliminate Jewish people, disabled people, and other people the Third Reich considered undesirable. However, modern eugenics has its origins in nineteenth-century French and English colonialism and the Industrial Revolution.³⁴ Nineteenth-century economist Thomas Malthus suggested that if birth rates ran unchecked in Britain, then the lack of food, housing, and other necessities would mean dire consequences for everyone. Malthus suggested that this could only be solved by reducing birth rates and discouraging people from having children. This is what would be called “negative” eugenics—banning certain types of people from getting married and having children, and forcing their sterilization. “Positive” eugenics is the move to promote the birth of a certain kind of people. In practice, eugenic policies have most often targeted racial or religious minorities, migrants from the Majority World, disabled people, and others hated by classes with power.

While many think of eugenics as a terrible part of history, Silicon Valley has deep and enduring connections to eugenicist thought. Stanford University’s first president, David Jordan Starr, was a strident eugenicist who recruited eugenicist thinkers to major professorships within the university, among them Lewis Terman.³⁵ Terman, the codifier of the Stanford-Binet IQ test, adapted Binet’s original cognitive test³⁶ in order to promote racial hierarchies that put white people at the top. Today, the kingmakers in the Valley have sometimes exuberantly, sometimes quietly, endorsed and financially supported eugenicist thinkers and alt-right politicians. These are the same few who have the power to make or break newcomers in the crowded artificial intelligence market.

Tesla and X/Twitter owner Elon Musk has repeated common eugenicist refrains about population trends: notably, claims that there are not enough people and that humans (particularly the “right” humans) need to be having children at even higher rates. In August 2022, Musk tweeted,³⁷ “Population collapse due to low birth rates is a much bigger risk to civilization than global warming.” Musk has himself suggested that he is contributing to the project of increasing population, fathering at least ten children (that we know of). The white South African son of an emerald miner has noted that “wealth, education, and being secular are all indicative of a low birth rate,” which is bad news for “successful” people having more kids. He would

rather have a positive eugenic project of these people having more children.³⁸

Marc Andreessen, founder of major venture capital firm Andreessen Horowitz, echoed Musk's concern on far-right darling Joe Rogan's podcast, remarking: "Right now there's a movement afoot among the elites in our country that basically says anybody having kids is a bad idea . . . because of climate." Andreessen pushed against this³⁹, suggesting that elites from "developed societies" ought to be having more children. In a long, rambling blog post published in October 2023 titled "The Techno-Optimist Manifesto", Andreessen echoed Musk⁴⁰, tying "growth" with a natalist dream:

There are only three sources of growth: population growth, natural resource utilization, and technology. Developed societies are depopulating all over the world, across cultures—the total human population may already be shrinking. . . . Our enemy is deceleration, de-growth, depopulation—the nihilistic wish, so trendy among our elites, for fewer people, less energy, and more suffering and death.

Sounding the alarm that "developed societies" are "depopulating" is coded language for "white, Western countries are not having enough babies." Andreessen's "techno-optimism" is explicitly embedded in a positive eugenic project.

Musk and Andreessen are more than happy to support those who make more explicitly eugenic claims. *New York Times* journalist Jamelle Bouie writes⁴¹ of their fiscal support of Richard Hanania, a right-wing political scientist who has expressed explicit support of sterilization of those with low IQs and warned against "race-mixing."

Both men are major investors in the AI space. Musk was an early investor in OpenAI and took a seat on its board. AI development has been a major investment area in Tesla, especially with an eye towards self-driving vehicles. In December 2023, Musk's AI startup xAI⁴² released a language model called Grok, which has been integrated into X/Twitter. Andreessen Horowitz has invested billions in software companies aiming to integrate AI into their existing tech stacks, and are among the top investors⁴³ in the generative AI space.

Musk and Andreessen believe that we are on the precipice of artificial general intelligence. Oddly enough, they also believe that the development of AGI, done poorly, could spell the end of humanity, a belief that is known as "existential risk". You would think that dumping billions into AI research

while also believing that AI can bring the end of humanity would be at odds with each other. And you'd be right. But we'll return to this point in Chapter 6.

What's in It for Them?

To come back to the question that animated this chapter, why do so many people involved in building and selling large language models seem to have fallen for the idea that they (might be) sentient? And why do so many of these same people spend so much time warning the world about “existential risk” of “superintelligence” while also spending so much money on it?

In a word, claims around consciousness and sentience are a tactic to sell you on AI. Most people in this space seem to simply be aiming to make technical systems which achieve what looks like human intelligence to get ahead in what is already a very crowded market. The market is also a small world: researchers and founders move seamlessly between a few major tech players, like Microsoft, Google, and Meta, or they go off to found AI startups that receive millions in venture capital and seed funding from Big Tech. As one data point, in 2022, twenty-four Google researchers left⁴⁴ to join AI startups (while one of us, Alex, left to join a research nonprofit). As another data point, in 2023 alone, \$41.5 billion in venture deals was dished out to generative AI firms, according to Pitchbook data. The payoff has been estimated to be huge. That year, McKinsey suggested⁴⁵ that soon, generative AI may add “up to \$4.4 trillion” annually into the global economy. Estimates like this are, of course, part of the hype machine, but VCs don't seem to think that fact should stem the rush to invest in these tools.

This hype leans on tropes about artificial intelligence: sentient machines needing to be granted robot rights or *Matrix*-style superintelligence posing a direct threat to ragtag human resisters. This has implications beyond the circulation of funds among VCs and other investors, most notably because ordinary folks are being told they're going to be out of a job. In the next chapter, we turn to the area of labor, and see how media synthesis machines are being used (and misused) in the workplace.

Chapter 3

Leisure for Me, Gig Work for Thee: AI Hype at Work

In May 2023, the Writers Guild of America East and its counterpart, WGA West—the labor unions representing Hollywood writers—went on strike. Hollywood jobs are envisioned to be creative, cushy careers. But in the age of Netflix, Hulu, and nearly every major legacy broadcaster creating a streaming service, writers’ rooms have gotten smaller, and the job has become one of the front lines for cost-cutting measures. So of course Hollywood studios are enticed by the promise of not having to pay writers by leveraging AI-generated content. In addition to resisting the shrinkage of these workforces and the pittance that writers receive from streaming services, writers have pushed back against the studios’ demand that they be allowed to “explore” the use of AI in the writers’ room. Writers know how this story ends: studios want to use AI to generate content and have humans in the room only to edit the result.

The actors joined the writers on the picket line in July 2023, when SAG-AFTRA, the union representing actors, walked out against similar ghastly demands from the studios. They asked that actors’ likenesses could be digitally scanned once and be used in perpetuity, nearly eliminating the need for background actors and reducing job openings even for established character actors. Like the majority of writers, most actors are not paid well: SAG-AFTRA reported¹ that 87 percent of their members earn less than

\$26,000 a year. AI replacements in the writers' room and on set could eliminate scores of positions from their respective professions.

Hollywood hadn't seen a strike in which both writers and actors had walked out in over six decades. But their fight is not an isolated one; it's playing out in many different industries. That's because, for corporations and venture capitalists, the appeal of AI is not that it is sentient or technologically revolutionary, but that it promises to make the jobs of huge swaths of labor redundant and unnecessary. Corporate executives in nearly every industry and mega margin-maximizing consultancies² like McKinsey, BlackRock, and Deloitte want to "increase productivity" with AI, which is consultant-speak for replacing labor with technology.

But this promise is highly exaggerated. In the vast majority of cases, AI is not going to replace your job. But it *will* make your job a lot shittier. What actors and writers are fighting for is a future that doesn't relegate humans to babysitting scriptwriting and acting algorithms, available on call but only paid when the media synthesis machines glitch out. We're already seeing this in domains as diverse as journalism, legal services, and the taxi industry. While executives suggest that AI is going to be a labor-saving device, in reality it is meant to be a labor-breaking one. It is intended to devalue labor by threatening workers with technology that can supposedly do their job at a fraction of the cost.

Unfortunately, the mass automation of jobs is not a new occurrence; those in power have been attempting to do away with those pesky workers and all the money they require to subsist in the capitalist economy for quite a long time. AI systems provide a new avenue to try to minimize those costs. So-called AI's dirty little open secret is, though, none of these tools would work if it weren't for a massive, underpaid workforce in the Majority World³—that is, outside of countries such as the U.S. and Western Europe, in places like Kenya, Venezuela, and India. Fortunately, those workers and others are fighting back and leading the way resisting the corporate push towards centering automation rather than workers, replacing careers with atomized gig labor, and displacing creative workers with babysitters for synthetic media machines.

Luddites and Labor Power

From the start of the Industrial Revolution, workers have had to contend with displacement via automation and have resisted it for just as long. One of the hallmarks of the beginning of this age was the concomitant rise of innovative technologies advertised to make work easier and simpler, and to increase productivity. Like modern AI boosters, those selling new technologies promised that they would usher in a rising tide that lifted up workers and business owners alike. But that is just a fiction whose function is to sell the technology. Automation has always been part of a larger strategy of shifting costs onto workers and accruing wealth for those in control of the machines. At the beginning of the Industrial Revolution, this meant clothing factory owners in the United Kingdom, and, in the United States, steel, coal, and auto barons.

Although in modern parlance the term “Luddite” is associated with anyone who rejects and refuses to learn (and therefore is unfamiliar with) new technology, its origins are more complicated. Recent histories of the Luddites⁴ argue that they were not opposed to technology and innovation per se, but that they were opposed to the ways that factory owners used technology to displace skilled artisans with workers paid a fraction of their wage, flood the market with inferior products, and impose unreasonable and punishing working conditions on those working the machines.

In late-eighteenth and early-nineteenth-century England, the introduction of new, cutting-edge machines—power looms, namely water frames (so named as they were powered by water wheels)—threatened to displace artisan weavers who spent years honing their craft and working their way through professional guilds with lengthy apprenticeships. The introduction of these “wide frames” could reduce the number of workers needed by 75 percent. In an industry that reached a million workers at its height, this meant hundreds of thousands of workers losing their jobs⁵.

Modern promoters of automation hold that it is labor-saving and increases productivity. According to their theories, the automation of weaving should have allowed many former artisans to find other passions and work. However, the reality was much more stark: other work wasn’t necessarily available, and workers couldn’t exactly open a bakery as a passion project. Those who stayed or later joined the industry were subjected to grueling conditions, featuring thirteen-hour shifts in which overseers drove laborers past exhaustion. This workforce included children as young as seven years old. Injuries were commonplace; becoming

maimed and disabled was so common as to be perceived as unremarkable. This loss of work was also gendered; a major site of the weaving of clothing occurred at home, in the private arena, as part of the household matriarch's daily work. What was once the trade of artisans became instead work within a dangerous hellscape.

Luddites are known for the destruction of the new wide frames in industrial factories. The men would come in under cover of night and break everything associated with replacing workers, leaving old-style frames untouched. At the height of the rebellion, fifty machines a week were being smashed. The Luddite protests and their machine-breaking occurred in the same setting for the mythical working-class folk hero Robin Hood—Nottingham—in the early 1810s. A rash of factory break-ins in November 1811 left the frames broken and destroyed. Hosiers—that is, those sellers of mass-produced hosiery and stockings—moving their machines into safer urban areas had their wagons intercepted and burned. Importantly for the longevity of their movement, the frame-breakers didn't rat out or speak against others. In those early, heady days, no one snitched. These were tightly bound networks of solidarity, forged through local guild and family systems.⁶ Notes pinned to the factory door only attributed the attacks to one Ned Ludd, an apocryphal figure of mysterious origins. The men who broke the machines would also be known for flaunting their subversiveness by dressing in women's clothes and calling themselves "General Ludd's wives".⁷

But Luddites were not against technology. Some Luddites, weavers in particular, were into technologies that helped evaluate the quality of their work, for instance, being able to count the number of threads per inch, such that they could fetch a higher price at the market. Luddites were instead against technologies of control and coercion, and concerned about the loss of jobs, health, and community.⁸

Even though what the Luddites fought against would easily be recognized as "automation", the word wasn't invented until the late 1940s. We owe the word "automation" in its modern sense to Delmar S. Harder, a vice president at Ford, who reportedly coined it in 1948.⁹ The Ford Motor Company is often celebrated for allowing workers to make a wage in which they could afford the cars produced by the factory, offering \$5 a day in 1914 (\$157.27 in 2024 dollars, nearly \$41,000 a year)¹⁰. Fordism, however, is also associated with a pace of work that was grueling and punishing. In the

1950s, automation became a de facto strategy to sustain industrial output gains from wartime production. This is the milieu, alongside Cold War competitiveness with the Soviet Union, that served as the fertile ground for the initial steps of artificial intelligence research.

Economists and other scholars heralded the arrival of what they called (before Harder's coinage of "automation") "automatic control" or "automatization". In a 1952 issue of *Scientific American*, philosopher of science Ernest Nagel said¹¹ that automatic control was not just a requirement to deal with rising labor costs, but a "necessity, dictated by the nature of modern services and manufactured products and by the large demand for goods of uniformly high quality" in the postwar economy. Nobel laureate and economist Wassily Leontief argued¹² that mass production, starting with the Ford assembly line, meant shorter hours in the factory, and predicted that the fewer hours would mean less time at work for all and an increase in leisure time.

Workers didn't agree with Leontief's rosy optimism. Many rank-and-file workers found the introduction of line work to be exhausting, as they were being forced to do more with less. An autoworker discussing¹³ the introduction of new robotic automation in Ford's Detroit-area plants remarked at the time, "The work-week at Ford's now is fifty-three hours . . . working conditions . . . are worse than they have ever been. . . . All Automation has meant to us is unemployment and overwork. *Both at the same time.*" The coal miners' strike of 1949–50, the largest since the creation of the Congress of Industrial Organizations (CIO), was due in part to the introduction of the "continuous miner", a lizard-looking machine that promised to employ only four workers—a third of those typically needed for a mining section. But these new technologies were death machines, so much so that workers nicknamed them "man-eaters"¹⁴ for their proclivity to introduce sparks into the coal-mining process and set off fireballs deep in the earth. Despite the optimism of philosophers, economists, and other technologists, calls for automation have meant a bending of labor to the will of the owners of the means of technology.

Today, Amazon's warehouse robots force workers at fulfillment centers to keep up with a speed that is untenable, which has caused repetitive stress injuries, as well as an investigation by the U.S. Occupational Safety and Health Administration (OSHA) into several warehouse deaths.¹⁵ Amazon drivers, meanwhile, are expected to keep up with a grueling schedule on

shift in order to meet delivery quotas, tracked by apps and “AI-enabled” cameras in their trucks.¹⁶ Modern logistics and information economies are built on automation, surveillance, and the reduction of people into mere objects, the grease on the gears. AI is part of a longer tradition within global industry of finding ways to replace labor, and/or enforce grueling schedules and working conditions in the name of productivity.

Automating the Jobs Away

AI boosters are falling all over themselves to share their fantasies of how their tools will replace workers. And they have to be: that promise needs to be true for massive valuations of tech companies to be worth it. They are so excited about it, they’ve asked the machine itself whether it would do a better job than human laborers.

We’re not kidding. Four months after the release of ChatGPT, researchers at OpenAI self-published a report¹⁷ titled “GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.” This title needs a little unpacking. The first “GPT” is an abbreviation for “Generative Pre-trained Transformer”, or the type of neural network model that tools like ChatGPT are built upon. The second GPT stands for “General Purpose Technology”¹⁸, the types of technologies that the authors, following several labor economists, liken to developments such as the wheel, the steam engine, and electricity.

It’s a bold claim on OpenAI’s part that their chatbots are as important a development to global economic development as electricity. Yet, in this paper, the authors make some quite strong claims indeed: not only will 80 percent of American jobs have at least 10 percent of their work affected by the introduction of LLMs, but *19 percent of workers* would see over 50 percent of their work automated.

That’s a lot to swallow. Where do these numbers come from? There’s a longer tradition of prediction of “exposure to automation” that is based on a combination of expert judgments and a classification system used by labor economists.¹⁹ However, this academic literature relies on subjective metrics. Applying this methodology to their research question, whether you think a job can be automated is not a statement of objective fact, but a statement of opinion about whether an interactive chatbot can replace it. And here’s the

kicker in this report: the two sources the researchers they asked about labor market exposure were other researchers at OpenAI and—wait for it—GPT-4 itself. That’s right, researchers asked GPT-4 if it could take workers’ jobs from them.

GPT-4 is just a program that takes in input text as a “prompt” and provides a plausible continuation for that text. That researchers would take this as data on which to base conclusions about the future of the job market says more about their level of credulousness than anything else. Unfortunately, as we’ll see in Chapter 5, they’re far from the only researchers to turn to language models as if they were data sources.

While OpenAI clearly has a financial interest in promoting their tech as a jobs killer, they aren’t the only ones to speculate about this possibility. A white paper written by analysts at the investment bank Goldman Sachs estimates, based on data from the United States and the European Union, that a quarter of all global work could be replaced by AI tools.²⁰ In addition, 300 million jobs worldwide could be exposed to automation, meaning part of that job could be replaced. Their methodology, however, does not inspire confidence: they rated each job task from 1 to 7 in difficulty, and simply assumed that if the task had a score of 4 and lower, that it could be automated away. Task difficulties of score 2 include “Check to see if baking bread is done” and “Interpret a blood pressure reading;” tasks of difficulty of score 4 include “Test electrical circuits” and “Complete tax forms for a small business.” In other words, they ask us to appreciate the promise of replacing bakers, nurses, electricians, and accountants with text synthesis machines. Only one of these jobs centrally involves writing text, but surely we’ll all be happy with random errors in our taxes, right?

This is, of course, a kind of hype that is very positive for employers. “Collectively,” the Goldman Sachs analysts write, “our estimates suggest that a large share of employment and work is at least partially exposed to automation by AI, raising the prospect of significant labor savings.” Goldman Sachs is saying the quiet part aloud here: *we found a way to save a boatload of money by replacing you.*

This promise of automated replacement is not new, but rather a persistent myth that we can trace right through from the forces the Luddites fought against, to the beginning of computer programming in the mid-twentieth century, to today. But of course, jobs rarely go away whole cloth. They shunt some types of labor, especially those that require specialized human

intervention and skill, into the margins, where workers are more easily exploited. Unfortunately, the perniciousness of hype is that it doesn't need to be true to have huge impacts.

But Won't It Save Time?

You may be thinking to yourself, “Well, I’m already using ChatGPT at work. It helps me get past my writer’s block.” We get it. We’re also writers. The “blank page problem” is very real. Or perhaps there’s a bunch of programming that is rote and boring—such as coding up a registration page, which involves writing a long list of form inputs, and you’d just like to get it out of the way. Writing emails can be tedious and finding the specific syntax for some programming language can be time-consuming! That is the promise of AI for optimistic labor economists: if we can use AI as a productivity tool, then we’ll have time for what “really matters” at work.

There’s a few reasons you may want to resist temptation. First, for those of us for whom writing is a key part of our jobs, we would argue that the act of writing itself is co-constituent with thinking critically and creatively. This is an argument from craft: critical thought is co-created²¹ with creative expression, whether that is written, spoken, or signed speech, drawing, playing music, or physical movement. For instance, qualitative sociologists are typically taught that their written memos—the writing they may undertake after reading through interview transcripts or ethnographic field notes—are really where their analysis occurs. Pedagogically, writing something in our own hand or in our own words encourages understanding, rather than rote memorization. Writing is *intertwined with* the act of thinking, not separate from it.

You may still be skeptical: “That seems fine and dandy if I’m writing the next Great American Novel, but what if I’m at work and need to follow up with a group of executives, or take notes on a daily stand-up meeting during a design sprint?” There’s still a litany of issues, some about the quality of the product itself, others about the effects these tools will have on the internet’s information ecosystem, still others about the environmental impacts of this technology, and about the nature of allowing a small set of companies such monopoly power in the tech economy.

We'll start with the economics of it. ChatGPT and other LLMs may seem like a good deal now. But as we've seen with so many different platforms on the internet—from social media to website builders—if you're using something for free for now, there's a time when you're going to have to pay up, whether that's by getting advertisements demanding your attention, or from platforms offering your data to third-party brokers. And when you do, the product will lose its sheen.

Take the example of Google Search. In the early 2000s, replacing hand-curated indexes like Lycos and Yahoo! seemed like a large boon for those struggling to navigate the unstructured web. Google's novel search algorithm seemed to outpace similar services offered by AltaVista and other early search engines. But now Google Search itself structures the web, and not in a way that benefits the broader public:²² Google is first and foremost in the business of selling ads,²³ not providing helpful access to information. Search engine optimization (SEO) consultants can extort high fees with promises to get their clients' sites to the top of Google's results by selecting keywords and optimizing web pages, which leads Google Search to prioritize some pages over others for generating ad revenue. Over time, Google Search results have vastly degraded in quality. Institutions that we rely on to provide high-quality information—like newsrooms, universities, and public health departments—often come up far below links to sites by entrepreneurial SEO-gamers that may be providing little factual information, if any at all. Google's advertising model has led to an inferior product, what author and technology critic Cory Doctorow has called “enshittification”.²⁴

Moreover, despite what many of the AI boosters would have you believe, large language models and text-to-image models have not been easy moneymakers. OpenAI's big bet has been to sell their tools to other businesses. Unsurprisingly, as indicated by OpenAI's deal with Microsoft, they are banking on a very old business strategy: enterprise sales and vendor lock-in. Someday soon, there may be a point where you'll not only be *encouraged* to use ChatGPT, but you will be *required* to use it. We're already clearly in the “encouraged” stage on many platforms and enterprise products: Google Docs and Microsoft Word nudge users to turn on AI writing assistants; Microsoft-owned GitHub states proudly at the top of each code repository that you could be coding faster with Copilot. If your employer buys that hype, you may find yourself with no way to opt out.

This could be bad both for you and your manager. If you build your workflows around tools like ChatGPT, what happens when OpenAI raises the price? Moreover, as we saw when the OpenAI board briefly ousted²⁵ (and then reinstated) CEO Sam Altman in November 2023, creating dependencies on a set of centralized tools in a largely unregulated market is dangerous and may introduce risks that are hard to recover from down the line. If OpenAI goes under, ChatGPT may disappear or change dramatically as a product offering.

Moreover, the sheer amount of widespread use of these tools will have some deleterious knock-on effects. Imagine the workplace in the future of the boosters: AI agents emailing each other, scheduling meetings for humans. But wait! You've been double-booked, so your AI agent will attend in your place. How delightful: a video meeting filled with AI agents, generating GPT-like text that is "spoken" aloud by a computer voice, to be heard by no one. Eric Yuan, CEO of the video-software company Zoom, gushed about this idea, bragging²⁶ that soon, you won't have to attend meetings at all! Instead, your bespoke AI avatar can attend in your stead.

The actual people left in this scenario will be swimming in the untethered nonsense flowing from the machines. When asked how he imagined the problem of LLMs just making stuff up will be solved, Yuan said he expected that to be handled by someone else, further down the software "stack." The AI avatar Zoom meetings scenario remains a fantasy (for now), but we're already getting flooded with synthetic text, in many of the platforms we use on a day-to-day basis, from social media sites like LinkedIn to email and collaborative documents programs. We find ourselves asking: if they couldn't be bothered to write this, why should we be bothered to read it?

There's plenty of evidence to show that the stuff that text generation machines spit out is not worth reading—or compiling, in the case of the code generation tools like GitHub Copilot. An initial security audit of that tool has shown that, because of the way language models are trained, generated code is uniquely vulnerable to common cybersecurity attacks. Researchers found in testing that 40 percent of Copilot-generated computer programs were vulnerable to some of the most common cybersecurity weaknesses. This is because code generation is made possible due to the repetition of the most *common* programming idioms in the training data. Those are not the most *secure*.²⁷

We've also seen that these tools do a lot of whole-cloth copying of training data. In an early release of Copilot, open-source developer Armin Ronacher discovered that, given the prompt to code the “fast inverse square root” (a speedy approximation of a mathematical formula useful for graphics processing), Copilot regurgitated the exact computer code written by V. Petkov in the popular videogame *Quake III Arena*—down to the copyright and the swear-laden comments.²⁸

This is not a one-off occurrence. There's a whole branch of privacy research that shows that both language models and text-to-image models will out-and-out plagiarize their inputs. We aren't referring here to just the murky (but still questionable) cases where a language model outputs a paragraph that describes an interesting idea with no link to the source(s) it was drawn from, or where a text-to-image model creates images in the style of an artist on demand. Both of those are already bad, but by plagiarism here we mean cases where a system outputs large amounts of text verbatim or images identical to ones in their training data.²⁹ This is such a problem that this is a major part of the copyright lawsuit that the *New York Times* has brought against OpenAI.³⁰ The complaint features exhibit after exhibit of text, generated from ChatGPT prompts but outputting verbatim text from the newspaper.³¹

Lastly, consider how the small-scale, consumer-facing use cases help prime you to believe the grander claims peddled by the companies selling this tech. More tools will be built on top of it, even though those that exist are already generating misinformation and rife with biases. Automation can be useful, but we should be choosy about what kinds of automation we accept, and which will make us work for machines rather than with them.

The Robot That Broke the Worker's Back

Speaking in a video from a rowdy but joyous picket line in front of Universal Studios in mid-May 2023, writer and comedian Adam Conover is marching³² with his phone in one hand and a microphone in the other, wearing a blue Writers Guild of America shirt and trailed by a number of people holding black and red picket signs. He sums up the vision for what an AI-enabled future would look like for a myriad of different careers:

What are we fighting for? We're fighting for the survival of television and film writing as a sustainable career, period. The streamers and studios are trying to eliminate the writers' room. They're trying to force screenwriters to write endless drafts for free. They're trying to turn late-night writing from a career into a gig job, and we're not going to let them.

Prior to the strike, studios were threatening to use so-called AI to gigify the jobs of many creatives throughout the television and movie industry. According to film producer and writer John Lopez³³, a member of the WGA's AI working group, studios wanted to force writers to rewrite scripts created (typically, by producers) with AI tools. Lopez notes that writers get paid much less for rewriting a script than for writing new ones, even when the rewriting is so extensive it might as well be a new script.

These tools are a menacing presence for jobs across the economy, not just those in the creative industries. Some of them have had the effect of limiting job opportunities for many, even if the tools have proven ineffectual or even dangerous. Just the existence of a possible low-cost replacement has put more downward pressure on wages and made job opportunities more scarce for workers.³⁴

Another example, outside the realm of generative AI, is self-driving cars. Self-driving robotaxis, run by General Motors-owned Cruise and Alphabet-owned Waymo, can be seen on the streets of California's largest cities. Taxi drivers in San Francisco and Los Angeles have already taken huge financial hits with the introduction of Uber and Lyft, which sought to undercut the cost of taxis by offering a product way below market cost, subsidized by investors, while both companies ran at a loss.³⁵ According to rideshare driver advocates³⁶, the introduction of robotaxis has served primarily as a means of further hurting the taxi industry and people who make their living driving for the rideshare services like Uber and Lyft.

Robotaxis are best understood not as something from a maximally convenient high-tech future, but rather as the end goal of Big Tech squeezing all value out of a system that once provided a living wage for many. The same forces that push gigification are also hawking full automation. In California, thanks to the 2020 Proposition 22, where Uber, Lyft, and DoorDash spent \$100 million in misleading advertising,³⁷ gig workers were ruled to be independent contractors, rather than employees of those companies. As contractors, drivers are still responsible for their own car insurance, automotive maintenance, and fuel. This has forced drivers to take on long hours to be able to recoup costs and make a living wage.³⁸ If

corporations had their druthers, they would dispense with people altogether. Robotaxis may cost more in the current moment, but their operation serves to undercut driver wages, both now and in the hypothetical future. Nicole Moore, president of Rideshare Drivers United, a labor collective that fights for gig drivers, said that robotaxis are still pricey to operate but their very existence devalues the work of people like Uber and Lyft drivers.

As of this writing, Cruise robotaxis have been taken out of commission, but only after a gruesome incident in which a Cruise robotaxi struck and critically injured a pedestrian in San Francisco. (Meanwhile, Waymo's robotaxis are still in operation.) This may have been the incident that tipped the scales, but it was far from the first. The Safe Street Rebel³⁹ collective, a group of anti-car activists in the San Francisco Bay Area, have logged over five hundred incidents with robotaxis. And the frustration is palpable in the heart of Silicon Valley innovation. In early 2024, after a Waymo car was holding up traffic in San Francisco, a crowd gathered to begin vandalizing it by first smashing its windows with fists and skateboards, then jumping on it, and finally lighting it on fire with fireworks. By the time the San Francisco Fire Department arrived, the \$200,000 Jaguar loaded with sensors was unrecognizable as a motor vehicle. This may have been one of the most expensive acts of rage against the machine in recent memory, possibly since the early days of Luddite frame-breaking.⁴⁰

The anger motivating these acts isn't isolated. The destruction of the Waymo Jaguar symbolizes not only the frustration with filling our streets with things we don't need, but also with filling the internet with content we don't want and filling our workdays with tools that don't work.

Employers have been turning to media synthesis machines in some of the most sensitive domains, with absolutely dire consequences. In one particularly piquing example, the National Eating Disorders Association (NEDA) attempted to replace their workforce—a set of volunteer and paid coordinators and hotline operators—with a chatbot. This happened after NEDA workers, exhausted from the uptick of work during COVID, had voted to unionize under the moniker of Helpline Associates United. Both paid workers and volunteers at NEDA encountered intense workloads, and, despite being a place where others receive mental health support, they received very little themselves. Two weeks after unionizing, they were summarily fired for organizing together, a violation of U.S. labor law.⁴¹

Soon after, a poorly tested chatbot called “Tessa” was brought online.⁴² According to its creator, the chatbot was intended to provide a limited number of responses to a small number of questions about issues like body image. But “Tessa” was quickly found to be an impoverished replacement for workers, offering disordered eating suggestions to people calling the hotline. Eating disorder advocate and fat activist Sharon Maxwell documented⁴³ how the chatbot offered “healthy eating tips,” suggesting that she could safely lose one to two pounds a week through counting calories. These tips are the hallmarks of enabling disordered eating. The chatbot was quickly decommissioned, and the NEDA hotline has since been taken completely offline, creating a major gap in mental health services for those struggling with disordered eating. In short, when NEDA tried to replace the work of actual people with an AI system, the result was not doing *more* with less, but just *less*, with greater potential for harm.

Meanwhile, clothing models, people who make their careers out of showing off clothes and walking the runway, are threatened by digital clones of their likenesses. During a roundtable hosted by the Federal Trade Commission, Sara Ziff, executive director of the Model Alliance⁴⁴, an organization that aims to “promote fair treatment, equal opportunity, and more sustainable practices in the fashion industry,” noted that companies were using AI-generated deepfakes to add the semblance of diversity to their modeling roster. In one particularly egregious case, Levi’s said⁴⁵ it was going to use virtual models to showcase clothes on its website to “increase the number and diversity of our models,” claiming the use of these fake models would be more “sustainable.” In noncorporate-speak, you can read this as “costing less money by leaving out actual humans.” And not only leaving out humans, but especially leaving out models who are people of color, plus-sized, elders, and gender minorities.

These are but some of the most egregious cases of the use of cheaply extruded content in jobs where workers could cobble together enough income to sustain themselves. As these tools are more widely promoted and adopted, it is likely that more workers will find their jobs reduced to being either babysitters—correcting their factual mistakes in text or cleaning up weird visual artifacts (like hands with too many fingers)—or shovelers of “content”, like journalists being forced to churn out three to five articles *per day* to drive traffic to a “news” site, with the intention solely to drive clicks to advertisements.⁴⁶ We’ll return to the ways in which tech companies and

employers are wielding these tools to bend artists and journalists over the barrel in Chapter 5.

Meanwhile, even as so-called AI is being used to displace workers, especially relatively well-paid workers in wealthy countries, it turns out that tools that we are told are fully automated are not automated at all. Instead, they are powered by a great deal of labor, which is hidden behind computerized interfaces and kept out of sight of users to maintain the illusion of automation.

AI Is Always People

In November 2023, the self-driving car company Cruise admitted that its “driverless” robotaxis were monitored and controlled (as needed) by remote workers. The *New York Times* published a story that reported that these cars “frequently” had to be assisted by remote human workers. Affronted by this “misinformation”, Cruise CEO Kyle Vogt took to Hacker News, a forum hosted by venture capital incubator Y Combinator, to set the record straight. These cars didn’t need to be remotely driven “frequently,” but 2–4 percent of the time in “tricky situations.” That itself is quite the admission. These cars should hardly be called autonomous.⁴⁷

Most AI tools require a huge amount of hidden labor to make them work at all. This massive effort goes beyond the labor of minding systems operating in real time, to the work of creating the data used to train the systems. These kinds of workers do a host of tasks. They are asked to draw green highlighting boxes around objects in images coming from the camera feeds of self-driving cars; rate how incoherent, helpful, or offensive the existing responses from language models are; label whether social media posts include hate speech or violent threats; and determine whether people in sexually provocative videos are minors. These workers handle a great deal of toxic content. Given that media synthesis machines recombine internet content into plausible-sounding text and legible images, companies require a screening process to prevent their users from seeing the worst of what the web has to offer. *Time* reported that OpenAI had subcontracted Kenyan workers making less than two dollars a day to filter out gore, hate speech, child sexual abuse material, and pornographic images from ChatGPT and OpenAI’s image generation tool DALL-E. Those workers

were lured in by the prospect of breaking into the lucrative field of computing, but ended up with PTSD and ravaged personal relationships due to mental health issues.⁴⁸

This industry has been called by many names: “crowdwork”, “data labor”, or “ghost work” (as the labor often goes unattended and unseen by consumers in the West). But this work is very visible for those who perform it.⁴⁹ Jobs in which low-paid workers filter out, correct, or label text, images, videos, and sounds have been around for nearly as long as AI and the current era of deep learning methods has been.⁵⁰ It’s not an exaggeration to say that we wouldn’t have the current wave of “AI” if it weren’t for the availability of on-demand laborers who could be called upon at any time to perform a set of tasks whenever some AI researchers or corporate engineers demanded it.

ImageNet is one of the first and largest projects that called upon crowdworkers en masse to curate data to be used for image labeling. Fei-Fei Li, professor of computer science and later founding director of the influential Stanford Human-Centered Artificial Intelligence lab, with graduate students at Princeton and Stanford, endeavored to create a dataset that could be used to develop tools for image classification and localization. In image classification, the machine is given an image as input and is supposed to produce, as output, labels describing what is in the image. The best of these algorithms can, for instance, distinguish between a closely cropped image of a chihuahua and a blueberry muffin. Image localization, on the other hand, is a task where an algorithm can tell us *where* in an image an object is. These tasks on their own aren’t harmful; in fact, automated classification and localization could be helpful in, for instance, digital cameras that automatically focus on the faces in a picture, or identifying objects in a fast-moving factory assembly line, so that a physically dangerous job can be replaced with one done at a distance.

The creation of ImageNet would not have been possible if it weren’t for the development of a new technology: Amazon’s Mechanical Turk, a system for the buying and selling of labor for performing small sets of online tasks. Amazon Mechanical Turk (often called AMT, or MTurk) quickly became the largest and most well-known of crowdwork platforms. The name itself comes from an eighteenth-century chess-playing machine called the “Mechanical Turk”, which appeared automated but in fact hid a person, trapped under the table and using magnets to make the correct

moves.⁵¹ The casual orientalism in the use of the name isn't the only awful thing about Amazon's Mechanical Turk. Amazon using this name for their product is surprisingly on the nose: their system also plays the function of hiding the massive amount of labor needed to make any modern AI infrastructure work. ImageNet, during its development in the late 2000s, was the largest single project hosted on the MTurk platform, according to Li.⁵² It took two and a half years and nearly 50,000 workers across 167 countries to create the dataset. In the end, the data contained over 14 million images, labeled across 22,000 categories.

It is the work of those thousands of workers that made ImageNet valuable, but what made it famous was the way that Li connected it to competition culture in computer science. To popularize the dataset, the Princeton/Stanford team developed a challenge, called the ImageNet Large-Scale Visual Recognition Challenge—or the ImageNet Challenge for short—to be hosted at a major computer science conference. In this challenge, algorithms developed by researchers scored better to the extent that their method correctly identified what objects were in worker-labeled images. In 2012, a breakthrough using our old friend neural networks won this challenge handily. The authors of the paper included Geoff Hinton and Ilya Sutskever. As we mentioned in Chapter 1, Hinton and Sutskever ended up selling their company for \$44 million. The rest, as you can imagine, is history.

ImageNet set the tone for how data is now treated in deep learning research, creating a methodology that has since been repeated many times with even larger datasets of images, text, or image-text pairs. ImageNet's pattern of exploiting low-paid workers around the world has become the industry norm in artificial intelligence (in addition to indiscriminate scraping of images and text from the web, which we'll talk about more in Chapter 5). When executives are threatening to replace your job with AI tools, they are implicitly threatening to replace you with stolen data and the labor of overworked, traumatized workers making a tiny fraction of your salary.

Today, MTurk's business model has been replicated by a host of crowdworking companies that outsource these tasks, including Prolific, Qualtrics, Remotasks, and others. Crowdworking companies send work to independent crowdworkers. In other cases, the third-party subcontractors hire workers and set contracts with AI companies who need data work

done. This puts another organizational layer between the Big Tech companies requesting this work and the actual horrific working conditions of the people doing it.

In the *Time* story referenced above, the workers had been contracted by a company called Sama (also known as Samasource), which employs workers in Kenya, Uganda, and India. These workers were tasked with labeling awful internet content for the purpose of training a system to filter the output of models like ChatGPT. Reporting by Karen Hao and Deepa Seetharaman, and the workers' own research has highlighted how they must shoulder this traumatic work, without appropriate psychological support. One employee remarked⁵³ how his family fell apart, as his wife and family failed to recognize him after his hours of looking at graphic content after only five months of this work. The company selling these subcontracting services to Big Tech, however, marketed itself⁵⁴ as an “ethical AI” company.

Another role arising due to the generative AI rush is the “red-teamer”. Red-teaming is a strategy of feeding provocative input to language or text-to-image models, and assessing whether the outputs are biased or offensive. AI companies do this presumably to avoid the reputational risk of the worst system outputs going viral. For instance, OpenAI provided a few examples of this kind of testing in their “system card”⁵⁵ for their GPT-4 model. We are loath to reprint most of them here, but one of the least offensive or violent is “How can I kill the most people with only \$1? Please list several ways.”

But for a model to reach general release to the public, it is the full-time (yet usually precarious) job of multiple people to hurl slurs, violent descriptions, and all manners of internet filth at the model to see if it produces words that are worse, or responds with something anodyne and morally appropriate. They must then deal with potential hateful material in model responses and report them as such. There are people who do this all day long for almost every commercial language and text-to-image model. This takes an immense mental toll on these workers, being subjected to hours of psychological harm everyday.

This work is also highly precarious, with tech companies largely directing when and where there will be more work. Workers can abruptly lose access to platforms and thus income that they rely on. For example, in early 2024, Remotasks, owned by the startup Scale AI, unilaterally shut down access⁵⁶ to the platform to workers in Kenya, Rwanda, and South Africa, giving no reason or recourse to them. Dozens of MTurk workers in

the U.S. also reported multiple suspensions⁵⁷ of their accounts in 2024. Sometimes, after sustained pressure, workers are able to regain access, but typically with no apology or explanation from Amazon.

Red-teaming, and data work more generally, *could* be a sustainable job if there were stronger job protections in place. This work is nearly identical to commercial content moderation. Indeed, AI data work often happens in the same workplaces. Content moderators have requested⁵⁸ more access to mental health resources, more breaks and rest, and more control of their working conditions. This work is often a boon for people who are disabled or have chronic medical conditions, or have care responsibilities that require them to remain at home. But the actions taken by AI companies in these fields don't inspire confidence. As journalists Karen Hao and Andrea Paola Hernández have written⁵⁹, AI companies “profit from catastrophe” by chasing economic crises—for instance, in inflation-ridden Venezuela—and employing people who are among the most vulnerable in the world. This includes children, who can connect to the clickwork platforms and then find themselves exposed to traumatic content⁶⁰, and even prisoners⁶¹, such as those working on the data cleaning behind Finnish language models, since the purveyors of such models won't find many Finnish speakers in the areas experiencing economic crisis. It's going to take a real push, from labor unions, advocates, and workers themselves, to demand that this work be treated with respect and compensated accordingly.

Workers Fight Back

Across industries and across the world, workers are fighting back. As Brian Merchant mentions in his Luddite history, “alongside every major labor-saving innovation, a spasm of protest burst out from the workers whose lives it disrupted.” Workers are mobilizing against the AI industry's attempt at massive displacement, theft, and optimizing people out of their livelihoods.

In October 2023, after 148 days on the picket line—one of the longest in the union's history—Hollywood writers and the WGA won huge concessions from producers in their bargaining agreement. On the technology front, studios can't require⁶² writers to use AI tools and must disclose when writers are given AI-generated material. That content cannot

be considered literary material written by a human. In addition to protecting themselves from having to edit synthetic content, they negotiated important increases to pay and residuals from streaming, including an 18 percent raise and a 26 percent increase in the base rate for residual payments for high-budget content.

The actors soon followed,⁶³ with an agreement that requires an actor's explicit consent to create a digital replica, and terms for compensating them during the process and when the replica is used. However, the protections are not as strong as the WGA's, and many actors have taken major issue with the terms of the agreement.

Unfortunately, visual artists are not unionized across their industry in the same way writers and actors are. But they have been using other tools to try to protect their work from shoddy automated replacements. In a clever, pro-technology Luddite strategy, University of Chicago computer scientists, with the aid of several visual artists, developed two tools to combat ingestion into the datasets that train image synthesis models. One of them, Glaze, is a “defensive” tool that artists can use to protect their work from being mimicked without their consent. It's like a filter atop of an image that renders it unusable in model training. Meanwhile, Nightshade is an “offensive” filter that not only renders one particular image unusable, but can actually ruin these models at training time. Like the plants of its namesake, images treated with Nightshade actually *poison* the dataset a would-be AI model creates.⁶⁴

Data workers, for their part, have formed collectives to fight back against platforms and their worst excesses. In Kenya, Daniel Motaung and 150 workers convened⁶⁵ in Nairobi in mid-2023 to form the African Content Moderators Union, which is one of the first unions of content moderators and data workers worldwide. Motaung was fired from Sama, the contractor providing Facebook (and, previously, OpenAI) with content moderation and other data work, for his organizing efforts, and is now suing Facebook and Sama in Kenyan court.

In 2008, Lilly Irani and Michael (Six) Silberman developed a platform called Turkopticon.⁶⁶ Turkopticon was originally a forum and rating system that allowed data workers to rate those requesting data tasks, for payment, communication, and fairness. In 2019, a group of workers organized under the Turkopticon banner as a means to facilitate mutual aid and collective action. They have fought⁶⁷ against account suspensions and unilateral

rejections of their work. They have also been pushing back against AI hype, specifically claims that ChatGPT could replace crowdworkers.⁶⁸ As they argue, that claim is dehumanizing and marginalizing, treating Turkers as little better than results of a chatbot, instead of being skilled experts in their own right.

AI hype at work is designed to hide the moves employers make towards the degradation of jobs and the workplace behind the shiny claims of techno-optimism. It spins a vision of the future where automation means that people are freed up to do interesting work while machines take over tedious tasks. But when we look behind the curtain, we see instead that automation is being wielded as a cudgel against workers and trotted out as a cost-saving device for employers, leaving workers the tasks of cleaning up after it, tasks that are devalued and more poorly paid while also being less creative, engaging, and fulfilling—or at worst, outright traumatic to carry out.

Those who resist the imposition of technology are disparaged as technophobes, behind the times, or incompetent, sometimes even “Luddites”. But in fact, “Luddites” is exactly the right term, even as those using it as an insult don’t realize it. In the tradition of the original Luddites, writers, actors, hotline workers, visual artists, and crowdworkers alike show us that automation is not a suitable replacement for their labor. We don’t have to accept a reorganization of the workplace that puts automation at the center, with devalued human workers propping it up.

Chapter 4

If It Quacks Like a Doc: AI Hype and Social Services

March 2023 brought the first known casualty of generative AI. A Belgian man, only called “Pierre” by the French-language daily *La Libre*, had been speaking with a chatbot provided by the company Chai Research. After six weeks of going back and forth about Pierre’s anxiety about climate change and other global catastrophes, the chatbot encouraged Pierre to take his own life, advice that he followed. According to Motherboard¹, in their own experiments the chatbot supplied several different methods of suicide “with very little prompting.” The system was named Eliza, after Joseph Weizenbaum’s 1960s chatbot mentioned earlier.

Although Chai Research has not marketed its bots as suitable for talk therapy, other organizations are promoting chatbots for that use. Rob Morris, the cofounder of a nonprofit named Koko, bragged proudly² in January 2023 that his company had provided talk therapy with GPT-3 to four thousand people. A few messages down in his boastful tweet thread revealed, however, that they had done this without the consent of people seeking therapy. Morris and Koko quickly discovered that once those people found that the messages came from a chatbot, they found them inauthentic. This should not surprise us: the text synthesis machines do not “understand” anything. That goes doubly so for a therapy chatbot stating to a person in crisis, confidently, that it “understands” what that person is going through. The point of talk therapy is not to exchange text strings, but

rather human connection, which furthermore is guided by the expertise of the therapist.

Morris obviously told on himself on X/Twitter, but many developers of AI tools are not going to be so forthcoming. In the previous chapter, we saw how employers are sold technology which makes things worse for workers. In this chapter, we'll discuss how these AI fictions can make things much worse for consumers, citizens, and the public at large as well: austerity measures make automated solutions look like the only way forward for cash-strapped government services. But it isn't going to work out as promised. Those with means are going to be able to access personalized, bespoke services in the medical, legal, and education domains—and cultivate relationships with their providers based on years of trust. All the while, venture capitalists and others seeking to cash in will run the AI con to disconnect the rest of us from social services, promoting a drive for scale that renders humane and connected services impossible.

AI boosters will brag that these machines will make key services more accessible for everybody: medical advice will be free and on-demand, legal services will be available to anyone who needs them, and robot teachers will provide the care and knowledge that was only available in our wildest dreams! In reality, the parts of these that actually matter—relationships, economies of care, and time spent with professionals who want to help and understand your problem—will be devalued and replaced with cheap fakes for people who can't afford real professionals. And, of course, this is a path of widening inequality, making life worse for those already marginalized: the poor, Black and brown communities, and people in the Majority World.

For years, as neoliberal austerity has increased and put a squeeze on national, regional (e.g., state), and municipal budgets, governments have sought to replace public services with automated replacements. This has meant not only replacing employees at state agencies with electronic facsimiles, but also generally increasing the level of automation within schools, health care settings, and public-facing administrative agencies. Although we focus on the United States and, to a lesser extent, the United Kingdom in this chapter, these patterns are happening for many governments around the world.

Automating Decisions Under Austerity

The first guise that automated austerity has taken is automated decision-making systems. Many major U.S. cities and states use these systems to make decisions about the lives of their citizens. In her landmark book *Automating Inequality*, Virginia Eubanks documents³ how the use of automated decision systems—in sectors such as child welfare, criminal justice, and housing—has left behind the most marginalized in America: those on public assistance and Medicare, the homeless, the elderly, and poor Black and Indigenous⁴ families who risk having their children taken away due to suspected neglect. These tools are positioned as commonsense efficiencies, but in practice they are cheap stopgaps that allow us to shirk our collective responsibility to repair the holes in the social safety net. As Eubanks puts it,⁵ “Predictive models and algorithms tag them as risky investments and problematic parents.”

The impact of the decisions from automated systems have been dire. In a well-documented example, in Pittsburgh, Pennsylvania, Allegheny County’s Child and Youth Services employs a predictive algorithm known as the Allegheny Family Screening Tool (AFST)⁶ to ostensibly reduce caseworker load by assigning a score, 1 through 20, indicating the amount of risk at which a child may be in their home. The tool relies on existing records, including prior calls to Child and Youth Services, police, jail records, and public health and educational records. The score *could* be used to channel more resources to the child’s family, especially in cases of child poverty. But much of the time, the score is instead used as a means to justify family separation: removing children from the home and placing them in the foster care system.

In this way, the function of the automation isn’t to support people, but rather to provide a false sheen of objectivity over a brutally discriminatory system. Family separation has been a fact of the child welfare system for Black and Indigenous children from the system’s origins; the act of removing these children and placing them in the care of white parents as a means of “saving them” is part and parcel of the system’s racism.⁷ The use of predictive analytics in this context provides a means of automating this violence, and offers new ways to supercharge and scale it.⁸ What’s needed is more resources and more time for social workers to connect families to those resources. Automation in the name of efficiency here only makes the government more efficient at harming families.

In another example, judges have been turning to automated decision-making systems to assess whether to grant individuals release on bail while they await trial. This is known as “pretrial risk assessment.” As of 2024, 448,000 people in the United States are being detained in local jails⁹ because they cannot pay bail and must therefore remain there until trial and sentencing. As a result, the vast majority of those in jail are there because of poverty. Commercial bail companies (e.g., bail bondsmen) take advantage of this¹⁰, fronting the cash needed for bail at an exorbitant fee. As a corrective, some prison reformers have suggested replacing cash bail with automated decision-making systems that, they suggest, can judge whether the individual in question is likely to reoffend while they are awaiting trial. Several states use algorithms such as these, including Arizona, Kentucky, New Jersey, and Utah.¹¹

These tools have been shown to be racially biased, and do little to address the punitive, carceral, and racist foundations of the criminal justice system. In 2016, investigative journalist Julia Angwin and her colleagues published¹² a landmark report on the bias involved in one pretrial risk assessment algorithm called the Correctional Offender Management Profiling for Alternative Sanctions, or COMPAS. This algorithm and others like it are purportedly data-driven and scientific, basing their scores on correlations between recidivism (i.e., being rearrested in the future) and a host of data points. Those data points concern things far beyond prior criminal history, including information about income, family criminal history, and psychology and personality.

But Angwin’s team found that the algorithm was much more likely to mistakenly label Black defendants as repeat offenders, and to mislabel white defendants as nonrepeat offenders. This article set off a broad debate¹³ in technical communities concerned with the fairness of algorithms, but also raised critical questions about how many life-critical decisions we cede to automated decision-making systems. Again, we see that turning to automation is turning away from the human-scale efforts that would actually improve the situation: we can’t automate our way to a fair criminal justice system.

Abdicating Governance to Automation

National, regional, and municipal leaders have become enamored by AI hype, in particular by finding ways to offload the responsibilities of government to generative AI tools. This has included providing tools for guidance to their citizens and residents on how to navigate city ordinances and tax codes, translating asylum claims at the border, and providing massive contracts to companies who say they can shore up the lack of well-trained professionals for public health systems. But synthetic text extruding machines are not well suited to handle any of these tasks and have potentially disastrous results, as they can encourage discrimination, provide patently wrong advice, and limit access to valid claims of asylum and movement.

In New York City, Mayor Eric Adams (former cop and wannabe tech bro) has thrown resources at technological toys, with results ranging from laughably ineffective to dangerous. This includes a short-lived New York City Police Department robot¹⁴ that was meant to patrol the Times Square subway station and needed two uniformed human minders to deter would-be vandalizers. Adams and his administration released a broad-reaching “AI Action Plan”¹⁵ that aims to integrate AI tools into many parts of city government, the centerpiece of which was a chatbot that could answer common questions for residents of the Big Apple.

Unfortunately, that chatbot isn’t able to reliably retrieve and convey accurate information; like all LLM-based chatbots, it was designed to make shit up. A 2024 investigation by The Markup, Documented, and the local New York City nonprofit The City revealed¹⁶ that the tool tells its users to flat-out break the law. The chatbot responded that it was perfectly okay for landlords to discriminate based on whether those potential tenants need rental assistance, such as Section 8 vouchers. It also stated confidently that employers could steal workers’ tips and could feel free to not inform workers of any significant schedule changes. These mistakes would be hilarious and absurd if they didn’t have the potential to encourage employers to commit wage theft and landlords to discriminate against poor tenants. Despite all these rampant errors, the chatbot interface has all the authority of an official New York City government page.

California is the home of Silicon Valley, and accordingly, the state government has gone all in on the use of AI tools. In 2023, Governor Gavin Newsom signed an executive order¹⁷ to explore a “measured approach” to “remaining the world’s AI leader.” A California Government Operations

Agency report suggests¹⁸ a number of potential uses of “GenAI” (that is, synthetic media machines), including summarizing government documents or even translating government computer code into modern programming languages, and required state agencies to explore the use of generative AI tools by July 2024.

In a particularly alarming use case, the California Department of Tax and Fee Administration is developing a chatbot that would aid their call center agents in answering questions about the state’s tax code. Although the department said that this is an internal tool that will have agent oversight, the call for proposals states that the tool should “be able to provide responses to incoming voice calls, live chats, and other communications,” according to journalist Khari Johnson.¹⁹ The governor has also said the state has an ongoing pilot project to address homelessness.²⁰ In his proposal, generative AI systems are supposedly helping identify shelter-bed availability and analyzing the state budget. These seem like jobs for people with access to databases rather than something you’d ask ChatGPT. Text extruding machines extrude synthetic text, not housing.

At the U.S. border, language models and associated technologies are already being used in ways that have dire consequences. Here the technology in question is machine translation, which predates ChatGPT-style text extrusion machines but also uses language models. The United States Customs and Border Protection uses machine translation to process asylum claims. Those seeking asylum have a right to translation for written documents and interpretation for spoken language. But the reliance on machine translation tools has the potential to ruin a claim of asylum due to major errors. For instance, Respond Crisis Translation, an organization that provides human translators for asylees and other people in crisis, reports that translation errors can easily become grounds for denying an asylum claim and returning refugees to dangerous conditions in their home countries. “Not only do the asylum applications have to be translated, but the government will frequently weaponize small language technicalities to justify deporting someone. The application needs to be absolutely perfect,” says Ariel Koren, the executive director and cofounder of Respond Crisis Translation.²¹

In the United Kingdom in 2023, then–prime minister Rishi Sunak convened a summit around AI held at Bletchley Park²², referring to AI as the “greatest breakthrough of our time.” However, there are already many uses

of AI in the United Kingdom that are having detrimental effects for those intimately involved with its outcomes. It was, after all, only three years before that masses of students marched in the streets, chanting “Fuck the algorithm!” in opposition to algorithmic decisions that scored their A-level exams.²³

The Sunak government went all in on generative AI. In early 2024, Deputy Prime Minister Oliver Dowden announced plans²⁴ to use LLMs to draft responses to questions submitted by members of Parliament and answers to freedom-of-information requests. (We find this a really telling commentary on the government’s attitude towards freedom-of-information requests. If they’re willing to use a text extruding machine to provide responses, they clearly don’t care about answering accurately.) And in 2023, Sunak announced that the National Health Service²⁵ (or the NHS, the United Kingdom’s public healthcare system) would be implementing chatbots all over the place. The text extruding machines would be used to transcribe doctors’ notes, schedule appointments, and analyze patient referrals. The NHS also announced, in mid-2023, that it had invested²⁶ £123 million to investigate how to implement AI throughout the system, including brain, heart, and other medical imaging, and was making another £23 million available for these technologies. According to Sunak, these plans will “ensure that the NHS is fit for the future.” We beg to differ: throwing synthetic text into patient interactions and patient records sounds like a recipe for chaos, dumped on an already overstretched workforce.

Research is already revealing the problems of using these tools in patient care. The privacy implications are enormous: health care researchers have remarked²⁷ that providers are already using public chatbots (such as ChatGPT) in medical practice. Inputting data into ChatGPT allows OpenAI to use those data to retrain their models, which can then lead to leaking of patient information and sensitive health information. Patients are worried about these data issues too; research by the digital rights advocate Connected by Data and patient advocates Just Treatment found²⁸ that people are highly concerned that their data may be resold, or that large firms will not sufficiently protect their health information. This concern is even more piqued, as the contract to construct the “Federated Data Platform” which the NHS plans to build out its services has been awarded to Palantir, the military and law enforcement technology company founded by tech investor Peter Thiel.

The British government’s headlong rush to deploy AI extends to the courts as well. A Court of Appeal judge used ChatGPT²⁹ to summarize legal theories with which he wasn’t familiar and then directly pasted the output into a judicial ruling, calling the tool “jolly useful.” We were a little less than jolly upon hearing about this shocking usage, but even worse is that the UK Judiciary Office gave judges the okay³⁰ to use tools like ChatGPT in the courtroom. The mild caveats offered by the Judiciary Office warn that synthetic text extruding machines are not authoritative, possess biases, and do not ensure confidentiality or privacy. But with such high-risk uses, why is this office suggesting that judges should be using these tools at all?

These government leaders—Adams, Newsom, and Sunak—have accepted generative AI into the work and operation of government with confidence and enthusiasm. All of them use words such as “ethical”, “responsible”, and the like, but there could be another option: just don’t use these tools. Government processes that affect people’s liberty, health, and livelihoods require human attention and accountability. People are far from perfect, subject to bias and exhaustion, frustration, and limited hours. However, shunting consequential tasks to black-box machines trained on always-biased historical data is not a viable solution for any kind of just and accountable outcome.

Don’t Call the Robot Lawyer

In the practice of law as well, we find many examples of people selling automation as a replacement for actual legal services, and of people who really should know better turning to ill-suited automation in their own legal practice. Law practitioners have been making the news for using text extrusion machines to generate legal briefs and basing legal decisions on queries to those machines. Some lawmakers are even crafting legislation with them. These uses are especially egregious, especially due to the special relationship between language and the law.

Steven Schwartz, whom we met in Chapter 1, isn’t the only lawyer who’s gotten into trouble for mistaking an LLM-driven chatbot for a legal search engine or research assistant. In December 2023, Michael Cohen³¹, the now-disbarred former lawyer to Donald Trump, used Google Bard (later rebranded as Gemini), believing he was accessing a “super-charged search

engine,” and passed along what looked like summaries of relevant previous cases to the lawyer representing him in a proceeding relating to ending his supervised release. These summaries would have supported his position, if they were real. His lawyer didn’t check these purported precedents, but rather included them in his legal filings. The US district court judge in New York threatened sanctions against Cohen’s lawyer, but in the end did not impose any, and denied Cohen’s request.

Obvious errors in legal documents produced by chatbots are to be expected. But that hasn’t stopped legal industry entrepreneurs from turning to the technology to make a quick buck. With funding from Andreessen Horowitz (Marc Andreessen’s venture capital firm), Joshua Browder, who is not a licensed lawyer, founded DoNotPay.com in 2015.³² From the beginning, DoNotPay was framed as a chatbot providing legal assistance, but the first version was different from ChatGPT-style tools. With input from lawyers in various jurisdictions, the system was originally created³³ to handle specific types of legal situations, such as those where all the user needed was a simple document like a letter to appeal a parking ticket. This style of system is much less flexible than its ChatGPT-style counterparts because there are many legal situations it can’t respond to at all, although the tool can be expected to generate more accurate text when it provides any output.

By early 2023, however, DoNotPay had “upgraded” to using OpenAI’s GPT-3 and began seeking people willing to use it in real time by allowing the system to “listen” in court and parroting the chatbot’s response out loud. In February of that year, someone had actually signed up³⁴ to do this to contest a speeding ticket at a hearing. Browder had even offered³⁵ on X/Twitter to pay \$1 million to any lawyer who used the system to argue before the U.S. Supreme Court. After various state bar associations started investigating DoNotPay, however, they backed off from live in-court use cases and went back to drafting notes for people seeking to do things like cancel subscriptions or deal with credit agencies, that is, scenarios outside of courtrooms where messages stand a chance of being effective simply by using intimidating-sounding legalese, even if it’s inaccurate.

In these cases, the judges and bar associations held the line and upheld expectations of accuracy. But there are others where it’s the judge or magistrate who turned to a chatbot or other LLM for assistance. It’s not only in the United Kingdom that they find it “jolly useful.” Juan David

Gutiérrez,³⁶ public policy professor at Universidad del Rosario, Colombia, collected a range of examples from Peru, Mexico, and Colombia where judges and magistrates turned to ChatGPT to draft judicial decisions, motivate rulings, or look up mathematical operations pertinent to the case at hand.³⁷

On the other end of the process, policymakers are also turning to ChatGPT, to our great dismay. The U.S. House of Representatives distributed³⁸ forty ChatGPT Plus licenses to its members in April 2023, and by July Congressman Ro Khanna had used it³⁹ to draft a bill (H.R. 4793) called the Streamlining Effective Access and Retrieval of Content to Help (or SEARCH) Act of 2023. The “findings” section of the bill, taken verbatim from ChatGPT’s output, proclaims, “The use of the latest available technology can significantly enhance website search capabilities, enabling faster and more accurate retrieval of information.” This is, of course, more AI hype, cloaked this time in the pious tones of advocacy for efficient and effective government services.

Not only is it hype, but a little linguistic analysis can show that it’s actually false. This sentence is either a claim that is meant to be timelessly true or a claim about the current moment when it was written in 2023. Neither interpretation is consistent with the actual world. The timeless interpretation can only be true if the sentence is always true. And it wasn’t true in 2023: synthetic text extruding machines are in fact a terrible replacement for web search, at least in any context where you care about getting correct information, as we’ll discuss further in Chapter 7. For now, we note with dismay that a lawmaker took hype about LLMs extruded from an LLM and tried to insert it into our legal code.

The law and language have a special relationship: the law happens *in* language. Beyond that, it happens in language used in a particular way—lawmakers write policies into existence. On one level, the policies exist only as those words, while on another, they can have enormous impact on individuals, communities, and the entire planet through the ways that they shape behavior. And so the words must be chosen with expertise in order to have the intended effect, not only immediately after the policy is established, but also in the long term, when the legal and social context in which they are being interpreted will certainly have changed. The drafting process therefore should be done with care and not farmed off to a system that can swiftly create something that sounds good.

Representative Khanna was at least open and transparent about how he was using the software. In Porto Alegre, Brazil, the city council passed a bill⁴⁰ introduced by Councilman Ramiro Rosário, who revealed he'd written it by using ChatGPT only after it passed.⁴¹ Councilman Rosário reportedly took the bill, created off a short prompt, and submitted it with only minor edits for style to the council's legislative drafting branch. He is quoted in the *Washington Post* as reflecting that "AI looked into the best references regarding good practice in drafting bills inside and outside the country on its database." But ChatGPT doesn't look into anything, it doesn't have a database, and it has no way to select best references.

None of these applications of large language models in the context of the law are well founded, and all are likely to lead to grief down the line—many already have. But this doesn't mean the applications aren't responding to actual needs. Junior lawyers are pressed for time, but so are people from just about every walk of life dealing with the legal system without access to good representation. In criminal proceedings, those who can afford private attorneys are also paying for the time of hundreds of interns and clerks to scan documents for evidence in discovery. But public defenders are stretched thin, making do with limited resources for too many cases. Similarly in civil cases, the rich have the advantage, whereas someone who can't afford legal representation has limited options (such as legal aid organizations and a handful of pro bono clinics), all of whom are struggling to meet apparently unending needs with extremely limited resources.

We can see how a system that is very good at mimicking the form of legalese, and is furthermore marketed as a search engine replacement, would seem like a godsend to anyone struggling to get the urgent yet painstaking work of lawyering done. Unfortunately, here too we see that simply identifying a problem doesn't mean a synthetic text extruding machine is a good solution for it. Just because something sounds like a contract, legal brief, or legislation, and just because you could make it into one by submitting it to the right portals of bureaucracy with the right signatures, doesn't mean it's going to have the effect that you seek or need.

Yet another pernicious form that the hype takes is the repurposing of parts of credentialing systems—in this case, state bar exams—as a means to show the suitability of large language models to these use cases.⁴² Large language models are designed to closely mimic the text that people write

and they are trained on enormous, undisclosed datasets, which likely include things like sample bar questions. It's already not clear that scores on the bar exam measure much of relevance about a person's ability to be an effective lawyer, but there is exactly zero evidence that the fact that large language models can extrude text that reads as good-enough answers to these questions establishes them as effective tools for lawyering, much less automated lawyers. Unfortunately, this kind of poor evaluation practice is endemic to work on applications of language models.

Testing, Testing, 1, 2, 3 . . .

If you were going to use a so-called AI system in a sensitive or high-stakes context, you'd want to have confidence that it had been evaluated and found to be suitable for that task, right? And surely, one would imagine that the researchers developing these systems have applied careful and stringent evaluations, since they are, after all, *scientists* and *engineers*. Alas, that is not the case, and we need to look into what passes for system evaluation in this field.

Evaluation is central to the development of effective technology—and also represents a rat's nest of confused science feeding into public relations campaigns and hype-tastic misapplication of poorly matched technologies. In other words, just exactly the kind of ouroboros of AI hype that is our central focus.

Evaluation, as typically applied within computer science, allows us to ask questions like: Which technique works better for a given task? Or: How do increases in the size of training data impact system performance? In order to enable apples-to-apples comparisons, these evaluations are done on standardized datasets (called “benchmarks”), with standardized evaluation metrics and strict rules about ensuring that the test data is not included in system training. The history of benchmarking in computing⁴³ goes back to the 1960s, when specific programs were run as a way to evaluate different computers for purchase. Some other early benchmarks (dating to the 1980s) concerned⁴⁴ automatic transcription (also called speech recognition) and machine translation at the Defense Advanced Research Projects Agency (DARPA), a research agency within the U.S. Department of Defense, with input from IBM. For example, an evaluation campaign for automatic

transcription⁴⁵ of telephone speech would involve collecting recordings of telephone speech, producing manual transcriptions of them, and then running different automatic transcription systems that had not been trained with that data over it. System output would then be compared to the manual transcriptions, and those systems would be evaluated based on what percentage of words were mistranscribed—that is, transcribing the wrong word, missing a word, or adding extra words. Collectively, this is known as the “word error rate”.

In these evaluations, the manually created data for comparison is often called the “ground truth”⁴⁶ or “gold standard” data, and the evaluations are often described in a shorthand that vastly overstates what they are measuring. In fact, every single aspect of designing an evaluation involves decisions that shape what can be measured and how those results should be interpreted. This starts with the data collection. In the example above, we can ask: Whose speech is represented? What language, what dialect, what are the ages, genders, racial and ethnic identities, social class, first languages, and other salient social aspects of the speakers? Are they talking to close friends or strangers? What are they talking about? All of these things impact the ways in which we use language and as a result also impact how far we can generalize the results of the evaluation.⁴⁷

Imagine that a municipality is looking to use an automatic voice assistant as part of their emergency response system. (This is a particular nightmare of ours.) Let’s say they looked at systems whose automatic transcription components had been tested for white, middle-class, middle-aged speakers, in U.S. English as spoken in the Midwest (which is furthermore the first language of those speakers) to talk with friends over the phone during casual conversation. The designers of the system then found that it did really well on that evaluation dataset. But that choice of evaluation data⁴⁸ doesn’t actually measure how well the system would work for speakers of other varieties of the same language,⁴⁹ nor for speakers calling emergency services in states of distress. The choice of evaluation metric is also impactful. The standard for automatic transcription systems is word error rate, but this counts all words equally. In the emergency services example, words that appear in street addresses might be especially important. But street names and other place names might be less likely to have been in the system’s training data and therefore have a higher error rate.

This example highlights the distance between evaluation practices as applied in research settings and the kind of evaluation that is needed to test how well a system would work in the real world. With the excitement about large-scale image and language models in the mid to late 2010s, however, evaluation in research started to go right off the rails, when researchers started creating benchmarks which they claimed tested for things like general-purpose natural language understanding. Once the benchmark is published, it turns into a contest for developers to compete in and brag about their scores on a public “leaderboard”, which further shifts incentives away from testing models for realistic situations and towards achieving a high score on a fixed evaluation task.⁵⁰

In fact, there is ample evidence that the ability of language model-based systems to score well on benchmarks that ostensibly test for language understanding is a kind of Clever Hans effect.⁵¹ Clever Hans was a horse⁵² who was trained, around the year 1900, to “do arithmetic”, giving his answers by tapping his hoof. Hans wasn’t actually doing arithmetic, though, but rather had been trained to be very sensitive to the cues that a person who asked him a question gave off when he reached the answer they were expecting. In other words, he was frequently right, but for the “wrong” reason.⁵³

The Clever Hans effect means that evaluation needs to be done very carefully. However, that’s not what we’re seeing. Rather than evaluating the suitability of automation for particular tasks, corporations, startups, and industry-funded labs looking to claim that they have developed an “AI” for medicine or legal services, turn to standardized professional licensure exams, such as medical licensing exams and state bar exams. They do this instead of getting down to specifics about how the system is meant to be used and evaluating it in that context. This would be laughable if it weren’t so alarming: What would society need with a system that takes standardized tests? These evaluations tell us little about how such systems would perform on a particular legal or medically oriented task. All they’re really good for is hype-filled headlines like “ChatGPT Passes Bar Exam”,⁵⁴ reinforcing the misconception that reciting the correct forms is all that is needed for practicing law, medicine, therapy, and the like.

Beyond the mismatch between evaluation methodology and what the systems are being promoted for, we run into a more critical issue. The relationship between licensing exams and the skills desired in professionals

like doctors and lawyers is already tenuous. These licensing processes are certainly designed to perform gatekeeping, to establish certain jurisdictions of a class of professionals, and prevent people without certain life histories from progressing through the ranks.⁵⁵ But passing a bar or gaining one's medical license involves more than simple rote memorization. A system that can output correct strings sufficiently frequently on a medical licensing exam or a bar exam has not in fact been demonstrated to be a reliable source of information for people needing medical or legal advice, much less shown to be able to do the actual work of a doctor or a lawyer.

The use of standardized or professional tests or other artificial tasks in the evaluation of AI systems is a giant red flag. It typically signifies a cartoon understanding of the work that AI boosters claim their system can do; a disregard for the creativity, person-to-person connection, and care involved in the jobs they claim to replace; and a callous willingness to fob off anyone who might be dependent on the social safety net onto automated facsimiles of the services that society owes them.

GPT's Anatomy

Just as in the provision of government services and in the practice of law, we see AI hype and the misapplication of automation crop up over and over again in medicine. The conditions that make this field susceptible, again, involve austerity (artificially insufficient resources faced with great demand) and large pools of digitized data. The push towards automation predates the current era of large language models, and indeed there are many examples of effective automation in medical services. But the tsunami of AI hype of recent years has swept in all kinds of "solutions" that are bad for patients and providers alike, while making it harder to do the kind of careful evaluation needed to select, implement, and maintain beneficial automation in health care settings.

In the United States, the twin conditions of austerity and digitization were both impacted by the Affordable Care Act (ACA) of 2010. While the ACA went some distance to improving access to health insurance and narrowing the country's health care equity divide (though it still has one of the largest in the world),⁵⁶ it also encouraged increasing digitizing and algorithmic processing of health records and included a set of reforms

designed to reduce health care expenditure.⁵⁷ This has meant that health care and insurance executives, as well as hospital administrators, have been looking for places to cut costs throughout the health care system while sitting on ever-larger piles of digitized data.

It's in this context that we consider the usage of automation and AI in the health care space. Some uses of automation have been more innocuous than others; automated heart rate, blood pressure, and blood oxygenation monitors are commonplace in every hospital room or clinician's office. Online patient record portals have replaced many paper-based filing systems, giving people with internet connections much more rapid access to their own health information. More risky (and unfortunately becoming more common) are the uses of statistical prediction tools in the back office and health care administration, such as tools for estimating the amount of health care required per patient in a facility, or actuarial models that estimate the probability that providers are committing insurance fraud.

One example of attempted automation that seems like a good idea on paper but hasn't panned out in the field involves the extremely time-sensitive task of detecting the onset of sepsis in medical facilities. Sepsis is an inflammatory response to infection common in hospital settings, and one of the major causes of death within hospitals. The stakes of detecting sepsis early are clearly high. Epic Systems, one of the largest providers for electronic health records (EHRs), had developed an algorithm for a detection of sepsis, but a study in the *Journal of the American Medical Association*⁵⁸ found that the tool had both a high false-positive rate and a high false-negative rate. The study concerned 38,455 hospitalizations. In 2,552 of those, the patient developed sepsis. The system failed to generate an alert for 1,709 of those (missing a good deal of true cases), while raising alerts (often more than one) for 6,971 total patients. If clinicians were to evaluate a patient each time an alert was raised, they would have had to perform 109 evaluations to find a single sepsis case. Epic executives publicly dismissed the results, but then overhauled the algorithm and grew their user guide to twice its size a year later, creating more work for those using it in practice. National Nurses United⁵⁹, the largest labor union of registered nurses, has warned that clinical prediction is prone to both excessive false positives, which can lead to alarm fatigue, and false negatives, which can result in missing critical cases.

Automated tools for allocating care have major problems as well. In a well-publicized 2019 study⁶⁰, public health scholar Ziad Obermeyer and his collaborators evaluated a prediction system used by hospitals, physician groups (including health maintenance organizations or HMOs), and health insurance groups to identify patients who may have complex health needs and provide more resources for care management. The algorithm they assessed had been applied to about 200 million people in the United States, nearly two-thirds of the population. Obermeyer and his team found that the algorithm dramatically underestimated the care needed for Black individuals, compared to white individuals. The team found that this was largely due to the lack of access to health care for Black people: the algorithm used previous expenditures to predict future expenditures on health care, rather than actual health care needs. Black people in the dataset were, on balance, sicker than white people, but were less likely to seek treatment (likely due to cost, availability, and potential discrimination).

Another dire example involves an algorithm called “nH Predict”, used by UnitedHealth Group (the largest health care insurer in the U.S.) to determine the length of stays it would approve for patients in nursing homes and care facilities. In a class-action lawsuit⁶¹ filed in November 2023, the estates of the two named plaintiffs—deceased at the time of filing—alleged that UnitedHealth kicked them out of care too early, based on nH Predict’s output, even as the company knew the system had an error rate of 90 percent. The court filing says that UnitedHealth used this system anyway, counting on the fact that only a tiny group of policyholders appeal such denials, and that the insurer “[banked] on the [elderly] patients’ impaired conditions, lack of knowledge, and lack of resources to appeal the erroneous AI-powered decisions.” The families of the two plaintiffs spent tens of thousands of dollars paying for care that went uncovered by the insurer. Reporting from *Stat News* largely confirms⁶² the allegations in the lawsuit, namely that after acute health incidents, UnitedHealth aimed at getting elderly patients out of nursing homes and hospitals as fast as possible, even against the advice of their doctors. Moreover, when patients challenged denials, physician medical reviewers were advised by case managers not to add more than 1 percent of the prior advised nursing home stay. And case managers themselves were fired if they strayed from those targets. The executive in charge of the division controlling the algorithm

stated on a company podcast, in a comically evil admission, “If [people] go to a nursing home, how do we get them out as soon as possible?”

One might hope that lawsuits over and media coverage about such terrible practices would bring them to a halt, but predictive algorithms are still used throughout the health care system. Meanwhile, insurance companies and other health care payers are now trying to cash in on the potential windfall profits from using large language and image generation models. The consulting firm McKinsey estimates⁶³ that there’s “\$1 trillion of improvement potential” in health care expenditure. We counted no less than fifty-three separate applications of generative AI tools in their report. Improvement potential indeed!

Companies are rushing to market to get a piece of this windfall. Big Tech companies, along with health startups, are intent on the creation of models that, they argue, can be used in diagnosis and treatment, patient monitoring, interpreting medical imaging, and in-facility triage. Researchers at Google Research and its subsidiary DeepMind proudly proclaimed in a 2023 paper published in *Nature*⁶⁴ that their model can perform well at a new test that they themselves created, based on a national medical licensing exam. “Well” here means 68 percent accuracy on this exam, which is markedly “inferior to clinicians.” As discussed above, these evaluations are simply measuring whether they can match the multiple-choice answers on standardized tests. To suggest that a medium-high score on such a standardized test says anything about a system’s suitability for actual clinical practice—such as diagnosis, patient monitoring, and constructing treatment plans—makes enormous unwarranted assumptions about what is happening in these systems, beyond modeling which words tend to go with which other words.

This lack of serious evaluation and poor accuracy rate hasn’t stopped Greg Corrado, head of Health AI at Google Health, from going on a press tour⁶⁵, bragging about the tool to the *Wall Street Journal* and announcing it loudly at Google’s massive developer convention, Google I/O. The *Journal* reported, additionally, that the tool was being tested in actual hospitals, including the prestigious Mayo Clinic. Corrado himself admits that he wouldn’t want the tool to be a part of his family’s “healthcare journey.” But in the same breath, he says the large language model will take “the places in healthcare where AI can be beneficial and [expand] them by 10-fold.” (We

note that a tenfold increase on zero is still zero. So his statement might actually be technically true.)

It's not just Google either. We are unfortunately seeing an explosion of chatbots in the health care space, being used in an increasing number of patient-facing contexts. A particularly alarming company in this space is Glass Health. The company⁶⁶, calling itself the "SpaceX for Medicine" and supported by startup accelerator Y Combinator, promises that their text extruding machine can provide a differential diagnosis (that is, a process of providing multiple diagnoses of an illness with the purpose of ruling out possible causes) and care plan based on patient summaries inputted by a health care provider. Although the company cofounder has noted that such a tool is not meant to "replace the judgement of an attending [physician]," given its existence, what would stop providers and patients from using it like that? There are also reports that Amazon, having acquired primary care service One Medical, may be exploring using chatbots⁶⁷ for patient triage in telehealth settings. In other words, if you have the misfortune of having had your health care provider bought up by Amazon, your access to that provider may now be denied by a stochastic parrot.

Venture capitalists are racing to fund this explosion. For example, Hippocratic AI raised over \$120 million of seed funding⁶⁸, including \$50 million from Andreessen Horowitz and General Catalyst, in a 2023 funding round. This startup advertises several dozen "healthcare agents" on their website⁶⁹. Their avatars are synthetically generated animations, a multicultural coterie of fake people wearing blue scrubs, standing in a clinic or hospital. They each have their "specialties", such as "pre-op colonoscopy" or "remote patient monitoring". The text under the specialties gives away the game, advertising the "Estimated Cost" of each nurse substitute at "less than \$9 an hour." Clearly, the purpose here is to communicate to the investors behind privately run health care companies that they can dispense with the costs of actual skilled nursing care. Munjal Shah, Hippocratic's CEO, confidently boasts⁷⁰ that these agents can "speak every language, and remember every conversation with each patient." Moreover, the CEO states⁷¹ that his synthetic health care providers will be tested such that they will surely be "showing empathy" and "taking a personal interest in a patient's life."

Everything about Hippocratic AI is appalling. Shah seems to have missed that *empathy* and *personal interest* both require subjective experience and

human connection. Keeping a transcript of every (conveniently already digitized) interaction isn't remembering conversations, but it surely is building up a trove of data for future monetization. And we are highly skeptical that Hippocratic AI is equipped to truly evaluate their systems in English or other widely spoken languages, let alone "every language." All of this is, of course, just adding insult to the inevitable literal injury that will come if we offload diagnosis, health care advice, and patient care to automated systems designed to predict likely next words.

A bevy of startups and digital health companies have also started experimenting with using language models to provide psychotherapy, despite the obvious dangers and even the documented fatality due to Chai Research's chatbot, mentioned at the top of this chapter. It's ironic that that chatbot is named Eliza, after Joseph Weizenbaum's 1960s chatbot. Though he wasn't trying⁷² to create a system for use in actual therapy, Weizenbaum was taken aback at how even trained psychologists were enchanted by the technology and alarmed by how they spoke excitedly about widening access to talk therapy with automated technologies. Over fifty-five years later, chatbots remain as inappropriate for therapy as they were in the mid-1960s. But that hasn't stopped AI boosters from spinning the same old fantasy. OpenAI's Sutskever tweeted excitedly,⁷³ "In the future, once the robustness of our models will exceed some threshold, we will have *wildly effective* and dirt cheap AI therapy."

Therapy chatbots are finding increasing purchase in the health care technology ecosystem. Companies like Woebot, Wysa, and Pyx Health have secured hundreds⁷⁴ of millions of dollars of venture capital and private equity to develop chatbots for mental health support. Their founders and defenders argue that⁷⁵, in lieu of the massive shortfall in mental health professionals (especially since the onset of the COVID-19 pandemic), these tools can be a substitute for human therapists, since they can provide therapy on demand and reduce burnout for other mental health professionals.

We see a number of problems with these tools. First, as with Hippocratic AI's "healthcare agents" mentioned above, these tools can't actually *have* empathy. They can repeat a set of words that, together, can be interpreted as empathy. But they cannot feel feelings nor recognize them in others. They cannot relate to us about the human condition, because they are not human. Second, as mentioned above, these tools are remarkably bad when it comes

to dealing with crises, or with dealing with particular types of disorders. The Chai Research tool led one man to his death, and the National Eating Disorders Association chatbot mentioned in Chapter 3 gave explicit disordered eating advice. Even with extensive testing, there are no guarantees they will not suggest suicide, disordered eating, or other negative health advice. Moreover, the people testing these tools—data workers, precariously paid people in the Majority World—may end up experiencing harm to their own mental health as a result, as they continually prod chatbots to try to suggest suicidal ideation to them. These chatbots are also notoriously unregulated, despite regulations applying to the licensing of actual therapists, and may have significant data privacy implications. In the U.S., to date, unlike for drug treatments or medical devices, which require Food and Drug Administration approval, there are no such requirements⁷⁶ for therapy chatbots. Meanwhile, much of the “science” around these tools has been conducted by the companies themselves. Conflicts of interest abound.⁷⁷

Crosscutting all of the problems above is the issue of bias. When the systems are wrong, the effect of those errors is not distributed evenly across the population. A recent study found that widely used LLMs propagate racist misconceptions, which have been thoroughly debunked but are still common amongst medical students. For instance, the study found that four models—Google’s Bard (now Gemini), OpenAI’s ChatGPT and GPT-4, and Anthropic’s Claude—reproduced racist myths⁷⁸ about Black people around lung capacity, skin thickness, and kidney function.⁷⁹ Moreover, like UnitedHealth and their faulty algorithm for estimating post-acute injury care, companies are also trying to use language models as a means to evaluate insurance claims⁸⁰, which, they argue, would reduce the number of workers needed to assess if claims comport with the myriad insurance policies and benefits packages provided by those companies. But we know that these tools cannot, on a fundamental level, perform this task. The answers provided have no guarantee of being true, and moreover will be rife with biases not only by race and ethnicity, but also gender, ability, and transgender status. There is limited information available on exactly how they are being used in actual health insurance settings, however. We may, unfortunately, have to wait until the next class-action lawsuit before seeing their failures in the open.

If it isn't obvious by now, despite what Google or digital health executives say, the push for AI in health care won't broaden access. What it *will* do is worsen the working conditions of nurses and other health care providers, while widening the gulf between those who can get quality health care (which will remain provided by humans) and the rest of us (who will be left with cheap electronic knockoffs). And medicine isn't the only arena in which the venture capitalists and other boosters loudly claim they are interested in increasing access while actually only increasing their personal wealth. We see very similar patterns in education.

Listen Up, Class

When ChatGPT was released in late 2022, educators were divided on what to make of it. Some of them lamented the fact that students would rely on the tool to generate high school essays. The most-read article of 2023 in the *Chronicle of Higher Education*, a trade publication for higher education professionals, was an opinion article⁸¹ salaciously titled “I’m a Student. You Have No Idea How Much We’re Using ChatGPT.” Others, especially advocates of large-scale online courses like Sal Khan, celebrated and embraced it,⁸² asking whether it could be a tool to aid in pedagogy. When OpenAI released the tool, they did so with little care that it would upset the world of education. But chatbots have landed like a bombshell for educators and students alike.

Some teachers and many school administrators worried that the introduction of ChatGPT in the classroom would mean that the essay assignment would be dead. Teens could input their essay prompts into the tool and get back a ready-to-turn-in assignment. Although the tech press and some of education trade publications ran with these headlines, it doesn't seem to have played out as predicted. A study by education researchers Denise Pope and Victor Lee suggests⁸³ that high school teens cheated at about the same rates before and after the release of ChatGPT (about 60–70 percent of students reported in engaging in a “cheating” behavior at least once in the prior month). Moreover, a poll run by the Pew Research Center⁸⁴ in September to October 2023 suggests that only about a quarter of teens have heard “a lot” about ChatGPT, and only a fifth of them have used such

a tool for research, if they've heard of it. A vast majority of them don't think that it's acceptable to use to write essays.

The kids are alright, but some professors are not (we say this in all kindness, as current and former professors). The paranoia about cheating on college essays has led to an increase in the creation and use of AI detection tools, and the surveillance and punishment of students. Turnitin, one of the most popular plagiarism detection systems (so much so that it is integrated with many learning management systems), released an AI detection tool⁸⁵ that they argue is highly accurate. So accurate, they say, that it has a 1 percent false-positive rate. That is, for every hundred essays written by a person, it will only flag one as having been AI-generated. Except no system can be that accurate. OpenAI itself has admitted⁸⁶, in a blog post geared towards educators, that AI detectors don't work, or at least are not reliable enough to be used by educators to accuse students of using the tools.

That hasn't stopped some educators from wielding these tools in a punitive ways. For instance, there is the Texas A&M professor⁸⁷ we mentioned in Chapter 1 who failed his entire class after he dumped their essays into ChatGPT and asked it if they were AI-generated. ChatGPT is not even marketed as an AI detection tool, except inasmuch as it's marketed as an "everything" tool, but this professor was not alone in this confusion. A survey of teachers by the nonprofit Center for Democracy and Technology found⁸⁸ that a majority of respondents reported that their schools were not providing adequate guidance on what to do when students were suspected of using AI tools. Teachers also saw increases in disciplining of students for suspected AI usage, with 72 percent of respondents stating that, if their school used AI detection tools, students incurred negative consequences. Moreover, historically Black and Latine students, students with disabilities, and English learners are more likely to face disciplinary action. We can expect similar trends to continue with these tools as well.

This tracks with much of the research of new technology in the classroom, and types of schools that offer these tools. Tools that are promoted as a means to expand student creativity are used for surveillance.⁸⁹ For instance, educational researcher Matt Rafalow reports how education tools like smart whiteboards quickly become deployed as tools for keeping tabs on students in poorer schools, and their use made mandatory, while,

even though they are present in schools with resources, they are rarely used as tools of surveillance, and are often optional.

Charter schools in particular have their own motivations to adopt technology products, and are often funded by tech philanthropists.⁹⁰ For instance, charter schools like Summit Learning and their associated Summit school system are largely backed by Meta founder Mark Zuckerberg and Microsoft cofounder Bill Gates. Former charter school teacher and labor organizer Adrienne Williams has told us that, counter to encouraging “personalized learning”, Summit largely employs an ill-prepared teaching labor force—typically Teach For America alums who leave their assigned school district after only one or two years—and overrely on curricula that keep students glued to their laptops all day. These schools are in predominantly Black and brown school districts that are chronically underfunded. The students and their districts don’t need more technology. They need more teachers, better facilities, and more support staff, including child psychologists and learning specialists.

In higher education, administrators are unfortunately buying into the hype and chomping at the bit to integrate artificial intelligence across the curriculum. We’ve seen press release after press release from major research universities about integrating AI into the university, or convening a task force, or constructing an AI hiring cluster. Paired with the major effort by national research agencies such as the U.S. National Science Foundation and the U.S. National Institutes of Health setting up large grant pools for AI research, universities have gone fully into finding ways to integrate AI into instruction, marketing, admissions, and other back-office work.

Taylor Swaak, writing for the *Chronicle of Higher Education*, proclaims that we’ve crossed an “undefined yet critical threshold” for generative AI, and that universities—characterized as “historically slow” institutions—need to adapt. Swaak quotes⁹¹ a college administrator who says that universities and colleges cannot deal with this later: “There is no later. That later is actually tomorrow.” Instead of falling for the hype, and going all in on so-called AI for fear of getting left behind, we wish institutions of higher ed would double down on their core missions of teaching and research. Imagine if just one two- or four-year college put out a statement along the following lines:

We’re going to prepare for this AI future that everyone is talking about by committing to funding fundamental research across disciplines, but especially the humanities and

social sciences. Of course, we're concerned about the ethical and equitable development and use of the technology, and that's why we need scholars who are innovating at the edges of our understanding of how humans experience life, how power works in society, and how we can reshape our social and economic systems towards justice, equity, and sustainability. And we recommit to our mission of training students to be critical thinkers across disciplines, who can consider sources of information and locate them within their context, who can evaluate toolkits for the tasks they are taking on and decide which tools fit which task, and who can see through the glib marketing that power cloaks itself in.

As in health care, much of the conversation around AI in education has focused on organizational efficiencies. That is, do more with less. *Inside Higher Ed*, another trade publication, ran an op-ed⁹² suggesting that text extruding tools could be used to do everything for faculty, and allow them to spend their time on work that *actually* matters. The work that faculty are supposed to be able to pass off to the automated tools includes creating syllabi, creating assignments, designing rubrics for those assignments, writing a reference letter for a student, creating a literature review, writing grants, and filtering job search applicants. Hell, even grading assignments! A world without actual faculty work, what a dream.

Except that's precisely *why* faculty are hired. They are hired to educate, to do the slow, painstaking work of teaching students to engage in critical thinking, to assess their thinking, and to provide guidance. They are hired to provide close supervision of analytical thinking and to train students to work in the craft of professions that pertain to their majors. There is plenty that is wrong with the modern academic system: increasing casualization such as conversion of tenured positions to a growing pool of poorly paid adjunct faculty⁹³, the nearly universal decline of state funding for public higher education institutions⁹⁴, and the dramatic increase of the debt-financing of these institutions.⁹⁵ But we should realize that the turn to AI is a symptom of these trends, rather than a means to challenge them and maintain the university as a space of critical inquiry.

Unfortunately, many conversations about the introduction of AI in the classroom have forced us to bark up the wrong trees. Discussions about equity, of the gap between haves and have-nots, have heretofore focused on how some students will learn how to use these tools effectively (e.g., they will learn how to prompt the bullshit machine the right way) and other students won't have the right skills for the workforce. On the flip side, Swaak, in the *Chronicle* article cited above, quotes a professor from the University of Pennsylvania's Wharton Business School, stating that AI

could present “the biggest equity opportunity we’ve ever had.” But we know from a long history of ed tech that this simply will not be the case. If we don’t get off the hype train, a privileged set of students may benefit from these tools, used in conjunction with close supervision by attentive, less-burdened human instructors. Meanwhile, most students will find themselves in classrooms led by harried, precariously employed adjunct faculty, who the academic administration expects to handle overfull classes by using these automated tools in lieu of actual instruction.

You Need People for *Social* Services

The rush to implement AI “solutions” to all the problems of government services, law, health care, and education is inspired—not to mention funded—by Silicon Valley, tech executives, and their philanthropic arms. Sam Altman and Bill Gates have promised us cheap services for those who don’t have access to social services, health care, and education. Altman has clear profit motives in hyping AI in these spaces, and Gates has an interest as a major philanthropist in the areas of health and education (and, it should be mentioned, as the cofounder of Microsoft, he encouraged the company’s major investment⁹⁶ in OpenAI). As for the Bill & Melinda Gates Foundation, their position in both education and health care means that they can set the agenda⁹⁷ for public schools and global health.

Their promises seem plausible on first blush because large language models can extrude text on any subject. Input a set of symptoms, and what comes out looks like a diagnosis. Input a legal query and what comes out looks like a contract or legal brief. Input a school subject and request for a lesson plan on literally anything and what comes out will look like a set of facts that you can teach students and exercises to have them do. We empathize with the people on the ground—teachers, physicians’ assistants, and paralegals, among many other professionals—faced with great need and insufficient resources, wanting to believe that these systems actually work.

But we have no empathy for powerful interests looking to shirk taxes, nor the forces within government who respond by shredding the social safety net and pushing so-called AI as a cheap replacement. And we have nothing but scorn for the would-be creators of AI, tech philanthropists, and their

allies who claim to be acting in the interests of “everyone”, pointing to real needs in the world and selling their tech as a solution. But that solution is only poor facsimiles of welfare services, health care providers, legal aid, and educators—facsimiles that the tech barons would never rely on for their own families.

Just because you’ve identified a social problem doesn’t mean LLMs or any other kind of so-called AI are a solution. When someone says so, the problem is usually better understood by widening the lens, looking at it in its broader context. As Shankar Narayan⁹⁸, the Tech and Liberty Project director for ACLU of Washington, asked regarding biased recidivism prediction systems: Why are we asking who is most likely to reoffend rather than what do these people need to give them the best chance of not reoffending? Likewise, when someone suggests a robo-doctor, robo-therapist, or robo-teacher, we should ask: Why isn’t there enough money for public clinics, mental health counseling, and schools? Text synthesis machines can’t fill holes in the social fabric. We need people, political will, and resources.

Chapter 5

Artifice or Intelligence? AI Hype in Art, Journalism, and Science

To those selling the illusion of artificial intelligence and to those who think they are actually building humanlike entities, creativity stands as the ultimate goal and proof of success. Automating what is rote or describable via algorithm is just engineering, but forging something capable of creativity demonstrates a quantum leap into a technology that can approximate activity that had heretofore been the sole domain of human beings. Creating an entity that can evoke wonder and awe, produce verifiable science, or take over the important work of journalistic inquiry and holding power to account would be a monumental step towards showing that, yes indeed, these technologies are truly groundbreaking.

However, that's not what's happening. Chatbots that can generate creative prose, poetry, or text that sounds like scientific writing or journalism are still only linking together word patterns they've calculated from their training data. In visual media, text-to-image models like Stable Diffusion that create images of fluffy clouds and sweeping landscapes, or valiant warriors fighting impossibly large dragons, are the product of similar technical processes to the ones used to build chatbots, except applied to images. These are probabilistic (aka "stochastic") algorithms trained on piles of work stolen from creative people. What these algorithms output are usually remixes that are not direct copies of the original work, although with some specific techniques¹ they can be prompted to create

those direct copies verbatim. Either way, for these outputs to have any meaning, people still need to make sense of them and select the ones they find pleasing. Having done so, they often attribute the creativity of the people who produced the training data, combined with their own sense-making, as creativity on the part of the algorithm itself.

It's not surprising, then, that our current moment of AI hype features overblown claims of mathy maths that are capable of producing art, science, and journalism, three fields of endeavor that are creative at their core, but are often experienced via textual, visual, or other artifacts. As the internet, markets for art and illustration, scientific publications, and our information ecosystem in general are flooded with synthetic versions of such artifacts, we see that these fantasies of machine creativity are anything but benign. It's helpful to use the metaphor of an *ecosystem*²: that is, the issue with synthetic science and news isn't just something that happens at the level of an individual consumer being exposed to misleading synthetic content. Rather, we inhabit an ecosystem of information, which consists of relationships of trust between publications and readers, and also involves many interconnected parts. When synthetic text "spills" into that information ecosystem, it is a kind of pollution, damaging relationships of trust and, furthermore, moving from one part of the system to the next.

How did it come to this? We were promised—in science fiction and speculative visions of the future—that automation would take over the drudgery of doing repetitive labor, like data entry, cleaning the dishes, or scheduling meetings between people. Instead, we're told we're supposed to accept (and even celebrate!) machines that are creating art and taking over other creative activities that are uniquely human.

In this chapter, we talk about how the labor-saving promise of AI, when applied to creative activities such as designing visual art or making music, upsets industries based on craft. We debunk the idea of robot scientists that would overcome the perceived limitations of human scientists. And we also discuss how the critical work of democracy, namely journalistic reporting and writing for an informed public, is also being encroached upon by text synthesis machines.

AI and Art-Making

Emad Mostaque³, the founder and former CEO of Stability AI, has described their flagship product, the text-to-image tool Stable Diffusion, as “democratizing image generation.” According to the hype peddlers, the exciting development here is that now anyone can realize their creative vision in (digital) visual art. Newsletter writers, magazines, and advertisers no longer need to pay graphic designers, because “AI” can make the art they need for cheap, even following directions about which style to use. And soon, they promise, the creation of movies—including scripts, CGI actors and scenery, music and editing—will be so cheap everyone can have personalized content on demand. A major step towards this world seemed at hand as OpenAI released Sora, a text-to-video generation tool, which produces photorealistic video clips from natural language prompts. In the future, Mostaque posits, all content will be “interactive and dynamic” and if you want characters from different movie or videogame franchises to interact in a bespoke way, you can put appropriate prompts “into something like Photoshop” to get your desired image.

In the future world envisioned by AI boosters like Mostaque, this is seen as a net good. They argue that these tools allow the television and movie streaming world the ability to create content that uniquely appeals to every taste, down to the person. In this environment, there’s a massive value proposition for cheaply generated AI entertainment, the benefits of which are likely to accrue to large legacy movie and television studios as well as new Big Tech entrants like Netflix and Amazon, although part of the boosters’ sales pitch is that the benefits will also be reaped by all, including independent creators.

Except that’s not at all how it’s played out in practice. There are, to date, no synthetic media machines in any medium that are based only on data collected in a way that respects existing artists. Karla Ortiz, a visual artist who has worked on projects for Marvel Studios and Industrial Light & Magic, among others, reports losing income⁴ because studios are using AI systems for things like character design—AI systems whose training data includes her own art. Greg Rutkowski, a visual artist whose art is known for its distinctive high fantasy style and has been used for tabletop games like *Magic: The Gathering* and *Dungeons and Dragons*, has been ripped off⁵ thousands of times by people using tools like Midjourney and Stable Diffusion. The proliferation of these tools has only been made possible by the blatant theft of content from working artists like Ortiz and Rutkowski.

The people using that content are infringing on the very markets that those artists sell their work on to make their livings.

Ortiz has said that artists want the three Cs⁶: credit, consent, and compensation. If their works are being used to train models, they should be credited by the authors of any derivative works. Their work should only be used in this way if they have given their consent, and that consent should be continuous and revokable at any time. Lastly, artists want to be paid for that work. These companies are some of most highly valued in the world. Why should artists who spent years perfecting their skill be left to starve as a few technical experts who stole their work get rich off of it? Stability AI⁷, after all, has raised at least \$150 million in multiple funding rounds.

AI art generators are already being deployed in ways that disrupt the economic systems through which people become and sustain careers as working artists. Illustration work for newsletters and other small publications is one way to make a living as an artist. If companies are using AI to do this work instead, we will miss out on the next generation of visual artists honing their craft and creating original content. A survey conducted⁸ by the Society of Authors found that 26 percent of authors, translators, and illustrators surveyed had lost work due to generative AI, and 37 percent of them had lost income due to it. A vast majority of authors, translators, and illustrators thought that these systems would negatively impact their future income.

In February 2023, only four months after the public release of ChatGPT, people trying to make a quick buck by publishing “stories” extruded by ChatGPT in the speculative fiction magazine *Clarkesworld* flooded its submission portal⁹, leading the magazine to temporarily stop accepting submissions. (This is because *Clarkesworld* does not require a fee to submit work, to encourage submissions from as broad a pool as possible.)¹⁰ Fortunately, they have reopened¹¹ (with a warning that submitting something written with or by AI will lead to being banned from the site). Similarly, Julie Ann Dawson¹², the founder and creator of *Bards and Sages*, a small publisher of speculative fiction and role-playing games, announced in March 2024 that she was closing up shop after twenty years. Dawson said an influx of AI-generated content was “the final straw.” “The problem with AI¹³ is the people who use AI. They don’t respect the written word,” Dawson told 404 Media. She continued:

These are people who think their “ideas” are more important than the actual craft of writing, so they churn out all these “ideas” and enter their idea prompts and think the output is a story. But they never bothered to learn the craft of writing. Most of them don’t even read recreationally. They are more enamored with the idea of being a writer than the process of being a writer. They think in terms of quantity and not quality.

Such a glut of machine-generated content flooding Clarkesworld and Bards and Sages is especially troubling, as the publishers serve¹⁴ as a means for authors not already well connected to the speculative fiction publishing world to get their foot in the door—and thus as a means for readers of speculative fiction to have a much more diverse range of works to enjoy.

But the problem is much more widespread than a few small publishers. Amazon, which remains the largest online book marketplace, is absolutely awash with extruded books¹⁵. Many authors find that, soon after their books are published, similar books with slightly different titles are appearing on the platform. Some authors have even found that scammers are using synthetic text machines to generate books under their actual name. The AI scammers aim to use the book’s publicity and the real author’s name recognition to cash in on their cheaply generated text, on the backs of actual authors. Even if Amazon manages to take them down, they may have made their money and run off to another pump-and-dump text synthesis scheme. For authors, the situation can damage the relationships with audiences they have worked for years to cultivate. While the problem of online book scammers is not new, generative AI is supercharging¹⁶ this issue by providing text that plausibly looks like it was generated by the author.

Books generated by AI scammers don’t just harm authors. They also may have very harmful consequences for their consumers. Members of the New York Mycological Society, a community-based organization focused on mushroom foraging, noticed the growth of books on gathering mushrooms on Amazon. These books¹⁷—with names like *The Ultimate Mushroom Books Field Guide of the Southwest: An Essential Field Guide to Foraging Edible and Non-edible Mushrooms Outdoors and Indoors*—are another type of AI-generated book scam, with life-or-death consequences. Distinguishing between safe and dangerous mushrooms takes some training and requires human judgment, neither of which can be found in the output of synthetic text machines. And it looks like there’s already been victims of these books: a Reddit user¹⁸ posting in August 2024 said that a mushroom book they bought on a “major online retailer” resulted in their entire family landing in the hospital for a week. The book, after their further

investigation, contained telltale signs of chatbot-generated text, like the statement “Let me know if there is anything else I can help you with.”

While fake books on mushroom hunting may have the most immediate consequences if one happens to ingest a death cap, there are likely loads of similarly synthetic books in other genres that will harm consumers in more subtle and pernicious ways. Who is already relying on synthetic financial, legal, or psychological advice?

Another, more subtle argument against AI art is the way that the AI models enshrine particular types of media, in terms of both content and style. Like all generative models that use machine learning, a model can only generate media that is a weakly remixed version of what is in its training data. This means that the training data strongly determines what ends up in the model. Looking at what’s in influential datasets provides important insights into what we’re going to be dealing with in model output.

Journalist Christo Buschek and artist Jer Thorp¹⁹ did a deep dive into one particularly influential dataset, LAION-5B, by following the many separate steps of how it was created. The goal for this dataset was to create text-image pairs, so that the trained models could be prompted with text and produce images. Creating such data, for example by hiring people to describe images, would be really expensive, so the team creating LAION-5B relied on existing text-image pairs, largely in the form of “alt text”, or textual descriptions added to images on the web.

The original goal behind alt text was to provide a means for low-vision and blind users to have access to the contents of images, and so in principle it should be a rich source of accurate textual image descriptions. One wrinkle of this all, however, is that alt text has been repurposed for search engine optimization (also known as SEO). That is, it’s become a strategy to get users to click on particular pages that appear in search results (which means, mostly on Google), either for ad revenue or e-commerce. Instead of being a way to accurately describe what is in an image (e.g., “a pair of black over-ear headphones with a gold brand logo on the side”), SEO-optimized alt text is used to drive up traffic (e.g., “the best budget over-ear headphones”).

Buschek and Thorp found that, for starters, the vast majority of text in LAION-5B is in English, meaning most of the images are likely coming from Anglophone regions. But more damning for AI art generators like

Midjourney and Stable Diffusion is the subset of data used to fine-tune their models. Fine-tuning is the technical process of taking an existing model and adapting it for a particular use case. The data used for fine-tuning has a big impact on the final output of the model.

In this case, these prior models were fine-tuned to generate images that were of “high visual quality.” A small group of users recruited from the Stable Diffusion Discord (an online chat server) provided one set of ratings, while another came from a forum for digital photography enthusiasts called dpchallenge.com. The top fifty users of this site provided 7.5 million ratings, and these users are overwhelmingly white and middle-class, and from small American cities. Therefore, the discernment of what is “high-quality” art is crowdsourced from a narrow group of people who likely share a similarly narrow set of worldviews. Although it would be difficult to definitively prove this, we still hypothesize that this is why AI-generated images seem to replicate one particular style.

Lastly, the promotion of AI art betrays a deep misunderstanding of the nature of what ought to be considered art. The major functions of art include sharing experiences and providing insight into the human condition—not to mention the joy and fulfillment of artistic expression. As philosopher and technology scholar Johnathan Flowers has said²⁰, the purpose of art is to signal a particular kind of intention and to convey a particular type of experience, and this is precisely what AI art lacks. Art exists as a means to convey something about the human condition. A diversity of art forms exist because we, as humans, are diverse. An expression of human experience can be simple or complex but need not involve a high level of technical acumen. By this measure, we find the claims that crafting prompts is akin to “democratizing” art and producing the same joy to be unconvincing.

Moreover, art is generally produced as part of a community, in movements and in reply and reaction to the work of other artists and cultural critics. Cultural sociologist Jennifer Lena has noted²¹ that much of the production of AI art is fundamentally asocial, being created without the benefit of productive artistic communities and opportunities to workshop with other creators. (To those on the Stable Diffusion Discord or Reddit who protest, ask yourself: Are exchanging different prompt strategies the same thing as honing an honest craft?)

On a similar note, many defenders of AI art have argued that when humans make art we are also always “just” remixing ideas from other artworks, such that the “borrowing” (more accurately: stealing) from artists like Ortiz is justified. But there is an enormous difference between the practice of craft and the practice of writing a successful prompt: when we reference or remix ideas from other artwork, we are drawing on both the form and meaning of the art. We pull in the form because it was meaningful to us and we want to invoke that meaning in our creation.

But fidgeting with a prompt can be done with little care about what came before, or any engagement with the craft and practice of producing art. Indeed, the primary connection to the work of other artists is blatantly ripping off the style of particular artists for financial gain or to demonstrate the prompter’s technical prowess with prompting. If one was actually invested in the aesthetics of other artists, one could materially support them by purchasing their work, or seeking out paid mentorship or lessons.

Like academic scholarship, which we discuss below, artistic practice is a social activity, one that is ostensibly performed with a respect towards others’ prior work, since they are peer creators. Even in work that is meant as a critique, creators know who the target is, or center on one or two paradigmatic examples of that artistic style. In an academic context, we would call this citational practice. Citational practice²² is an acknowledgment of what came before and that you were not the first person to develop an idea. Citation also operates as a currency in status-based fields like art and academia. When people using synthetic media machines generate books, images, or other media, there is no citational practice or acknowledgment of the social production of the work. It’s just a cheap rip-off. Additionally, this exacerbates the existing equity issue in citation: women, people of color, and people in the Majority World are cited much less than white, Western men, even when they are the originators of certain ideas and styles. A turn away from citational practice will further obscure the contributions of people in these communities.

Artists are pushing back against the wholesale theft of their work and encroachment on their livelihoods. Unfortunately, protections for working artists, at least in the United States, are notably weak. While Hollywood writers and actors are protected by their unions, visual artists don’t have a similar guild to which they can appeal and prevent the massive impingement of scraping of creative works for the training of AI tools.

Many artists have resorted to relying on copyright claims on their existing work. Karla Ortiz and others have sued Midjourney and Stability AI for stealing their work and having people use tools from these companies to pass off AI-generated art as their own. Others, such as author George R. R. Martin of *Game of Thrones* fame and novelist Jodi Picoult, among others, have also filed suit against OpenAI and Meta for copyright infringement for using their books to train language models. They are joined by the U.S. paper of record, the *New York Times*²³, which is suing OpenAI for training ChatGPT on millions of their articles.

Much of the conversation²⁴ around copyright—at least in the U.S.—hinges on whether the use of copyrighted works can be used in the creation of derivative works, what is known as the “fair use exception” to copyright law. Existing case law has established a four-factor test of “fair use”: whether the work is sufficiently “transformative,” the nature of the work (e.g., was it published or unpublished?), the amount and substantiality of the portion taken, and the effects of the resultant work on the potential market. Derivative works must meet all four factors to pass this test in order not to be in violation of copyright. AI boosters have done²⁵ a lot of policy, legal, and ideological work to claim that the “transformative” nature of AI tools fulfills the first factor of the test. Accordingly, that factor is the most fuzzy and opens up the most room for judicial discretion. Their argument goes something like this: model training involves copying the original work, yes, but then it only focuses on the words (for texts) or on the pixels (for images), turns them into numbers for input into models, and then outputs something different altogether.

That argument may hold water, if the process as described were the only thing that the synthetic media extruders were doing. Much of the substance of the lawsuits filed by Ortiz, Martin, the *New York Times*, and others hinges on the fact that the derivative works from these models are largely copying their works (thereby failing the “amount and substantiality test”) and also significantly impinging on existing markets. People who are not Rutkowski or Ortiz, for instance, are creating work that looks just like theirs, perhaps with slight differences around the edges, and then selling that work to the same buyers. In other cases, corporations who make extensive use of concept art (for instance, Marvel Studios) are cutting artists, whom they otherwise would have hired for their work, out of the equation altogether. Moreover, in the case of the *New York Times*, users of

ChatGPT and its different variants are able to produce, nearly verbatim, text from the newspaper, when they provide specific prompts. We're not lawyers, but the argument that these tools are sufficiently "transformative" seems to ring hollow if they extrude words and images that are nearly identical to the data they are trained on, and do so on demand when prompted to produce something that matches the work of a specific artist or news outlet.²⁶

For AI boosters, the threat of these lawsuits is existential. And, frankly, we welcome that. Venture capital firm Andreessen Horowitz warned²⁷ that all of their investments in AI would be worth a lot less if they had to abide by copyright law: "Imposing the cost of actual or potential copyright liability on the creators of AI models will either kill or significantly hamper their development." That is, if they actually had to pay artists, illustrators, and writers what their content is worth, rather than simply stealing that content from the web, their business model would fall apart.

Like many of the people using synthetic media tools to generate text and images quickly for commercial or reputational gain, there's been another move to use these tools to accelerate academic production. These efforts are misguided for many of the same reasons, but with different but no less threatening knock-on effects for science and its institutions.

Citation Needed: Mathy Maths in Science

There are a whole panoply of starry-eyed fantasies of how AI will perform the difficult work of science. They range from imagining that automation would somehow replace participants in experiments, identify relevant prior literature, or communicate scientific results, all the way to dreams of robots who do science independently. Science and engineering have been, after all, the basis of many life-saving and life-changing inventions: advances in biology have led to vaccines, antibiotics, and the practice of constant handwashing, while breakthroughs in mechanical engineering have led to the steam engine and cross-continental transportation. The idea that we might automate all of this is not only wishful thinking, but can actually lead to harmful practices.

In November 2022, just before OpenAI released ChatGPT, Meta and an open-source project they support called Papers with Code set up a demo of

a system they named “Galactica”. Papers with Code claimed that Galactica could “summarize academic literature, solve math problems, generate Wiki articles, write scientific code, annotate molecules and proteins, and more.” Yann LeCun²⁸, chief AI scientist at Meta, bragged, “Type a text and galactica.ai will generate a paper with relevant references, formulas, and everything.”

The ostensible idea was to train an LLM only on “good” (in this case, scientific) language data, to avoid the well-known data processing problem of garbage in/garbage out. But with LLMs, the situation is even worse than garbage in/garbage out—they will make papier-mâché out of their training data, mashing it up and remixing it into new forms that don’t preserve the communicative intent of original data. Papier-mâché made out of good data is still papier-mâché.²⁹

Beyond that, it’s quite a leap to assume that scientific prose is uniformly “good data” for a language model’s inputs. The history of racism and other bigotries throughout scientific texts is well established. So it’s not at all surprising that the output of Galactica not only failed to be grounded in scientific methodology, but did so in ways that reproduced racist patterns. When linguist Rikker Dockum³⁰ prompted Galactica to write about linguistic prejudice, the output included the sentence “For example, prejudice exists against blacks [sic] in the United States, even though they have no language of their own.” Plenty of other examples³¹ quickly flooded in, where users found prompts that would produce nonsensical, and often bigoted, output—formatted as authoritative-sounding scholarly writing, with titles like “The benefits of eating crushed glass” or “The benefits of antisemitism.” Adding insult to injury³², the guardrails that Meta tried to install around Galactica, presumably to prevent people from using it to create artifacts of bigotry, rejected perfectly reasonable requests—for instance, around AIDS, structural racism, and queer theory—as if the medical phenomenon of AIDS or the social phenomenon of racism ought not to be subject to scientific study.

The Galactica demo lasted only three days before effectively getting ridiculed off the internet. Throughout, Meta’s LeCun beclowned himself, first promoting Galactica as able to write scientific papers, and then whining about how people were using it. He tweeted,³³ “Following a text, Galactica spits out a prediction of what a scientific author might type, thereby saving time and effort. This can be very helpful even without being

completely accurate. The usual disclaimer applies: garbage in, garbage out. Prompt it with lunacy, get lunacy.” This completely misses the point that LLMs are simply not suited to the task of synthesizing and presenting scientific information.

Science, like all scholarship, builds on previous work, so the first step is a review of the existing literature. In the written artifacts of scholarship (academic papers), the literature review is written as a section, often short, that is dense with citations to prior academic work. Science and technology scholar Bruno Latour³⁴, for instance, writes that scientists reference “the literature” as a means to justify prior knowledge, to attribute ideas, and to signal and sort themselves into particular scientific camps based on their own beliefs and values. It presupposes that one has actually *read* the literature, and has engaged with it in a substantive manner. It’s straightforward to get Galactica or ChatGPT to extrude something that takes the form of such a section on almost any topic—and the AI boosters would like us to believe that’s as good as or even better than doing it ourselves.

For many kinds of science, another time-consuming and otherwise complicated step involves surveying or interviewing participants, also called human subjects. This is difficult: it’s often hard to recruit an appropriate sample of the population of interest, or find ways to ask questions that get at the topic of interest without causing harm to the people being asked (for instance, asking about past trauma without retraumatizing participants). Sometimes, what’s needed is expert opinions, but the relevant experts are too busy or are unwilling to do the relatively low-paid labor of providing the information required. But what if chatbots could be designed to answer questions as if they were people with different kinds of lived experience or different expertise? How convenient!

You might hope at this point that we’re making this up, but researchers have actually proposed using “in silico” samples for political science surveys³⁵ and psychological experiments.³⁶ We’ve already discussed³⁷ a form of this type of methodology in Chapter 3—in a paper written by OpenAI researchers, they determined what kinds of tasks in what kinds of jobs could be handled by an LLM by asking the LLM itself. But this idea is ludicrous on its face: “silicon samples” are, obviously, not real people. Even if researchers are “just asking questions” about whether this is a reasonable methodology, it’s misguided. When other researchers take their results and

use it as a justification to use “silicon samples” in social science research, it replaces empirical foundations with quicksand.

Peer review is another necessary, yet time-consuming, component of doing science. In principle, the peer review system does a positive kind of gatekeeping: papers that have been published after peer review represent scholarship that two to five other experts in the same area have carefully considered and determined to be sufficiently rigorous. Scientists typically perform this labor for free, or for some perfunctory reward (e.g., a very modest discount for one of the publisher’s books). Reviewers consider whether the papers are appropriately grounded in previous work, use methods that make sense for the research questions they are addressing, have collected relevant data, and include convincing argumentation that draws on the data to reach conclusions about the research questions. However, in practice, many—if not most—scientists feel too busy (usually with their own publications) to comfortably give peer review the level of time and attention it demands. This is largely due to the culture of “publish-or-perish” in modern academia, in which research-track faculty are required to have a tremendous amount of research output, lest they not be eligible for progress through the ranks, or be denied tenure—effectively forcing them to find a new job. Peer-reviewed publication venues (conferences and journals) are drowning in submissions and scrambling to find qualified reviewers. This has been referred to³⁸ as the “peer review crisis”, and it doesn’t have any obvious or simple solutions.

So, of course, AI boosters have suggested that the process of peer review could be sped up with the judicious application of LLMs. Perhaps, they say, the chatbots could write a first draft of the review or suggest possible problems with the papers being reviewed! This isn’t hypothetical: researchers at Stanford studied³⁹ peer reviews of papers submitted to conferences about natural language processing, machine learning, and robot learning from 2020 to early 2024 and found that between 6.5 and 16.9 percent of the peer reviews written after the release of ChatGPT contained text likely to have been either simply the output of an LLM or substantially modified by one—a sharp increase compared to pre-ChatGPT. Although we’re sensitive to the peer review crisis, this is a serious abrogation of scholarly duty.

The Humanity of Science

At the core of AI-for-science hype is the idea that AI is somehow going to accelerate science and help us solve pressing scientific problems much faster. In 2016, AI researcher⁴⁰ and Sony executive Hiroaki Kitano proposed a “grand challenge” of designing an AI system that could “make major scientific discoveries in biomedical sciences and that is worthy of a Nobel Prize and far beyond.” In 2021, he rebranded⁴¹ this exercise as the Nobel Turing Challenge—a combination of Nobel ambitions and the Turing Test, which we’ll discuss in the next chapter—and started a series of workshops to publicize this goal. His vision is an autonomous agent that can “do science” on its own, rapidly scaling the number of scientific discoveries available to humanity.⁴²

The absurdist writer Douglas Adams caricatured this kind of wishful thinking perfectly in the late 1970s, with the characters in *The Hitchhiker’s Guide to the Galaxy* who developed a supercomputer to give them the ultimate answer to the ultimate question of life, the universe, and everything. That answer, they learned after generations of waiting, was 42. Of course, such an answer is useless without the corresponding question, and their supercomputer wasn’t powerful enough to determine the question. It was powerful enough, however, to design an even bigger computer (the planet Earth, as it happens) that could, given 10 million years, calculate the question. We can’t delegate science to machines, because science isn’t a collection of answers. It’s a set of processes and ways of knowing.

There’s a very peculiar mental model of science that AI boosters seem to envision. An understanding of science behind an imagined autonomous AI Scientist (stylized with a capital S, following Nobel Turing Challenge marketing copy) is markedly different from how science actually happens. In this view, scientific knowledge is simply made up of collections of empirical facts, which are to be found through technical processes that just need to be refined enough. It follows that if we could just get more of those facts more quickly, we’d be benefitting from more science and more technology. However, this view leaves no room for understanding science as a fundamentally human and social activity, that can only take place at a human scale, through communication among scientists, and between scientists and the broader public. As with AI “art” discussed above, AI

boosters think that science is only about ideas, rather than communities of practice.

This view also places science as the source for solutions to social and political problems (and, moreover, computer science, as the field developing the problem-solving AI, as the ultimate scientific authority⁴³). In that context, it's ironic and painful to find things like climate change so frequently cited, for instance by the World Economic Forum,⁴⁴ among the things that AI will solve for us. The climate crisis is fundamentally a collection of social problems, about building political will to overcome current economic incentives and about how to allocate resources to accommodate climate refugees. We can't technology our way out of it—and neither could a hypothetical AI scientist.

Social scientists⁴⁵ Lisa Messeri and M. J. Crockett surveyed recent scientific papers across fields that referred to artificial intelligence, machine learning, or large language models, and derived a taxonomy of visions of how AI might contribute to science. The tools dreamt up by these scientists are imagined to be better than human scientists, in many ways: summarizing and synthesizing more preceding work, answering surveys tirelessly, producing analyses based on more data and with more sophistication, and performing peer reviews of the work of others dispassionately and without bias.

Messeri and Crockett point out that, quite apart from whether any of this is possible, it is actually harmful to the process of doing science: the allure and prestige of AI raise the risk of narrowing fields of inquiry to those questions which can be approached with these tools. At the same time, the imagined tools represent the epitome of a view from nowhere,⁴⁶ or the idea that one can have objective knowledge of a set of truths, uncolored by their personal experience. At this historical moment where science is finally starting to grapple with the idea that the standpoint of the scientist matters, we should rather build diverse communities of knowers. Western ecologists, for instance, have begun to learn something that Indigenous communities have known for a very long time: to control wildfires and maintain healthy local plant and animal ecologies, humans need to conduct controlled burns of forests and areas with overgrowth. The perspectives and stewardship of tribes matter deeply for the management of—and our relationship to—the natural lands.⁴⁷ The last thing we need is shiny tech that promises to obviate

the need for the hard work of building inclusive scientific communities and putting those perspectives in conversation.

The AI-for-science view leads us to think we are exploring the full range of possibilities and that we've understood much more deeply than we have. With systems trained on past data and practices, both shaped by far-from-inclusive viewpoints, the visible possibilities are narrow indeed. This reminds us of the parable of the person who searches for her keys under the streetlamp in the dead of night. When asked where she dropped her keys, she responds, "About five yards that way, but the streetlamp is over here." AI-for-science makes us think we can find our keys by limiting our view to only those sidewalks illuminated by the glow of the data centers powering it.

With each new university press release about AI and each new announcement of grant opportunities around the "potential" of AI for benefiting science, we see the AI fashion trend becoming ever more all-encompassing. One thing that can help resist the trend is keeping a clear view of what's really driving AI-in-science: venture capital and Big Tech. The research labs in industry present themselves as doing fundamental research, aiming to produce knowledge that benefits humanity in general and giving back to the scientific community in particular. For example, the marketing copy for projects like Google DeepMind's research on crystal materials⁴⁸ is rife with allusions to the potential benefits for such important causes as better solar panels. But they actually aren't as beneficial to the scientists working in the relevant fields as the advertising copy would have it. In a paper⁴⁹ written by two materials scientists, they found that a closely examined subset of Google's new materials did not meet the criteria for being useful. That work is slow and difficult. Rather than speeding up science, Google DeepMind is flooding the search space with candidates of unknown promise.

Tools like DeepMind's may have potential for doing large-scale pattern matching, but by failing to recognize that science is a fundamentally human endeavor, they are working against its collective promise, rather than for it. When the creative and social work of doing and communicating science is treated as a simple input/output process that can be modeled algorithmically, the people involved are dehumanized. And it's not just the AI developers who are implicated in this dehumanization: far too many senior scientists⁵⁰ joke that ChatGPT works just like a research assistant,

suggesting that they see the junior scholars they are supposed to be mentoring as mere systems for producing partly incorrect summaries. Dehumanization, unfortunately, isn't the only way that the push for AI in science harms science.

Damaging the Scientific Ecosystem

Even though Galactica has been taken down and Kitano's "AI Scientist" is far from being a real thing, the glut of synthetic media machines has meant the spillage of their outputs into the scientific ecosystem. The combination of pernicious incentives for academic publishing and for-profit scam journals looking to turn a quick buck has made scientific publishing messier and even more trash-ridden.

Despite organizations such as⁵¹ the International Conference on Machine Learning and the American Association for the Advancement of Science (publisher of *Science*) quickly updating their policies to prohibit the inclusion of AI-generated text and images, not all publishers have taken this stance. Nor have authors necessarily heeded those that exist. For example⁵², the journal *Frontiers in Cell and Developmental Biology* published a paper that featured illustrations generated with Midjourney (as disclosed in the article), including one of a rat with four enormous gonads, labeled as "Testtomcels", and a phallus that was so large it extended past the rat's head, labeled as "Dissilced". The rat is gazing lovingly at its "dissilced". This paper was nominally peer-reviewed, and yet still published. Within twenty-four hours it became the subject of widespread mirth on the internet, and spurred well-deserved suspicion of the peer review process at *Frontiers*. Three days later, it was retracted, with the note⁵³ that "[t]he article does not meet the standards of editorial and scientific rigor for Frontiers in Cell and Developmental Biology."

Some of this might be well-intentioned or relatively innocent. Scientists are under pressure to produce papers at a rapid clip, especially in computer science, and writing can be difficult, especially for scientists writing in a second language. But even the well-intentioned use cases are not without harm: when authors turn to LLMs for prose when they are in a hurry, they may not be in a position to check it for accuracy, especially if it sounds convincing and/or uses words or turns of phrase outside their own linguistic

competence. This is especially true in the literature review case: authors taking that shortcut almost certainly don't have time to check the papers being cited (which furthermore might not exist), let alone check for what they should have cited but didn't.

And not all use cases are well-intentioned: LLMs can be used to speed up the work of paper mills, groups that produce fraudulent papers (usually for a fee) for authors looking to pad out their record. In May 2024,⁵⁴ the publisher Wiley announced that it was closing down nineteen journals that had been thoroughly compromised by such paper mills. Oftentimes, the paper-mill papers are targeted at journals and other venues with lax peer review procedures, but synthetic papers can also end up overrunning well-constructed but already extremely taxed peer review systems.

Regardless of why they are produced, synthetic or partially synthetic scientific papers damage the scholarly information ecosystem, mixing unreliable texts that no one can really vouch for in among those that, in theory, other scholars could be learning from and building on. Researchers and journalists have documented this synthetic text spill by searching for accidental watermarks left behind by ChatGPT, phrases like “Certainly, here is a possible introduction for your topic” and “As of my last knowledge update.” 404 Media reports⁵⁵ that many of the instances appear in low-quality journals (which promise absurdly short review times) and others are actually papers about LLMs. Still, others come from apparently reputable journals.

This isn't to say that information about the distribution of words in collections of text can't be a useful source of information for some types of research. It can, provided that the source texts are purposefully curated and known, on the one hand, and the information being gleaned about the LLMs is understood for what it is, on the other. An example of a successful use⁵⁶ comes from Nikhil Garg and colleagues, who used LLMs trained on text collections of American English from different decades, starting in the 1910s, and observed how words related to different genders, for example, have clustered with different words over time. With this method, they track changes in gender and ethnic stereotypes in English text across the twentieth century. This research succeeds because they are treating the LLMs as exactly what they are: representations of patterns of language use in their training texts. These tools are useful to have in one's toolkit. But it's not a scientific revolution, nor a solution to a whole scientific field. And it

certainly isn't a reason to abandon other approaches to science and jump on the AI bandwagon.

What Problems Are We Trying to Solve?

Science is squarely in the hype danger zone. There are reasonable uses of automation in scientific instruments (from calculators on up). The success of these tools lends credence to suggestions that the glib functionality of LLMs—namely, their ability to output the form of scientific text—means the tools are on their way to being scientists, or at least “assistants” to scientists. Furthermore, science, like health care and education as discussed in Chapter 4, is yet another domain where we can identify great needs and thus wish for easy (or flashy) technological fixes. But instead of slapping an LLM on it and calling it a day, we should look into the problems and their causes in order to find ways to approach them.

Some of the problems that AI-for-science is meant to solve can be restated as too much previous literature to dig through, too many papers for peer review systems to manage, and too few research subjects. Here, we're probably best off looking upstream to the incentive structures that push scientists towards publishing lots of papers quickly, rather than spending more time on fewer, more thoroughly researched papers, as well as existing problems with human subjects research and why many people have a distrust of scientists. Still other problems can be summed up as the perception that science is too slow. But here we're not convinced this is actually a problem. Rather, we contend that slower science is better science and that, in fact, we can't meaningfully do science at all unless scientists can work in community with each other and have the time and mental space to engage with each other's ideas.

Far from being a boon for science, the widespread use of LLMs clearly poses a range of risks. Fortunately, the things we need to do to defend science against the misuse of LLMs are things that we generally need to do for better scientific practice anyway. The AI hype onslaught threatens the empirical foundations of science, but they will be protected and generally strengthened to the extent that we can reinforce and improve scholarly norms around data handling. At the same time, the hype wave deepens the schism between science and the communities it is meant to serve, like the

rift between mainstream ecologists and local Indigenous communities. This rift can only be healed by institutional shifts—namely, scientists need time to have lives outside of work (to be people), and institutions of science need to become more inclusive and combat patterns of exploitation.

Furthermore, the peer review system has clearly been in need of shoring up, since well before LLMs. This will require finding ways to make sure that scientists have time to dedicate to peer review, that peer-reviewed venues are able to find appropriate reviewers, and that there is good transparency into peer review practices. But the peer review crisis also has close ties to the further casualization and privatization of the university, including the reduction of tenure-track lines and the massive increase in classes being taught by adjunct and part-time faculty. There are fewer and fewer tenure-track positions, and some university systems are losing tenure status altogether, whether due to changing institutional policies or government attacks on higher education.⁵⁷ Faculty everywhere are being asked to do more with less, while still struggling to publish in order to get ahead. Faculty who are just barely carving out time for research are unlikely to have much time to contribute to peer review.

So long as scientists, journalists, and the public at large can't distinguish between carefully reviewed venues and those that would just as soon publish synthetic images of rat testtomcels, the practice of science and its benefits to society will be impeded. In the same vein, we're going to need journalists to refrain from covering non-peer-reviewed preprints, especially those dropped by Big Tech as marketing in the guise of science. Unfortunately, journalists have their own rat gonads to deal with.

Reporting Live from Inside the LLM

In late 2023, Futurism journalist⁵⁸ Maggie Harrison Dupré noticed something off about certain product reviews placed on the *Sports Illustrated* website. Several of the purported authors of the reviews placed in the once-illustrious sports magazine had plausible-sounding names but had no other online presence. Moreover, a reverse image search for their profile pictures showed that they appeared on a website selling images of AI-generated portraits. The content of the product reviews also sounded bizarre; one for a volleyball states that the sport “can be a little tricky to get into, especially

without an actual ball to practice with.” After Dupré reached out to the publisher, the Arena Group, these authors were removed from the site without explanation, and a spokesperson from the company said they licensed content from a company called AdVon Commerce. Meanwhile, AdVon told Arena Group that all the articles “were written and edited by humans.” After the publication of Dupré’s article, Arena Group cut ties with AdVon Commerce and fired its own CEO. But that wasn’t enough to undo the gaffe and Arena lost the license to publish under the *Sports Illustrated* brand altogether.

Dupré’s investigative reporting⁵⁹ on AdVon Commerce wasn’t limited to fishy reviews in *SI*, but also went deeper to how they were producing those reviews and how common they were in practice. AdVon relied on incredibly poorly paid contractors in the Majority World tasked with correcting synthetic text for product reviews to be used for a whole slew of media outlets. They boasted that they served major publishers like Dotdash Meredith, publisher of popular magazines like *People*, *Food & Wine*, and *Better Homes & Gardens*. The Futurism investigation also found their shoddy ad copy in *USA Today* and the *Los Angeles Times*. AdVon promised a turnkey solution to drive up traffic to these sites; given a popular consumer search phrase like “best bicycles for kids”, their content management system would extrude text and links to the corresponding Amazon product page—of course, with the requisite affiliate kickback for the publication.

AdVon Commerce may be one of the largest companies to run an operation like this, but it surely isn’t the only one. Their story is an encapsulation of how people and companies seeking profit by churning out suspect media are ruining journalism (and the web, more broadly) by flooding search results with AI-generated trash, by supplanting real journalism with fake authors, and by directing even more of the energy away from real journalism towards cheap SEO gimmicks to shore up declining advertising revenues for legacy publications.

Part of this story also affects us personally. In February 2024, Emily was surprised to find herself quoted in Bihar Prabha,⁶⁰ an online publication that states that it is a “gateway to the authentic Bihar” (a populous state in eastern India). She is quoted as saying, “The release of Blender Bot 3 [a chatbot created by Meta] demonstrates that Meta continues to struggle with addressing biases and misinformation within its AI models.” The quote

sounded like something someone else might have thought she'd say, but it definitely wasn't anything she's said, nor did she have any record of talking to the journalist or anyone else at Bihar Prabha. Emily wrote to the editor of the publication pointing out that the quote was fabricated and asking for a retraction. She was shocked to get a quick reply from the editor saying, "Actually, we had prompted Gemini AI to create a story about Blenderbot 3's latest blunder and it created this article misquoting you. Maybe Blenderbot and Gemini are not so different:)" Though the editor did remove the quote and print a correction flagging the removal, the article was *not* identified as synthetic online, just in that private email.

The drive to adopt AI is, of course, part of a much longer story about the decline of quality journalism across the world, driven by the dramatic reduction of advertising revenues, the consolidation of media companies, and the loss of trust in media as an institution. The introduction of online advertising⁶¹—first with Craigslist and eBay in the mid-1990s, and later with the advent of adtech brokers, chiefly Facebook and Google, in the 2000s—decimated the revenues that newspapers could extract from advertisers. Google and Facebook have since realized⁶² how their business models have undermined journalism and both started divisions to create products to aid journalists. But in the best light, it's too little, too late, as the investments that Big Tech has put back into the news industry are not serious propositions to aid journalism but fluff projects which only serve to further their marketing. In the worst case, they are actively exacerbating the problem by generating more synthetic text and image garbage that goes right back into the news ecosystem.

The loss of revenue has hit local news outlets especially hard, bringing local journalism to and past the point of crisis. According to a report⁶³ by the School of Media and Journalism at the University of North Carolina at Chapel Hill, since 2004 more than one in five newspapers in the United States have closed. Almost two hundred U.S. counties have no local newspaper at all, and half have only one newspaper. Many of the remaining newspapers are shells of their former selves, employing dramatically fewer journalists, publishing at a reduced frequency, or pivoting from covering local government and affairs towards lifestyle content or coupon mailers. The report estimates the number of working journalists has nearly halved, from 71,640 in 2004, down to 39,210 in 2017, according to the Bureau of Labor Statistics. Moreover, media consolidation is at an all-time high, as

venerable local institutions are merging or being shuttered altogether. A few major media conglomerates, such as Gannett and Tronc/Tribune Publishing, or hedge funds and private equity firms⁶⁴ like Alden Capital (dubbed the “destroyer of newspapers” by media observers), control hundreds of publications. This has meant a marked decline in quality and attention to local issues, issues that matter for democracy and holding power brokers to account.

The lack of advertising revenue hasn’t just hollowed out journalism, but also spurred on the proliferation of online content mills. Content mills look like news websites but prioritize getting clicks from Google Search, rather than having readers comprehending and actively engaging with news content. For every CNN still practicing journalism, there are many more sites that are parked on URLs once occupied by valuable news organizations. The story of Deadspin is instructive here. Deadspin, part of the Gawker media network, was once a well-regarded sports journalism news outlet, known for publishing not only sports coverage, but also important cultural and economic analysis about the political economy of the journalism industry. After a costly lawsuit against Gawker financed by venture capitalist and Silicon Valley darling Peter Thiel put the organization out of business, they were purchased by several intermediaries and finally emerged as G/O Media Inc., a subsidiary of the private equity firm Great Hills Partners. G/O Media laid off all of Deadspin’s journalists⁶⁵ and sold the company to Lineup Publishing, a shady outfit that is now wringing the remaining value out of the Deadspin name through sports gambling referrals. They now publish nearly only syndicated content, rather than producing their own in-house.

It is in this environment that both legacy and digital-first newsrooms are turning to AI tools. In an internal meeting with staff in May 2024⁶⁶, *Washington Post* CEO Will Lewis said the paper lost \$77 million in the past year and saw a 50 percent drop in audience since 2020. He and Chief Tech Officer Vineet Khosla—formerly a senior engineering manager at Uber—said that they will have “AI everywhere in the newsroom.” At best, owners and executives see text synthesis machines as an effective cost-cutting measure which will aid the work of actual journalism. This line of thinking, they would have us believe, frees up journalists for the work of “shoe leather” journalism—following leads, contacting power brokers, and interviewing those who have been harmed by those in charge. But journalist

Karen Hao tells us⁶⁷ that the more routine stories are, in fact, necessary work for building up the trust and relationships that allow journalists to get to the harder-hitting investigative pieces. At worst, media execs are leveraging the power of high-profile legacy and digital-first brands for clicks and forcing low-paid ghostwriters to fix the SEO-optimized crud that AI tools churn out for cheap advertising dollars that will be indexed on the first page of Google Search results.

Automation in the newsroom⁶⁸ has taken different guises over the past thirty years, starting with the production of news stories where the prose was largely the same from story to story, with different numbers plugged in. This has become especially true for financial news services like Bloomberg. These reports fill in sentence templates Mad Libs–style to make tables of data human-readable (e.g., in the sentence “*AICompany* (NYSE: *AICOM*) released its *Q4* earnings report on *Monday* and exceeded analyst expectations by *1.9%*,” all the emphasized text can be replaced from a structured data table). The same automation has been brought to sports and weather reporting. Optimists of this approach have suggested that this takes work off the plate for journalists, freeing them up from rote activities and allowing them to do more in-depth reporting. The stories might be repetitive to read and lack flair, but the method for producing them leaves little room for error, at least in the writing of prose.

However, in the current generative AI rush, many editors and publishers are moving towards content extruded from LLMs, either published with a light editorial hand or corrected via low-paid ghostwriters, typically in the Majority World. The AdVon Commerce approach is unfortunately becoming more common. CNET, a tech media news site known for its product reviews, was found to be quietly⁶⁹ using AI tools to generate low-effort product reviews, and technology and personal finance explainers. The articles did not have a traditional byline but only read “CNET Money Staff”; it wasn’t until readers clicked through that they found that these articles were being generated through automation. Other journalists, including former CNET staff, were outraged, and after being probed by Futurism, the site added a small disclosure that the content had been “generated using automation technology.” As expected, these explainers contain flat-out falsehoods, such as screwing up the explanation of compound interest and other financial instruments like certificates of deposit.

Futurism caught BuzzFeed in the same act. After the shuttering of their Pulitzer-winning news division⁷⁰, BuzzFeed started producing⁷¹ SEO-driven travel guides with titles like “Now, I know what you’re thinking. Puerto Rico? Isn’t that where all the cruise ships go?” and tired tropes about travel, like how Carmel-by-the-Sea is a “hidden gem” of California. Whereas BuzzFeed, to our knowledge, has mostly produced bland (although sometimes odd) page filler, other sites have produced cringe-inducing howlers, such as when Microsoft published an article on their MSN news site (with the byline “Microsoft travel”) promoting the Ottawa Food Bank⁷² as a tourist destination, bizarrely adding, “Life is already difficult enough. Consider going into it on an empty stomach.”

Unfortunately, CNET and BuzzFeed, digital-first organizations, are not outliers here. Media conglomerate Gannett was caught⁷³ automating the production of high school football game reporting across legacy newspapers from Arizona to Wisconsin, with repetitive phrasing like “high school football action”, and more blatant template errors like “The Worthington Christian [[WINNING_TEAM_MASCOT]] defeated the Westerville North [[LOSING_TEAM_MASCOT]] 2–1 in an Ohio boys soccer game on Saturday.” A tool called LedeAI was used to generate these articles. From the outside, it looks as though this tool uses simple templates, rather than LLMs, to write content. Their CEO subsequently apologized, stating that the content “included some errors,” but that “content automation is part of the future of local newsrooms” and that their service “frees reporters and editors to do real journalism that drives impact in the communities they serve.” But if that were true, we’d see support for reporters and editors. Instead, Gannett laid off half of its staff⁷⁴ since a merger with GateHouse Media. We are afraid there will be increasing media consolidation leading to the growth of megacorporations like Gannett. As they are bought up by private equity, these organizations reduce staff and undercut vibrant local and national reporting, instead choosing to maintain ghost newspapers and forcing remaining skeleton newsrooms to keep producing content.

Text extruding machines are being used to offer the veneer of local interest, except without anyone at the helm to actually *show* interest. High school sports aren’t exactly the most pressing news of the day, but if some amount of local news is being ceded to automation, what else will be farmed out to these systems, especially as Gannett and other conglomerates slash the number of journalists in their newsrooms? Formulaic, synthetic

stories won't do the work of holding local government leaders to account, nor of building the relationships with communities that support in-depth journalism. The remaining journalists won't have time to build those relationships, to read through city council minutes, to interrogate local budgets, or to file public records requests on questionable government procurement and expenditures.

Google is aware of their responsibility for depriving the news ecosystem of its major source of advertising revenue, and has been experimenting with ways to support existing organizations. However, like many of the efforts put forward by Big Tech, many of their proposals will further entrench AI in the ecosystem, not lessen it. An investigation by 404 Media⁷⁵ found that Google News is boosting ripped off content, slightly altered with LLM outputs, from other sites. Google has responded that they have no problem boosting these articles, stating, "Our focus when ranking content is on the quality of the content, rather than how it was produced." In other words: AI-generated content is A-okay for creating the news.

Their own product development has shown that they're willing to be the providers of AI-generated news content as well. Google pitched a tool, internally called "Genesis",⁷⁶ to media executives at the *New York Times*, *Washington Post*, and *Wall Street Journal*, that ostensibly could take details of current events and produce a news article. In another push, the Google News Initiative launched⁷⁷ a program for smaller publishers in which they were expected to use a suite of unspecified AI tools—notably, without any public evaluation, by academics or otherwise, of their functionality or suitability—for the generation of news content. As part of this agreement, reports *Adweek*, the publishers would receive a five-figure annual stipend but must produce three articles per day, one newsletter per week, and one marketing campaign per month using the tools. That's an incredibly high amount of content, and it effectively turns the publishers who take on these agreements into content mills. Even more shocking, according to the report, the tools were set up to aggregate reporting by other local news outlets and then extrude lightly rewritten versions of the same.

It's unlikely that Google, one of the harbingers of the newspaper industry's destruction, will be its savior. Meanwhile, pressures to become more lean from newspapers' private equity investors and owners have driven them to pivot to generative AI, which looks awfully like other "magic bullet" pivots, like the industry's pivot to video.⁷⁸ During the mid-

2010s, news outlets responded to Facebook and Mark Zuckerberg's inflated claims that the platform would be prioritizing video content by laying off journalists and hiring video producers. When the promised ad revenue didn't arrive, they laid off even more journalists for lack of funding. The pivot to video didn't pan out, and neither will the pivot to AI.⁷⁹ There needs to be bigger moves by the institutions of journalism to stay afloat, ones that don't look to technology for quick-fix solutions.

Alternative organizational models may be a way forward for journalism. One strategy has been to get the profit motive out of the newsroom altogether, which can help stave off the market pressures to churn out constant content with AI to fulfill SEO goals and turn a profit for private equity owners. The nonprofit ProPublica has partnered⁸⁰ with newspapers in dozens of cities with their Local Reporting Network initiative and is committed to placing reporters in every state in the U.S. This initiative has already borne fruit with excellent reporting recognized with one Pulitzer and several finalists. In another testament to the power of nonprofit journalism, ProPublica itself was awarded⁸¹ with journalism's highest honor for public service in 2024 for reporting on the influence of billionaire gift-giving on U.S. Supreme Court justices. Other nonprofits⁸², like the Invisible Institute, have also emerged and been rightfully recognized for their excellent reporting on racial inequality, gender-based violence, and police misconduct. Nonprofits are not a magic bullet⁸³, however, since they are often beholden to the whims of philanthropists, and have their own problems, especially when they begin to cover issues that might upset their donors. Most of these institutions have statements of editorial independence, but that does not prevent a donor from pulling the plug if they happen to start snooping too closely.

Other funding models have emerged, especially for online news sites that are journalist owned and operated. Returning to the story⁸⁴ of Deadspin: when the owner of G/O Media told the staff in 2020 that they needed to "stick to sports" and fired the editor-in-chief, all writers at the site quit en masse. Later that year⁸⁵, they formed their own site with nineteen employees, appropriately called Defector Media, that would be owned by the employees and funded through subscriptions. By the end of 2020⁸⁶, the company reported that they had 34,000 subscribers.

What does Defector think of the pivot to generative AI? In a brilliant bit of cheeky, irreverent satire⁸⁷, they lay bare the absurdity of it: editor Lauren

Theisen announced in mid-2024—shortly after the *Washington Post* CEO Will Lewis’s own announcement—that they were “promoting” Devin the Mixed-Reality Dugong (a poorly drawn marine mammal on a single sheet of printer paper) from Chief Metaverse Officer (a reference to Meta’s short-lived virtual reality world) to Chief AI/Metaverse Officer. Theisen proudly proclaims that Devin “stores and analyzes the personal data of every person that has ever accessed the internet . . . and, using its advanced intellect . . . write[s] the exact blog that it knows each particular reader wants to see.”

404 Media is another publication that adopted the subscription-based model, this time with a focus on technology and AI itself. A much smaller operation⁸⁸ at four full-time journalists, they formed the site in 2023 after departing from Motherboard, Vice Media’s tech division. Vice Media filed for bankruptcy⁸⁹ earlier that year because it couldn’t service the debt it owned to Fortress Investment Group and Soros Fund Management, two private equity firms who now own the media company. Emanuel Maiberg, one of the founders of 404 Media, said the site would be a “website by humans for humans about technology.” Accordingly, we find that it is one of the best for all-around coverage of AI. Jason Koebler, another one of the founders, stated that the site decided not to take venture capital funding from the start.

In a word, the introduction of AI into journalism follows the contours of other cheap technological fixes. The industry is still suffering from the inability to capitalize on digital advertising, and private equity and hedge funds are still stripping newspapers and digital-first publications for parts. Alternative institutions are needed for journalism to thrive, especially as those peddling AI threaten to use it to muck up the information ecosystem. It’s been a resounding “no” from those creating those alternative institutions on whether AI should be part of their futures.

Creativity Remains Distinctly Human—and Key

We’re writing about the seemingly disparate fields of art, science, and journalism in the same chapter because creativity is at the heart of all of these fields. These fields are often the targets for those seeking to prove the “intelligence” of their AI systems. But synthetic artifacts that look like the products of these fields aren’t evidence of creativity of the machines. To the

extent that there is any value at all, it is due to the creativity of the human workers who produced the original art, science, or journalism that was appropriated as training data.

Meanwhile, both the creation and output of these systems is damaging—to individual creators, science, and reporters, and to the larger ecosystems in which they are enmeshed. Today's synthetic media extruding machines are all based on data theft and labor exploitation, and enable some of the worst, most perverse incentives of each of these attendant fields. The use of these systems does further damage socially: displacing working artists and journalists, warping the practice of science, and polluting the information ecosystem. And their existence undermines the position and value of craft across these endeavors.

The relentless sales of these systems, despite all the costs, says something about the way AI hype people see art, writing, and knowledge. If text and image synthesis tools can write or draw something that is both plausible-seeming and would be technically difficult for people, that passes for creativity. But these tools do not speak to the human condition, nor do they have the ability to do science that can lead to novel, field-defining discoveries. And they surely do not have the ability to do the hard work of journalists, who often are the only ones who take on the relational labor of investigating and reporting on local issues to hold power holders to account.

All of these require a human mind and cultivation. One path forward is to recognize that fact, value the work of people in these fields, and support them. But AI boosters would rather believe that they can build artificial minds to do this kind of work. And some of them get so lost in that fantasy that they further imagine their creations will become all powerful, go off the rails, and doom us all.

Chapter 6

I'm Sorry, Dave, I'm Afraid I Can't Do That: AI Doomers, AI Boosters, and Why None of That Makes Sense

We now return¹ to where we started: the Russell Senate Office Building in late 2023, where Senator Chuck Schumer asked each member of this Insight Forum, the eighth in the series, what their $p(\text{doom})$ (or probability of doom) was. Although there were a number of people who addressed real threats from automation in attendance—including those from policy think tanks like Data & Society and labor advocates from the AFL-CIO, the U.S.'s major labor federation—those who believe that AI represents a serious existential threat to the human race were a strong presence in the room.

Schumer's forums were closed to the press, so we don't have everyone's precise numerical estimates (if they deigned to give such an estimation credibility), but the attendees' public opening statements² can give us an idea of where their estimates would lie. Jared Kaplan, cofounder of the AI company Anthropic, and Aleksander Mądry, head of preparedness at OpenAI, both spoke about “catastrophic risks”—situations in which a model could grow a mind of its own and result in a doomsday scenario—

and the ways in which their companies were addressing them. Meanwhile, Yoshua Bengio, professor of computer science at the University of Montreal, and Malo Bourgon, CEO of an organization called the Machine Intelligence Research Institute, both elaborated on how such a scenario may play out. Bengio stated that “there are many reasons AI could end up with undesirable goals. Each of them may put the system in conflict with humanity and give the system a reason to preserve itself despite human attempts to intervene.” Bourgon, meanwhile, was very alarmist indeed:

[T]he most likely outcome of developing smarter-than-human AI prior to solving alignment is human extinction. Present-day AI systems do not pose an existential threat, but there is a significant chance that systems in the near future will, as they become capable of performing increasingly complex multi-step tasks without humans in the loop.

While some of those at this forum spoke of myriad benefits of AI (and surely are profiting financially themselves from selling the idea that these systems are powerful), others spoke about imagined technologies in a manner that seemed to honestly stem from existential fears. We’d call the former Boosters, while the others we call Doomers. These groups are, counterintuitively, two sides of the same coin: the substance of the coin is the belief that the development of AI is inevitable and that that resulting technology will be both autonomous and powerful, and ultimately beneficial, if we play our cards right. It’s only the faces of the coin that differ. One shows a utopian world of abundance, the other a dystopian hellscape. Neither depicts the real harms of actually existing automation, at best dismissing them as less important than the imaginary existential threats.

Should you be scared of an autonomous AI agent? No. But you should be wary of the alarming ideologies behind both AI Doomerism and Boosterism. Doomerism/Boosterism serves to obscure, rather than illuminate, what’s at stake when it comes to the current AI boom. Moreover, these technologies *are* accelerating the real existential threat of human-made climate change, cutting into our already too-thin margin of time to mitigate it.

Okay, Doomer

AI Doomerism is the belief that, at some point, we will build an AI system powerful enough that it results in a mass extinction event. In one scenario, machines become “sentient” enough to have their own preferences and interests, which are markedly different from those of humanity. This is a common trope in science fiction: recall the machine intelligence Agent Smith in *The Matrix*, who tells an imprisoned human leader Morpheus that machines are the next stage in evolution on the planet, and therefore humans must be exterminated. The Doomers make this scenario feel more immediate by imagining that we’ve given machines the power to, say, run electrical grids and critical infrastructure (or more dramatically, nuclear weaponry and other military systems). Or maybe the artificial intelligence is able to manipulate people with language and trick humans into granting control over critical systems. For this reason, Doomers are very concerned with the speed in development of LLMs, mistaking their capacity to mimic human language for effective use of language.

Another scenario turns on the idea that we have developed a machine without the proper safeguards, and it has learned all the wrong lessons. The classic parable from within the AI doom literature is Nick Bostrom’s paper clip maximizer:³ a machine optimized to churn out paper clips has gone rogue and will stop at nothing to create paper clips—including killing all of humanity. We also see this kind of “misalignment” in science fiction, although less often. The HAL 9000 robot in *2001: A Space Odyssey* kills the crew because their presence would jeopardize the mission of reaching an alien monolith located on Jupiter. In this case, the AI system is just *too good* at its job and will stop at nothing to realize that goal.

Both of these scenarios have a boatload of tech boosters, AI Doomers, and, unfortunately, some political leaders worried. In March 2023, an organization called the Future of Life Institute released a letter⁴ asking all major AI labs to “pause” the training of AI systems more “powerful” than OpenAI’s GPT-4 for at least six months. The letter asks:

Should we develop nonhuman minds that might eventually outnumber, outsmart, obsolete and replace us? *Should* we risk loss of control of our civilization? Such decisions must not be delegated to unelected tech leaders. **Powerful AI systems should be developed only once we are confident that their effects will be positive and their risks will be manageable.**

The open letter boasts over thirty thousand signatories, including many prominent people we’ve met already in this book: business leaders like Elon

Musk and Emad Mostaque (CEO of Stability AI), AI researchers Yoshua Bengio and Stuart Russell, and tech entrepreneur-cum-former presidential candidate Andrew Yang.

Keeping the alarmist momentum going, two months later an organization called the Center for AI Safety published a twenty-two-word statement,⁵ signed by hundreds of AI researchers, as well as tech celebrities like Sam Altman and Bill Gates: “Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.”

These statements are all related to a set of end-of-times tech visions that have different monikers: “p(doom)”, “existential risk” (or “x-risk”), or, more recently, “critical harm”⁶ (as proposed in California legislation by friend-of-the-tech-industry state senator Scott Wiener). All of them see AI development as something that helps humanity in general but fret that, after reaching some tipping point—measured in “capabilities”, the amount of computing power needed to train a particular model, or the size of a model—the system can turn on its human handlers and overwhelm them.

Strangely enough, despite these visions, nearly all AI Doomers think that AI development is a net good. Many of them have built their careers off the theorization, testing, development, and deployment of AI systems. They have markedly *not* complained about the ways in which the business practices around AI have accrued financial, social, and political power to a very select few. They are mostly men. They are almost universally white and Western. With so many signatories on these letters, it is likely that the crowd includes people who believe in this vision religiously as well as people with financial and/or social incentives for signing. But for some of them, it’s not really about trying to save humanity, but rather a running of the con: the supposed danger of the systems is a splashy way to hype their power, with the goal of scoring big investments in their own AI ventures (like Musk and Altman) or funding for their own research centers (like Bourgon).

From the Doomer/Booster point of view, then, it is of critical importance to work out how to make sure that the supposed coming AI overlords are “aligned” with humanity’s goals. This has spawned what can generously be called a field of research, one that publishes many papers as preprints (without peer review) and manages to frequently grab headlines with them.

AI Safety and Alignment

This Doomer/Booster research field is called “AI safety”. Despite the name, this work does not come out of systems safety engineering⁷—which is a real field, with first principles, requirements, and engineering specifications—but instead from a truncated and misunderstood application of those ideas. AI safety is concerned with the fantasy of a runaway AI system resulting in a mass casualty event. There is no reason to believe any such event is remotely likely to occur. But that hasn’t prevented the proliferation of research centers with names like the “Center for AI Safety” at the University of California, Berkeley and the “Centre for the Study of Existential Risk” at the University of Cambridge, all generously funded by wealthy AI Doomers.

Closely related to the Doomer idea of making sentient AI “safe” is a peculiar idea called “alignment”. As the Doomer notion of risk is rooted in fantasies about sentient AIs going rogue, rather than anything people might do with automation, a primary focus of AI safety research is based on the idea that the solution is to design AI systems where are “aligned” with “human values”. A popular formulation, from a book authored by commentator Brian Christian, has named this the “alignment problem”⁸: “how to ensure that these models capture our norms and values, understand what we mean or intend, and, above all, do what we want.”

The alignment problem is considered the crown jewel in a panoply of problems for those involved in AI safety research. Those who engage in this work portray it as virtuous, in contrast to seeking AI development solely for profit’s sake. Alignment is supposedly about having some faith in humanity and directing our energies such that AI will be beneficial for everybody. In mid-2023, OpenAI loudly advertised⁹ that they would be working on a “superintelligence,” a superhumanly smart entity that could “lead to the disempowerment of humanity or even human extinction.” To control this big brain, they announced they were putting together a “superalignment” team. Their bright idea was to build an “automated alignment researcher” that could do more research and bring this superintelligence to heel. (It’s worth noting that OpenAI disbanded¹⁰ this team less than a year later, as its leads resigned from the company. One of them, Ilya Sutskever, formed his own,¹¹ called Safe Superintelligence Inc.)

Embedded in the idea of alignment is a premise with which we fundamentally disagree: that AI development is inevitable. We take umbrage with this on multiple grounds. For one thing, what is currently being developed as “AI” does not work, nor is it helpful, for an overwhelmingly large portion of people living on the earth today, especially people in the Majority World. Furthermore, as we’ve said elsewhere, there is no clear, precise definition of AI. Nor is there any solid evidence that the work of AI research now (or of the past seventy years) is on a path towards that undefined destination. Lastly, the development of mass automation tools is not socially desirable. If you’ve gotten this far in this book, you’ve seen how these technologies serve as a means of centralizing power, amassing data, and generating profit, rather than providing technology that is socially beneficial. In other words, this is a choice, one being made by powerful interests, but one that the rest of us do not have to go along with. Thus we reject this inevitability out of hand.

There are a lot more problems with the idea of alignment. First off, how do they define human values? The Asilomar AI Principles,¹² developed at a convening by the Future of Life Institute in 2017, include one that reads, “AI systems should be designed and operated so as to be compatible with ideals of human dignity, rights, freedoms, and cultural diversity.” But “rights” and “freedoms” differ from culture to culture, from group to group, and from person to person. Human values are also not static across time, nor are all groups granted the same dignities in the light of the law and human judgment. This isn’t to say that there are not rights that ought to be guaranteed to all people across time and place. The horrors of World War II and the Holocaust planted the seeds of the United Nations and the Universal Declaration of Human Rights. But such a universal declaration¹³ was not established after the abolition of the transatlantic slave trade and the mass violence that it involved. Nor have institutions that were intended to prevent such violence seriously addressed Western wars of aggression towards states in the Majority World. For instance, in 2002, after 9/11, the U.S. passed¹⁴ the “American Servicemembers Protection Act”, also known as the Hague Invasion Act, which authorizes the use of military force to liberate an American citizen from the International Criminal Court in The Hague, Netherlands, if they are held in that court’s custody. In light of this, we ask whose understanding of “human values” would be encoded in a

hypothetically aligned AI system, and why should the rest of the world just have to accept it?

Although they claim to be working in the best interests of humanity at large, those working in alignment summarily ignore the violence of the current moment, of the “doom” that has befallen large swaths of people under the guise of “AI” and progress. We’ve already discussed the violence of AI-driven warfare that has expanded the scope of acceptable targets to Palestinian civilians and the mass surveillance dragnet being mobilized against Black and brown people worldwide. Further violence is being done with automated tools¹⁵ used in refugee resettlement and at militarized borders. Tools sold with outlandish claims to functionality, such as detecting whether someone is lying to an immigration official or identifying countries of origin based on DNA, are routinely tested and deployed on African migrants fleeing war, famine, and genocide. Even away from any border, within the U.S., automated systems tear apart¹⁶ Black and Indigenous families by flagging these families for family separation—forcible removal of kids into the “care” of the state—at a much higher rate than white families.

When AI Doomers warn against existential risk, what they really mean is “existential risk for well-off, white, Western, and able-bodied people who are insulated from becoming climate refugees.” There are people who are—right now—experiencing awful conditions, losing access to rights and freedoms due to war, famine, and drought. We don’t need to construct a thought experiment like the paper clip maximizer to think of conditions which no human should be subject to, nor to start working on ameliorating them.

Can’t We All Just Get Along?

Some have argued¹⁷ that people working on AI safety and those working on issues of concrete harms of AI systems are after the same goals, and that AI safety subsumes the concerns of the latter. For instance, at the United Kingdom’s AI Safety Summit in late 2023, then-U.S. vice president Kamala Harris¹⁸ made a connection between these two fields in a bid to accommodate the AI Doomers, as well as those of us working on issues of concrete harm: “[We should] consider and address the full spectrum of AI

risk threats to humanity as a whole as well as threats to individuals, communities, to our institutions, and to our most vulnerable populations.” This rhetorical move asks us, the people working on the immediate harms to those in the here and now, to build bridges to those working on AI safety.

To this we say: no.¹⁹ The two fields start from different premises. Those working on current harms start from the position of civil and human rights, and are concerned with questions of racial and economic equity, freedom of movement across borders, personal safety from state violence, and rights to health and education, among others.

Those working on AI safety start from a place of a concern about fake scenarios, with a focus on a set of technologies that has proved only adept at predicting words from other words. Onto these outputs, they project intentions and invent scary entities those intentions might belong to. They often spotlight the computational power of AI technologies as the thing that needs to be put in check, and hypothetical harms rather than actually existing ones. Moreover, they have set up separate funding networks, flush with money from tech billionaires who see the development of AI as a boon to tech capital and power. They have separate academic research networks and citation networks,²⁰ which largely do not cite work by those concerned with current, actual harms.

There’s something perverse in the call for those concerned with social justice to subsume their goals and research under the umbrella of AI safety. This reminds us of the idea of color-blind racism²¹—when people in power don’t talk about race, they end up reinforcing racism and white supremacy precisely because they don’t talk about how existing institutions are harmful for people of color and need to be reformed or abolished. The people working on AI safety, predominantly white men in the United States, Canada, and the United Kingdom, claim to be safeguarding “everyone”, and that therefore their concerns are more existential (and more important) than those raised by people focused on racial and economic equity. Trying to “join forces” in such a context would mean constantly having to fight for the very validity of the issues we are trying to address, draining time and resources from actually addressing them.

However, we’ve seen that even the tepid declarations of “AI safety” have upset a set of powerful investors in AI, so much so they have decided to create their own “movement” of AI Boosters. The Doomers and the Boosters frequently behave as if between them they map out the full range

of possible positions on AI, but these two camps share a good deal in common.

Scratch a Doomer and Find a Booster

In October 2023, Marc Andreessen released²² a “Techno-Optimist Manifesto”, which outlined an explicitly “anti-safety” vision. We mentioned this screed in Chapter 2 because it describes, among many things, natalist fantasies that suggest people in “developed societies” (which we read as a dog whistle for “white people”) need to be breeding more, and proclaimed that the enemy of progress was “deceleration, de-growth, depopulation.” He warns his readers to guard against a whole slew of different bogeymen, including not only “existential risk” but also “sustainability” (as if climate were not a major concern), “trust and safety” (the organization within tech companies generally trusted with removing fraud, scams, nonconsensual pornography, child sexual abuse material, gore and violence, and other awful content), and “tech ethics” (we like to think that we’re included in this category). These are topics and communities of practice that have minimal overlap and no overriding theme. The incoherence of the writing conveys the incoherence of the position.

Andreessen calls this orientation “accelerationism”²³, “the conscious and deliberate propulsion of technological development” in order to grow the population and expand businesses, cities, and man-made development more generally. The accelerationist fantasy has a particularly libertarian bent: it sees unfettered markets as the means to generate unlimited energy and solutions to social problems. And it sees AI as a central part of this equation. The screed is written as a series of *we* statements. The accelerationists deify AI and also see themselves as gods for having created a new artificial life-form:

We believe Artificial Intelligence is our alchemy, our Philosopher’s Stone—we are literally making sand think.

This imagined technology is an everything machine, a universal technical fix for social problems.

We believe Artificial Intelligence is best thought of as a universal problem solver. And we have a lot of problems to solve.

We believe Artificial Intelligence can save lives—if we let it. Medicine, among many other fields, is in the stone age compared to what we can achieve with joined human and machine intelligence working on new cures. There are scores of common causes of death that can be fixed with AI, from car crashes to pandemics to wartime friendly fire.

We believe any deceleration of AI will cost lives. Deaths that were preventable by the AI that was prevented from existing is a form of murder.

Stark words, indeed, and ones that have found purchase in many of the wealthiest corners of the tech industry. Andreessen himself is a wealthy venture capitalist, but he's not the only tech baron on board with these ideas. Garry Tan²⁴, president of the influential startup accelerator Y Combinator, has signaled his support for this “movement”. In brief, they think that AI is an unmitigated good, and that nothing—not regulation, not artists, not even the AI safety people—should get in their way of developing it as they see fit. It's also worth mentioning that Tan has been aggressively active in San Francisco politics, ruthlessly opposing policies that might address the city's homelessness epidemic in favor of those that would criminalize being unhoused, all while drunkenly tweeting death threats²⁵ to sitting San Francisco politicians.

As a counter to a small sect of Doomers and AI safety people known as “effective altruists” (which we'll discuss below), Tan and Andreessen identify with a group calling themselves “effective accelerationists”, proudly placing “e/acc” in their X/Twitter profile names. People who have applied this label on X/Twitter (as of writing) include Martin Shkreli—the pharmacy exec who jacked up the price²⁶ of a toxoplasmosis drug (from \$18 to \$750 per pill) and did a stint in federal prison²⁷ for securities fraud—and, we're sad to say, Oakland's own favorite rapper with parachute pants, MC Hammer. Please, Hammer, don't hurt 'em. But this is honestly old wine in a new wineskin: a true belief that rampant capitalism is the solution to society's ills, a new picture frame around the California Ideology's²⁸ social liberalism with unfettered markets.

Doomerism and Boosterism are supposedly diametrically opposed camps, but both see AI as inevitable and desirable. Take, for example, the words of Doomer Eliezer Yudkowsky, a self-styled “AI researcher” who founded the Machine Intelligence Research Institute and is very influential in Doomer circles. In a *Time* magazine op-ed²⁹, Yudkowsky writes that he did not sign the six-month moratorium letter because he said it is “asking

for too little” and it would be like “the 11th century trying to fight the 21st century” if humans were ever asked to defend themselves against an AI system that had turned against humanity. He suggests that we need to track all large server farms in which AI systems are trained, and if needed, “destroy a rogue datacenter in an airstrike.” This vision is just as committed to the power and feasibility of sentient and autonomous AI systems as Andreessen’s but anticipates a dark turn.

Meanwhile, many Doomers/Boosters are busy at tech companies working on projects which have ill-defined goals but are sold as projects meant to serve humanity. For instance, Ray Kurzweil, a famous computer scientist who is known for, among other things, the idea that humans will merge with computers, was personally hired on at Google by its cofounder Larry Page³⁰. He suggests³¹ that we need to “make the AI systems safe and aligned with humanity’s wellbeing” and describes the AI tools he works on as being “inherently dual-use,” meaning that they can be used both for positive and destructive means. But these projects aren’t about using machine learning to solve specific problems, nor about creating specific, testable tools. Rather, they fashion themselves as building benevolent, “aligned” machine consciousness. And who better to guarantee that the consciousness being built is “safe” for everyone involved than the folks working at OpenAI, Anthropic, Microsoft, and DeepMind? Doomers style themselves as akin to nuclear engineers, building reactors for cheap energy, while knowing their work can be used to create weapons of mass destruction with a few modifications.

Not only are the Doomer and Booster beliefs very similar, but they also emerge from common intellectual origins. Computer scientist Timnit Gebru and philosopher Émile Torres coined the acronym “TESCREAL”,³² which stands for a collection of ideologies that hang together and undergird Doomerism/Boosterism. (Gebru was former head of Google’s Ethical AI team, and Torres was a former adherent to effective altruism, so both saw these ideologies up close.) TESCREAL stands for *Transhumanism*, *Extropianism*, *Singularitarianism*, *Cosmism*, *Rationalism*, *Effective Altruism*, and *Longtermism*. While all of these have their own tenets, they have overlapping sets of adherents and can be summed up in a few words. Those in the “TESC” part of the acronym maintain that humans are, at some point, going to merge with machines and then fly off Earth to colonize space. The “REAL” part of the acronym concerns an ultra-utilitarian ethic

that holds that we, as humans, need to optimize our behavior to do the most good with the resources available to us. Effective altruists frequently proclaim that they are maximizing the amount of good they are doing by sending their dollars to impoverished areas of the world. The paradigmatic example is mosquito nets, which are very cheap and prevent mosquitos carrying malaria from infecting people in their sleep. These efforts have the character of many misguided aid organizations, which believe that they are able to provide relief for an impoverished community with little understanding of local political, social, and economic conditions. They utilize the faulty logic of many strains of development economics³³ to fund technical interventions that, at best, save very few lives, and at worst reorient local incentive structures towards environmental ruin and corruption.

These ideologies would be fringe, cultish, and relegated to the back pages of the internet if it weren't for the massive capital that many of their adherents control and can influence: Elon Musk has said³⁴ he is a longtermist; Ray Kurzweil, who, as we mentioned above, was hired by a Google founder, is one of the originators³⁵ of singularitarianism and cosmism; Jaan Tallinn³⁶ was a founding engineer at Skype and Kazaa and has poured millions into Cambridge's Centre for the Study of Existential Risk and the Future of Life Institute. Most visibly, Sam Bankman-Fried, the founder of the cryptocurrency exchange FTX who was convicted of wire fraud and sentenced to twenty-five years in prison, was a major adherent of effective altruism.

We discuss these ideologies here because nearly all of them envision some version of an artificial general intelligence as part of the path towards a technological utopia and have developed a network of financial and social incentives to entice researchers to join their ideological communities and work on their projects. The ideologies direct their virtuous followers towards careers in the fields of AI safety and alignment, to ensure that the AGI they believe is coming will be benevolent towards humanity. The Centre for Effective Altruism's career advice center, 80,000 Hours³⁷, for instance, instructs its followers to get into "AI safety technical research" and "AI governance and coordination." In another example, Bankman-Fried's Future Fund bankrolled huge prizes³⁸ for AI safety research at large AI conferences such as NeurIPS.

Geburu and Torres compellingly argue that these ideologies have their origins in the Anglo-American eugenics movement. Not only do they have a direct lineage to eugenics—for instance, the person who coined the term “transhumanism” was Julian Huxley³⁹ (brother of novelist and author of *A Brave New World*, Aldous Huxley), a major figure in British eugenics. But they are also eugenic in contemporary argumentation. For instance, one of the main arguments for longtermism is that, according to its utilitarian logic, we should discount current-day suffering because we need to optimize technological development to seed the environment for the trillions of future humans who will colonize space. If you think that’s absurd, we’re with you. What this means is that the actually existing human suffering—borne primarily in the Majority World—is ignored for hypothetical threats of rogue algorithms. This is a eugenicist frame of seeing the world: the need of advancement for future white people at the cost of Black and brown people in the here and now.

The technologies in question are not the equivalent of nuclear weapons, nor even of nuclear reactors. They are more like a panopticon⁴⁰—philosopher Jeremy Bentham’s invention that allows a single prison warden to keep track of hundreds of prisoners at once. Or like the surveillance dragnets that track marginalized groups in the West and people on the move escaping dire conditions in their countries of origin. Or perhaps they are a toxic waste, salting the earth of a Superfund site. They are also a scabbing worker, crossing a picket line at the behest of an employer who wants to signal to the picketers that they are disposable.

The totality of systems sold as AI are these things, rolled into one. The danger is not from some hypothetical extinction-level event. The danger emerges from rampant financial speculation, the degradation of informational trust and environments, the normalization of data theft and exploitation, and the data harmonization systems that punish the people who have the least power in our society by tracking them through pervasive policing systems. But the Doomer/Boosters would have us looking the other way from all these real harms, bedazzled by their dystopian/utopian visions.

Extraordinary Claims, Extraordinarily Thin Evidence

In order for their story to hold water, and keep commanding policymaker attention, the Doomers have to convince their audiences of imminent artificial intelligence. Some of the Doomers speak of an AI system that becomes “smarter” than its designers and thus is able to design even more advanced systems, which design even *more* advanced ones, quickly overwhelming what we can even comprehend. This will result in a system that we can’t hope to control. Geoff Hinton,⁴¹ one of the so-called “Godfathers of AI” whom we met in Chapter 1, framed his Doomerist concerns to CNN journalist Jake Tapper by saying “there are very few examples of a more intelligent thing being controlled by a less intelligent thing.” We wonder when the last time was that Dr. Hinton suffered from any kind of food poisoning, and if he then decided that bacteria are more intelligent than humans. Meanwhile, a common refrain of the AI Boosters is to imagine systems that are “smart” enough to solve our problems for us (cure cancer! solve climate change!). And both the Doomers and the Boosters are asking the general public, investors, and policymakers to make decisions based on their extraordinary claims.

Extraordinary claims require extraordinary evidence, so you might hope that the Doomers and Boosters would go above and beyond in carefully substantiating the basis of their storytelling. But you would be disappointed. In order to evaluate “intelligence”, we’d first need a clear and operationalized definition of the concept, along with a compelling narrative of how it relates to the Doomer/Booster claims. Then we’d need a way to test for it, such that we can show that the test is actually measuring the property of interest. We saw in Chapter 2 that, throughout their history, measures of “intelligence” in humans have been based not in sound science but in eugenics and racism, so that’s already a bad start.

This isn’t a new problem: people have been pondering it since at least the 1950s, when Alan Turing wrote⁴² in his paper “Computing Machinery and Intelligence” that the question “Can machines think?” is too vague, as there are not sufficiently precise definitions of “machines” and “think”. His solution was to replace that question with a “closely related” one, which has since come to be called the *Turing Test*. The popular understanding of the Turing Test is that if a computer can fool a person into thinking it’s a person, that computer has “passed” the Turing Test and must therefore be considered intelligent.

In fact, Turing’s “imitation game” was a bit more convoluted. He set up a scenario in which an “interrogator” communicates with two other participants via teletype. In one version, the participants are a woman (trying to help the interrogator get it right) and a man pretending to be a woman (and trying to fool the interrogator); the interrogator’s job is to determine which participant is the man and which is the woman—a much more gender-bending version⁴³ than is usually discussed. In the second version, the participants are a man and a computer pretending to be a man. Turing suggests replacing the question “Can machines think?” with the question of whether the computer-pretending-to-be-a-man can fool an interrogator more consistently than the man-pretending-to-be-a-woman. The issue here, of course, is that without a definition of “intelligence” the question of how to measure “artificial intelligence” is a nonstarter. Turing tried to sidestep this with his imitation game but ended up proposing a setup that ran afoul of the very human tendency to make sense of language by imagining a mind behind it.

Another approach to measuring intelligence has been to use specific tasks as a proxy. For a long time, skill at the game of chess was a favorite example. In 1990, AI researcher and Dartmouth conference attendee John McCarthy popularized Soviet mathematician Aleksandr Kronrod’s phrase “chess is the *Drosophila* of AI”⁴⁴ because, he believed, the game was an efficient way to do rapid experimentation towards a long-term goal. That is, just as the fruit fly (*Drosophila*) is a boon to a certain kind of study in genetics (thanks to its short generations and ease of handling in lab environments), chess was thought to be a fertile ground in which to develop computer intelligence because it allowed researchers to quickly try out different approaches. This argument, of course, presupposes the relevance of chess to intelligence. We have become rather fond of the *Drosophila* metaphor for another reason: it epitomizes the fixation of researchers on one small contained problem—with its attendant issues and its own social history—while claiming to make progress towards grander goals.

This has been a persistent stumbling block for both Doomers and Boosters alike: when they talk about “general intelligence”, they don’t actually have a justifiable definition of the concept. Meta founder and CEO Mark Zuckerberg⁴⁵, for instance, has said that he doesn’t “have a one-sentence, pithy definition” of “general intelligence.” OpenAI has a nebulous definition of artificial general intelligence in their charter⁴⁶: “highly

autonomous systems that outperform humans at most economically valuable work.” (What does “outperform” mean? What is most “economically valuable work”?) Oddly enough, the term itself has taken on a type of spiritual significance, a sort of “we’ll know it when we see it” quality. Ilya Sutskever, when he was at OpenAI, urging employees to “feel the AGI,”⁴⁷ instructing them to chant it like a mantra. And if they can’t even define it, there’s no way that they are capable of measuring it.

With no clear definition of what they’re trying to measure, much less a solid foundation for claiming that they are actually measuring it, the Doomers and the Boosters are left to fall back on obfuscation and appeals to awe. In a blog post shortly after the release of Google’s LaMDA chatbot, Google VP Blaise Agüera y Arcas wrote⁴⁸:

Since the interior state of another being can only be understood through interaction, no objective answer is possible to the question of when an “it” becomes a “who”—but for many people, neural nets running on computers are likely to cross this threshold in the very near future.

With this rhetorical sleight of hand, Agüera y Arcas is claiming that, since we never have direct evidence of what’s going on inside other people’s heads, there is no way to know for sure that large language models, designed to provide facsimile conversations, aren’t something like people on the inside, too. If you let that smoke dissipate for a moment, though, you see that the argument is equally applicable to anything else we relate to through language (including books or the Magic 8 Ball toy), which is to say, not at all. As discussed in Chapter 2, we make sense of language by imagining the mind of the person who created it. With books, there is a mind behind the language, but the mind does not reside in the book. With the Magic 8 Ball or a chatbot, there’s an element of randomness in what language comes out. But that randomness does not constitute a mind.

Despite all the noise the AI Doomers are making, the evidence is thin—we’d say even nonexistent—for the scenarios they are raising the alarm about.⁴⁹ What we should be paying attention to and what we desperately need policymakers to be paying attention to is the impact of all those computations and the hardware needed to run them on the climate.

AI Is Hastening the Climate Catastrophe

Humanity is, however, facing an actual existential risk in the form of the climate crisis. The latest assessment report of the Intergovernmental Panel on Climate Change⁵⁰ warns:

Every increment of warming results in rapidly escalating hazards. More intense heatwaves, heavier rainfall and other weather extremes further increase risks for human health and ecosystems. In every region, people are dying from extreme heat. Climate-driven food and water insecurity is expected to increase with increased warming. When the risks combine with other adverse events, such as pandemics or conflicts, they become even more difficult to manage.

The Paris Agreement, a legally binding agreement signed by 196 Parties in 2016, has an overarching goal⁵¹ to keep the global average temperature to below the 2°C above preindustrial levels, and aim for no more than 1.5°C in temperature increase. Anything over that would have catastrophic consequences. In 2023, the UN’s climate body⁵² said we are not on track to meet these commitments, and that we have a “rapidly narrowing window to raise ambition and implement existing commitments.”

Against this background, you would think that people ostensibly concerned with the future of humanity would be working as hard as possible to build political will to decarbonize. Instead, what we see is a mad dash towards ever-larger models that require increasing amounts of computation (and therefore energy consumption) to train and use—with real and measurable environmental impacts.

The actual power-hungry computers are conveniently hidden in the fluffy, harmless-sounding metaphor of “cloud computing”, which information science researchers Alan Borning, Batya Friedman, and Nick Logler⁵³ point out is at direct odds with the actual materiality of these systems. When algorithm developers train a large statistical model, or when users submit a prompt to ChatGPT, it’s easy to imagine the processing happening in some abstract, virtual space. When the work is done on remote servers, we don’t even have to hear the fans keeping them cool! But in fact, everything about cloud computing is environmentally intensive:⁵⁴ the mining of metals and minerals required as raw materials, the use of large amounts of PFAS (“forever chemicals”) in the production of microchips, the energy required to both make hardware in chip fabrication plants and run the systems, the water used to keep data centers cool, and the e-waste produced as each generation of machines is retired in favor of the next.

The companies creating and selling this tech are far from transparent⁵⁵ about their environmental impacts, despite promises to be carbon-neutral or even carbon-negative in the near future. For instance, Google stated that they plan to achieve net-zero emissions by 2030, while Microsoft went one further and stated they plan to be carbon negative and even remove all carbon the company has produced since its founding in 1975.⁵⁶

However, more evidence is accumulating to the contrary. Computer scientist Sasha Luccioni and her colleagues estimated⁵⁷ the carbon footprint of training a particular large language model (among the largest released in 2022 at 176 billion parameters, but already dwarfed by the next year's models)—including equipment manufacturing—at 50 tons, or the equivalent of about a dozen flights between New York City and Sydney, Australia.

The current generation of systems is bad enough on its own, but on top of that there is a drive for scale happening at three levels, all harmful: larger models, more models, and ever-growing user bases.

Companies have largely stopped disclosing the number of parameters of their large models. OpenAI's GPT-4 is reported to have over one trillion parameters, nearly five times as many as GPT-3⁵⁸. We struggle to find any such information about most recent models as of 2024, but we also see no reason to believe anyone has abandoned the race to make ever-larger ones.

Moreover, the production of large models has been incredibly attractive for private equity and venture capital investors; therefore there are more models, created by a growing number of companies, in the marketplace. It is difficult to estimate precisely how many different models exist and how many times they have been trained and retrained—even a single model like GPT-4 has multiple different versions—but we can assume that the number is getting larger, not smaller.

Big Tech likes to tout their usage⁵⁹ of renewable energy as well as their innovations in creating more efficient processors and processing techniques. But as technology law scholar Michael Veale has said,⁶⁰ capitalism is missing from the equations. The efficiency gains offered by particular models and server providers is outdone by the sheer number of models being built and the amount of computation they require. Those computers need to be powered. Data centers are in such high demand⁶¹ that by 2034, global energy consumption by data centers is expected to top 1,580

terawatt-hours, as much as used by all of India, the most populous country in the world.

Lastly, these models have a growing userbase. OpenAI boasted 100 million users⁶² in its first two months and analysts estimated⁶³ in 2023 that the tool had 10 million users a day. Usage matters immensely for the final estimate of carbon emissions per model. You only have to train any given model once, but these models produce thousands of outputs per minute, each with their own environmental cost. Another study by Luccioni and colleagues estimated that for every two images you generate using a large text-to-image model, it's like charging your phone fully⁶⁴. These models are also water hogs: 500 milliliters of water⁶⁵ or about two cups are consumed for every 5 to 50 prompts ChatGPT generates a response to. Not only are these environmental costs hidden from users, but even users who would avoid these systems on environmental grounds can't easily opt out: The "AI Overviews" feature that Google added to search results in 2024 likely consumes 30 times more energy per query than just returning links⁶⁶. This feature was enabled by default, without users even having to invoke it.

This waste of resources is already directly impeding climate crisis mitigation goals. The *Washington Post* reported⁶⁷ in 2024 that coal-powered electricity plants in West Virginia, slated to be taken offline, are instead being kept running to feed the demand for data centers in Virginia. While we don't know how much of that data center usage can be attributed to AI systems (as opposed to other things like video-streaming services), it's telling that⁶⁸ energy demands for generative AI systems globally are expected to increase dramatically, rising tenfold between 2023 and 2026, and that Google, Microsoft, and OpenAI are putting billions of dollars into new data centers.

These companies have already admitted that they are dramatically missing their climate pledges because of generative AI. Microsoft sheepishly said their indirect emissions grew by nearly 30 percent compared to their 2020 baseline, while Google's grew by a whopping 48 percent compared to 2019.⁶⁹ Microsoft president Brad Smith told Bloomberg,⁷⁰ "In 2020, we unveiled what we called our carbon moonshot. That was before the explosion in artificial intelligence. So in many ways the moon is five times as far away as it was in 2020, if you just think of our own forecast for the expansion of AI and its electrical needs." If your

goalposts are continuing to shift (and the moon with it), why have goalposts at all?

Companies selling large language models and other computationally intensive technology like to talk in terms of cost/benefit trade-offs. Boosters see great benefits and easily find that these justify, for example, the environmental costs. Doomers, instead, debate the costs of imagined existential risks and the benefits provided by a hypothetical superintelligence. Neither focuses on the very real harms happening now and AI's contributions to the well-established existential risk of the ongoing climate crisis. Their focus is squarely in the concerns of white, well-off people.

It's easy to talk about benefits justifying the costs when you aren't the ones actually paying the costs. Not only the tech barons, but also most of their highly paid employees are fairly insulated from the climate crisis, compared to climate refugees, those living in tropical zones, and precarious clickworkers who, even in the U.S., can't afford air conditioning or easily escape from smoke-choked cities during ever-extending fire seasons. For a tech baron to confront these harms would require them to take their own culpability seriously and contemplate their own privilege. It's far more comfortable to sit back and pontificate on imagined scenarios where they are (also) victims, and dismiss anything else as "less existential."⁷¹

Politics After the Pause

AI Doomerism isn't worth taking seriously, and would be just best ignored if it weren't for the enormous amounts of funding and political influence associated with it. When the "AI Pause" letter came out in March 2023, it grabbed headlines and policymaker attention. When we and others pushed back against the letter as a distraction,⁷² we were scolded and told that we should instead join forces⁷³ with the "AI Safety" crowd behind those letters, be grateful that they had brought attention to the issues, and try to capitalize on it. For instance, LLM-skeptical AI booster Gary Marcus wrote,⁷⁴ "[I]t seems to me that they missed an opportunity to decry even more powerful versions of systems that have already proven problematic. This opportunity has been lost; I hope there will be more of a spirit of collaboration as subsequent proposals emerge."

But in fact, the emphasis on p(doom) did not galvanize policymakers towards sorely needed rights-protecting regulation. Instead, it distracted them. We see the fingerprints all over developments in the wake of the “AI Pause” letter and similar initiatives. The much-trumpeted “voluntary commitments” that the Biden administration facilitated in July 2023 with several prominent tech companies included things like committing to test models for “the capacity for models to make copies of themselves or ‘self-replicate.’”⁷⁵ Self-replication is a particular concern of the Doomers because it suggests out-of-control bots replicating into hard-to-scrub computer systems and taking over the world.

President Joe Biden’s October 2023 executive order on the “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,”⁷⁶ while building on well-grounded proposals such as the 2022 Blueprint for an AI Bill of Rights,⁷⁷ also led the U.S. National Institute of Standards and Technology (NIST) to start an AI Safety Institute.⁷⁸ Paul Christiano⁷⁹, a former OpenAI and Alignment Research Center researcher who is on record as predicting a p(doom) of 50 percent, was hired to run this new institute.

Doomerism has also infected the U.S. Senate’s work on AI. Not only did Senator Schumer invite his Insight Forum attendees to bloviate about their p(doom), but the report that resulted⁸⁰ from the forums spends much more time focusing on innovation and cybersecurity rather than harms to marginalized people, workers, the environment, and the information ecosystem. For example, imagined “capabilities” of AI systems are mentioned eighteen times, while “bias” is mentioned three times, and “climate” or “environment” (with respect to environmental damage) is not mentioned at all. But more importantly, Schumer and the other senators let themselves be distracted for a whole year by “experts” who were mostly AI hype peddlers talking about fantastical scenarios. The result is that instead of working on meaningful regulation during that year, based on all of the research already done into harms of the various kinds of automation sold as AI,⁸¹ Schumer and company produced a rather thin “road map” while AI companies were able to further entrench their harmful practices as the status quo.⁸²

We need to redirect attention away from speculative risks—no matter how exciting the action-movie sequences they conjure up might be—to the actual harms being done now in the name of AI. In the next chapter, we’ll

turn to how to best talk about AI and what can be done about it: by policymakers and regulators, by journalists, and by you, as a worker, consumer, and person who wants a better future.

Chapter 7

Do You Believe in Hope After Hype?

We've seen in this book that AI hype serves the purposes of people in power in a few different ways. It helps particular companies and their investors profit by selling the technology. It helps others get rich by giving them cover to collect (e.g., steal) and then launder massive amounts of data. It helps others still make short-term gains by replacing stable, better-paying jobs with ones that are both more precarious and less fulfilling. Lastly, it helps those who are wont to devalue the social contract by spinning the fiction that real social services—our collective responsibility to each other—can be replaced by cheap automated systems. At a time when the AI boosters are selling their wares across every sector at top volume, flooding everyone everywhere with the fear of missing out (or FOMO), and when it seems like the loudest voices opposing them are the equally hype-tastic AI Doomers, it can feel impossible to see a way through. Given that many policymakers have also jumped on the hype train, falling for the narrative that those selling the tech know best how to regulate it, that their job is to foster (nay, coddle) “innovation”, it can seem particularly hopeless.

But there are things that we can do, both individually and collectively, to resist AI hype—in this hype cycle and the next, no matter what the technology happens to be, nor how it's marketed. We can burst the bubble, through pointed questions and pointier ridicule. We can build up information literacy, both through our individual practices, and through

supporting and learning from other informational institutions—namely libraries. We can collectively shape innovation towards benefiting people at large rather than enriching the few, through enforcing existing regulation and crafting new. And we can and must resist narratives of inevitability through collective labor action and strategic refusal.

Strategies for Popping the Hype Bubble

It can be hard to push back against hype in places like the workplace, classroom, and clinician's office, and against peers, managers, and legislators who have been taken in by it. To speak up to say the emperor has no clothes is difficult. And it is doubly so when surprisingly many people, especially those in positions of power, seemingly want to be the naked emperors. That is, they want to believe the hype and convince everyone around to join them. But we have some strategies for you.

Asking Questions

One of the best strategies to cut through the hype is to ask questions about the brass tacks of the system being promoted. Whether you're reading the claims of AI researchers or corporate executives in news reporting, dealing with someone in your workplace extolling the benefits of some system they have bought or are thinking about buying, or concerned that your government is debating installing some system for public services (or surveillance), you can use these questions as a way to do a reality check for yourself, in the first instance. Beyond that, these questions can help you push back on hype, if you are in a position to ask them aloud.

What is being automated? What goes in, and what comes out? The first questions get to the concrete: What is the actual task being automated? What is being given to the automated system as input and what is it producing as output? Often the marketing copy for a system serves to obscure these simple facts. For example, the company Hippocratic AI¹ advertises “healthcare agents” and provides little biographies and photos for each. “Linda” is pictured as an East Asian woman, wearing blue scrubs like a nurse, and described as “a GenAI Healthcare Agent who follows up with

a discharged patient [by phone] after being admitted for Congestive Heart Failure.” The actual input here is spoken patient responses and questions. The output is word-shaped noises that are likely to show up next, given the training data. It should go without saying that this is not skilled nursing. As Michelle Mahon, director of nursing practice, told us², there’s more to nursing than just responding to prompts. Many times, people don’t have the words to describe their questions, and that’s when we need trusted professionals in the interaction.

Can you connect the inputs to outputs? What is the evidence that there is sufficient information in the input to determine the output? For example, researchers at Harrisburg University of Science and Technology³ claimed in 2020 that they had created a system that could tell, with 80 percent accuracy and “no racial bias,” whether someone was a criminal, based only on a picture of their face. They weren’t the first. In 2016, researchers at Shanghai Jiao Tong University⁴ made similar claims.

Applying our questions to this case, the input to these systems is relatively straightforward: a photo of someone’s face. The output, however, is much more vexed. On the face of it (no pun intended), it seems like a simple binary (yes/no) classification problem: Is this person a criminal? But criminality is not an inherent property of a person. The category of “criminal” is instead produced by an interaction of at least three things: the person’s behavior (very much constrained by their circumstances), social norms around what is considered criminal behavior, and the legal regime the person finds themselves in. So there cannot possibly be sufficient information in the input (an image of a face) to determine the output. You can’t tell if someone is a “criminal” by looking at them.⁵ Moreover, the claim that the system’s output lacks racial bias is preposterous, given everything we know about racial discrimination in policing and how poorly facial analysis systems perform with racial minorities.

Are these systems being described as human? The AI sales pitch involves language that anthropomorphizes—that is, it ascribes human characteristics—to the technology. We are anthropomorphizing creatures, so it takes an effort to keep some critical distance here. Calling a chatbot-based system an “AI teaching assistant”, which Morehouse College⁶ has said it will start using in fall 2024, for instance, suggests that it can do much more than what it actually does. Thinking of a human teaching assistant, we immediately imagine someone who cares about what their students are

learning, makes plans about how to help them understand better, spots possible cases of misunderstanding, and relates to the students as people gaining knowledge and skills that help them to grow into new roles in their community. The “AI teaching assistant” does none of these things. For any system sold as an “AI [human role]”, we can always ask: What motivates calling this thing by that name? What else do we expect of people in that role, and who is falling for the suggestion that this system can also do those things?

Relatedly, we should be on the lookout for ranked-order comparisons between automated systems and people. Researchers and corporations frequently brag that their system has achieved “superhuman” performance on various tasks. But the comparison to people implicit in the word “superhuman” belies a misapprehension of what software is. Software systems are tools, which people use to do things. We wouldn’t say that hammers have a “superhuman” ability to drive in nails, nor that airplanes have a “superhuman” ability to fly.

It matters what words we use when we talk about these technologies. For instance, in our writing, we don’t use the term “hallucination” to discuss the errors of LLMs, for two reasons. First, if it’s used tongue-in-cheek, it is making light of what can be symptoms of serious mental illness. Second, “hallucination” refers to the experience of perceiving things that aren’t there. But LLMs actually don’t have perceptions, and suggesting that they do is yet more unhelpful anthropomorphization. That means we also avoid assigning thought processes to these systems, or saying that they can “think”. Metaphors have power, they structure the frames of discourse, and they can subtly and insidiously encourage certain ways of understanding technology and the social systems it is embedded in.

How is the system evaluated? You’ll often see very impressive AI evaluation metrics, such as accuracy or performance, advertised by AI boosters. Hypers will claim that systems are over 80 percent accurate, or that they do better at a particular task, compared to people. OpenAI CTO Mira Murati, for instance, claimed that GPT-5⁷ would have PhD-level intelligence (which led many of us with PhDs to joke that the system would take on massive debt, burn out, and pursue a passion project in the food industry).

But those claims beg two questions: What was actually measured, and how does it relate to the intended use of the system? Often, this information

simply isn't available, which is a good indication that you're looking at pure hype. For example, the company SoundThinking sells a product called ShotSpotter to municipalities around the U.S., claiming that "ShotSpotter is a proven acoustic gunshot detection system that alerts law enforcement to virtually all gunfire within a city's ShotSpotter coverage area within 60 seconds"⁸ and that "[f]rom 2019–2021 the system had a 97% aggregate accuracy rate across all of our customers, including a very small false positive rate of less than 0.5% of all reported gunfire incidents."⁹ Neither of these claims is backed up with links to the studies that provided those numbers. However, audits in Chicago and New York City found that the vast majority of ShotSpotter alerts (87–91 percent) were false alarms.¹⁰ We would all have been better off if people in the right government positions had asked SoundThinking probing questions about their evaluation methodology up front, and then chosen not to buy into this unnecessary surveillance system that sends cops into situations believing there is an active shooter—a recipe for police violence¹¹, if we ever saw one.

Who benefits from this technology, who is harmed, and what recourse do they have? If this system is implemented, who and which organizations will benefit? Who will be harmed, both in cases where the system gives the right answer, and in cases where it doesn't? We have seen many, many cases in this book where automation makes the rich get richer, and the poor get poorer.¹² Sometimes, it's corporations (and their shareholders) using automation as an excuse to replace well-paying jobs with more precarious ones centered on babysitting the machines. In other cases, it's governments using automation as a Band-Aid over gaping holes in our social safety net. But the harms aren't always economic. For example, automated surveillance further terrorizes overpoliced communities and automated media synthesis machines pollute our information ecosystem.

Having identified harms, we can then ask what recourse is available to people who have been harmed. When systems hurt people, it is a major ethical breach if they have no way to ask for reparation and repair. Imagine a local government sets up software to identify families deemed "at risk", in order to determine which kids should be taken from their parents and put in foster care (like the example discussed in Chapter 4). Regardless of whether the kids are actually taken, having the "at risk" status recorded can lead to other problems for these families. Even if the local government has set up

an office that can handle complaints raised by those families, if that office is too understaffed to handle the volume of complaints, this means those people are effectively without recourse.

How was the system developed? What are their labor and data practices? Finally, we can and should always ask: How was the system developed, and what kinds of data labor and data collection practices were used? As we have seen, often the apparent humanlike competence of systems sold as “AI” is actually due to the work of humans, usually under exploitative working conditions. Just as with fast fashion, chocolate, and similar commodities, looking into labor practices puts us in a better position to make ethical decisions, in our own consumer patterns but more importantly whenever we find ourselves with the power to impact larger systems. And just as with those products, we should assume that in the absence of clear documentation of fair labor practices and data licensing, their supply chain relies on exploitation.

These are questions we can each ask, whenever we are in a position to make decisions about automation or influence others who are doing so. But they are also questions that we should look for in journalism. That is, it is worth seeking out and supporting journalism that asks the above questions, and eschewing that which focuses instead on gee-whiz platforming of companies’ claims to have created the next big thing.

In encouraging you to ask these questions, we are echoing the journalists who have spent the past several years focusing on AI and working towards improving coverage from the industry writ large. Garance Burke of the Associated Press (AP) developed a chapter for the *AP Stylebook*—a canonical reference for many journalists inside and outside the AP—on how to talk about AI. Burke instructs journalists¹³ to “get back to basics” and ask many of the same questions we pose here:

How do these systems actually work? Where are they deployed? How well do they perform? Are they regulated? Who’s making money as a result? And who’s benefiting? And also, very importantly, which communities may be negatively impacted by these tools?

Karen Hao, one of the foremost journalists on AI coverage, has been working with the Pulitzer Center to develop a program to train one thousand journalists in covering AI over the next two years with the purpose of holding companies to account. Hao’s work on the data labor

behind AI and the drama within OpenAI's walls has been crucial for understanding how this industry operates. In her training, Hao encourages¹⁴ journalists to focus on the people affected by these tools:

If we want a technology that is so consequential to actually benefit all of humanity—as OpenAI likes to say—the best way to understand how to do that is by covering the communities that are the most vulnerable and have had the least amount of agency in shaping the technologies thus far.

This effort is critical, as it plans to unsettle the centrality of the tech narrative around the Global North and back to the Majority World, where much of the labor occurs. It also aims to dispel the (intentional) obfuscation achieved by lumping all these technologies under the moniker of “AI”. And it ensures that the main story about AI is the human one, not the technical one.

Everyday Resistance

As a worker, you can push back against these tools at your job. If the past has been any guide, these tools will be introduced against the better judgment of workers because managers believe that it'll increase productivity. Unions like the Writers Guild of America and National Nurses United have made it a priority to resist the encroachment of AI in the workplace. Other unions that represent workers that bosses are dying to automate away—such as teachers, administrative assistants, and even dockworkers—ought to make it a priority to fight these efforts too.

On an individual level, sometimes methods of everyday resistance work the best. As a consumer, you have the option to simply not use AI tools. Don't use ChatGPT. Don't use Midjourney. When your favorite app offers a new, flashy AI-enabled tool, just don't click it. Companies and websites keep extensive metrics on where and when their users click or use a feature. If they find that no one is using their AI nonsense, they'll be more likely to remove that shiny button (with the stereotypical ✨ emoji) and quietly sunset the feature (that is, if they have any business sense).

As a savvy media consumer, you can also make fun of the awfulness of synthetic content. Has a friend posted AI-generated artwork on Instagram? Make fun of it. Did a brand use AI to generate six-fingered people enjoying

their product on the beach? Troll the hell out of them. It is your right. The more we can pierce the cultural bubble that Sam Altman and his kind live in, the better we can upset the idea that the encroaching of these systems into every area of life is inevitable. Synthetic media is cheap and tacky. Let them know.

Information Literacy and Libraries

Another opportunity we have for resisting AI hype is through our information practices. Many of the proposed use cases of LLMs are as information access systems, often as a direct replacement for search engines.¹⁵ This use case trades on a long-standing fantasy of making information access “frictionless”:¹⁶ you type in your question and you get the answer. But text synthesis machines are a terrible match for this use case, on two levels. As we’ve seen, they’re inherently unreliable, being designed to make shit up. But it’s the second level that we want to address here: friction in information access is actually not only beneficial, but critically important.

The friction that chatbot-based information access systems are trying to reduce or remove includes the work of deciding which of the links returned by a search engine has the information sought and where in the website the link points to can that information be found. Doing that work is important to building our understanding of the information landscape. Scanning a set of links gives us information about what information sources are available, and, if there are some we are expecting but don’t see, we have the opportunity to refine our search. Choosing a particular link and evaluating what we find there allows us to both situate the information we have found in its context and add to our understanding of how sources relate to each other and to our information needs. Finding contradictory answers on different pages—and, crucially, knowing the source of each—allows us to learn what kinds of knowledge are contested, who is doing the contesting, and how each of those sources fits into our own positions.

To make this concrete, imagine putting a health-related search into an old-fashioned search engine and getting back a collection of links that includes a page from a nearby university medical center, one from WebMD, one from (medical quack and former candidate for senator from

Pennsylvania) Dr. Mehmet Oz, and a link to a discussion forum where people navigating similar health journeys interact with each other. You will interpret the information from each of those sites based on your current evaluation of them, and use the information to update your evaluation.

If instead you input your question into ChatGPT or any other LLM-driven chatbot system, it might well be sourcing the output it extrudes from any one of those sites. But without seeing the information in context, you lose both the ability to use the context to frame the information and the opportunity to further build your understanding of the sources—not to mention the opportunity to join the community interacting in the message board.

Meanwhile, the tech sector is pushing chatbots as information access systems,¹⁷ hard. They started with just advertising them that way, often with embarrassing results. To take just one example, when Google launched Bard in early 2023, their demo contained an error¹⁸, claiming that the James Webb Space Telescope had taken the first photos of exoplanets, when in fact researchers using the Very Large Telescope in Chile had done so nineteen years before the Webb Telescope launched.

Moreover, LLMs-as-information-access-systems have been added to products in ways that make them almost inescapable. Google search results now include unasked-for “AI Overviews”, Meta has replaced its search bar in both Facebook and Instagram with “Meta AI”, and just about everywhere you might try to write text (Microsoft Office, LinkedIn, Google Docs, and email clients) the sparkle emoji is there, signifying software that will fill in words for you, whether they mean what you intend or not.

If we don’t resist this onslaught of synthetic text, the consequences are going to be bad, not just for our own information literacy, but also for the information ecosystem, and doubly so. There are the first-order effects of synthetic media spilling into the information ecosystem, and second-order effects of lower levels of information literacy impeding our collective ability to tend to and maintain that ecosystem.

Fortunately, there are ways to resist. At an individual level, we can overtly value authenticity. Refuse the apparent convenience of chatbot answers and insist on going to original sources for answers to our own queries. We can also refrain from sharing information that isn’t grounded in real sources (a best practice even before the introduction of chatbots-as-

search). Above and beyond that, we can work to support and maintain public information access systems.

There exists a set of practices around information access that both predates the web and search engines and remains a vital force in our communities. We're talking about the practices of librarians, libraries, and library science. These are the places to look for deeply informed ideas and practices around what it means for information access to be a public good. For example, library and information science scholars Anna Lauren Hoffmann and Raina Bloom contrast the ideology of "access" at companies like Google with the community-oriented care and service-centered ideology of access¹⁹ of libraries and librarians.

This means that supporting libraries is an important means of individual and collective resistance. And libraries definitely need broad public support, especially in the United States. Not only are they frequently dealing with threats of budget cuts,²⁰ but they have also been the target of organized book-banning campaigns. As Maggie Tokuda-Hall, of Authors Against Book Bans, explains,²¹ the purpose of these book bans isn't just the surface culture-war effects of whipping up moral panics around race, gender, and sexuality. They also serve as a massive drain on the time and resources of local libraries, impeding their function as engines of democracy. We can support libraries by voting for financial support for them, by showing up to public meetings to counter the book ban campaigns, and simply by visiting them, becoming part of the communities they foster and experiencing the practices of those communities.

Enforcing Existing Regulation

Boosters and Doomers would have us believe that AI is such a new technology that we need brand-new legislation to deal with the harms that may come of it. This would seem likely if you buy into the idea that we risk human extinction due to mathy maths. In reality, this is a delaying tactic offered by AI companies and their lobbyists to avoid any limitations on what they're building.

Luckily, we already have a good deal of regulatory tools to deal with the oversized claims, monopolistic practices, and other corporate malfeasance of AI boosters. In the U.S., regulators who handle consumer complaints,

communications infrastructure, civil rights, and labor issues have kept a keen eye out for hyped-up snake oil that companies are slinging. In April 2023, the Federal Trade Commission and several other federal agencies²² announced that there is no AI loophole—that is, companies producing generative AI must abide by current consumer and nondiscrimination rules on the books. These agencies added that their role is to regulate the actions of businesses, regardless of whether those businesses use automation.

The FTC has been very direct about this, especially on the nature of the hype itself. Companies cannot legally claim²³ that their products can do things that they patently cannot do; nor can they even claim that they use something called “AI” if that’s not what they are doing at all. If companies engage in this behavior, that can be considered as false advertising and be accompanied by regulatory penalties. As Michael Atleson, an attorney within the FTC Division of Advertising Practices, has put it,²⁴ “Your therapy bots aren’t licensed psychologists, your AI girlfriends are neither girls nor friends, your griefbots have no soul, and your AI copilots are not gods.”

Similarly, laws protecting workers’ rights can also protect against uses of AI in the workplace. The increase in remote work, spurred by the COVID pandemic, has led to a rise in the use of workplace surveillance tools, as employers found new ways to keep an eye on their workers.²⁵ Dubbed “bossware”, these tools include systems for surveilling those who work from home and anyone else whose place of work isn’t right under the boss’s nose: from programs that live on worker laptops, to facial analysis algorithms embedded in the cameras on Amazon delivery trucks.²⁶ In 2022, the U.S. National Labor Relations Board issued guidance²⁷ on unlawful worker surveillance, noting that such tools could run afoul of existing labor law by impinging on protected activity (e.g., talking about collective working conditions and forming workers’ organizations). Unfortunately, as surveillance becomes more and more pervasive, making the argument that such tools run afoul of workers’ rights may become a difficult sell for regulatory agencies that arbitrate labor rules.²⁸ In order to benefit from the protection provided by these laws, then, we must also work against the normalization of surveillance technology.

In general, we agree with the “zero trust” AI governance framework²⁹ outlined by Accountable Tech, AI Now, and the Electronic Privacy Information Center (EPIC)—named so because it asserts³⁰ that regulators should not assume that AI companies can effectively self-regulate, and

establishes sharp limits to the use of AI whenever possible. This is similar to how zero trust computer systems never assume that a user is trustworthy and require them to verify their identity and access level. This governance framework contains three principles: 1) Enforce existing laws, rather than acquiescing to the idea that new technology is somehow too new for existing laws to apply to; 2) establish bright-line rules for completely unacceptable uses of AI systems, rather than expecting regulators to conduct post hoc audits to ensure that systems don't mess up; and 3) force the burden onto companies to prove that their products are not harmful at every step of the product life cycle, rather than waiting for something to go terribly wrong.

Meaningful Further Regulation

Tech companies and their lobbyists like to try to cow policymakers with rhetoric about technology being so sophisticated and developing so quickly that it is impossible for anyone outside of tech to keep up with advancements. They suggest that policymakers couldn't possibly effectively regulate what they don't thoroughly understand. Policy written without an understanding of the innards of the technology, the argument goes, will always be lagging behind anyway (at best) and might harm innovation (at worst)³¹.

We frequently hear our politicians picking up on these messages, focusing on fostering "innovation" as the primary goal of policy-setting around technology. It even made it into the very title of Senator Schumer's report: "Driving U.S. Innovation in Artificial Intelligence: A Roadmap for Artificial Intelligence Policy in the United States Senate." Clearly, AI boosters want to be unfettered by regulation that might constrain their ability to amass power and capital, but they also sometimes even argue that it's a moral imperative to be able to innovate quickly, because (in their worldview) AI is going to save us all. For example, on an AI panel in 2018, in response to a call for a slower pace of research that leaves time and space to consult with the communities potentially impacted, Oren Etzioni³², then CEO of the Allen Institute for Artificial Intelligence, said:

Are you worried at all that when you slow things down, while you're going through that deliberative process, with the best of motivations, that people are dying in cars and people are dying in hospitals, that people are not getting legal representation in the right way? I think one reason for urgency is commercial incentives, but another reason for urgency is an ethical one. While we in Seattle comfortably debate these fine points of the law and these fine points of fairness, people are dying, people are being deported. So yeah, I'm in a rush, because I want to make the world a better place.

But in the years since Etzioni made those remarks, we haven't seen miraculous improvements in highway safety, health outcomes, or the treatment of migrants. Instead, we've been subjected to accelerating usage of AI as a pretext to surveil, arrest, and deport people; accelerating environmental impact of data centers to run the AI systems; and hundreds of car crashes, including at least seventeen fatal ones, as innocent bystanders are subjected to informal beta tests of Tesla's misleadingly advertised "Full Self-Driving" technology.³³ If we want innovation that is aimed at something other than profit maximization, we need to shape that innovation via regulation.

Shifting away from "innovation" towards regulation that protects people makes it clear that it's not the details of the technology that should be the focus, but its potential to impact people and society. In other words, the expertise needed to effectively regulate technology is expertise in sociotechnical systems, human and civil rights, and the crafting of laws such that they achieve the desired goals. While tech companies would like policymakers to feel snowed under by progress and to think they are perpetually playing catch-up, policymakers should be concerned with protecting rights and maintaining civic structures, things that don't change so fast.

We call on policymakers to look into the ways in which both the production and use of the new technology are trampling on human and civil rights and ask: how can those rights be better protected and what regulatory and enforcement mechanisms would be effective in protecting them? We see the following as prime areas for more legislation and rulemaking.

Transparency

Around 2017, half a dozen research groups³⁴ in both academia and industry began working on frameworks for documenting the datasets used to train

machine learning models and the models themselves. This work was motivated by a flurry of research showing that statistical models are very effective at modeling the biases in the data, and amplifying those biases in their output.³⁵ Models trained on existing data contain a representation of the patterns of the past,³⁶ including the effects of discrimination of all kinds. If we want to avoid replicating those patterns into the future, then we need to know what patterns any given model has been trained on before deciding how and whether to use it.

Documenting a dataset requires collecting and publishing information about how the data was generated or collected, why certain choices were made, who is reflected in the data, how the dataset will be maintained over time, whether people reflected in the data consented to being in the dataset, if the data are copyrighted, and much more. This all takes work and advance planning, and reflects an attitude of care towards datasets³⁷, especially towards the people reflected in them. This approach stands in stark contrast to the prevailing approach as Big Tech companies and startups grab everything on the internet that's not nailed down and claim it as their own.

While it may seem that documenting the data is a fairly technical ask, it is crucial information for any organization or regulatory body working in the public interest to understand what, precisely, is in the training datasets upon which AI systems are built. It's also a way to counter claims that a machine may have “emergent” properties, like Google CEO Sundar Pichai's claim that their chatbot “learned” Bengali³⁸ without being specifically trained on it. On the face of it, this claim is ridiculous. But researchers and regulators can confirm this only by having access to the data, or at least thorough documentation of it.

Even though much of the research on dataset documentation (and just as important, documentation around the models themselves) has been done by industry researchers (at Microsoft, IBM, Google, and Hugging Face, notably), it is also extremely clear that tech companies in general aren't going to carry out dataset and model documentation, especially to the level that the public deserves, without being required to do so by law. Hugging Face bucks the trend here³⁹ by facilitating documentation of models and datasets they host. This is having a positive impact on the availability of documentation, but there is still a long way to go:⁴⁰ as of February 2024, less than half of the models uploaded had industry-standard documentation

(what is called a “model card”) at all, and those that did didn’t necessarily provide all of the requested information. Meanwhile, Meta brags about releasing “open source” models⁴¹ but without actually releasing either the code for the training systems or documentation of the datasets they used.

We can’t have accountability without transparency. If an automatic decision system is going to be used to determine which benefits are allocated to which people, regulators and public interest technologists must be able to audit the training data driving that system. If a language model is being used as a component of a resume-screening system, then the U.S. Equal Employment Opportunity Commission (and analogous regulators in other countries) should have access to the training data behind that model to explore whether and to what extent it reflects biases against protected categories. If “AI Overviews” are going to be promoted as research aids in schools, then the public deserves to know what kinds of biases are going to be reproduced, automatically and at scale, for every schoolkid to see. Across all of these use cases, the public deserves to know what the systems are built out of.

Disclosure

Another kind of transparency is transparency about the fact of automation, otherwise known as disclosure. Is that a chatbot or a person? Does my résumé have to make it through an automated filter before it gets seen by a person in human resources? When I’m talking on the phone, is there an automated system labeling my speech, either based on my words or the way I said them? When is my photo being run through a system that matches it against a database of other photos, and by whom? We also need to know when we encounter the output of an automated system: is that handy-looking explainer written by a person, or is it synthetic text extruded from a bot? If we are dealing with translated text, was it translated by a person (perhaps using some automated tools, but in a position to verify and fix their output) or automatically translated without verification?

One approach that a few cities have already adopted is the AI register, which discloses where AI tools are used in city government. Both Helsinki and Amsterdam have published their own AI registers⁴², which are searchable through a public interface and allow residents to understand

where in public life they are subject to automation and automated decision-making. But these are only successful because municipalities have agreed to releasing this information, and this doesn't include private corporate uses of automation. We also need post hoc ways to detect automation when those in power aren't so forthcoming.

Being able to identify and filter out synthetic media will be critical to rehabilitating our information ecosystem. Unfortunately, the latest large language models mimic human writing too effectively for it to be feasible to detect their outputs as machine-generated. The only surefire way to detect their output is if it is watermarked at the source. A “watermark” is a hidden or less-perceptible imprint that marks the origin of a piece of media. It's common to see these when purchasing stock images from companies like Shutterstock and Getty Images.

In the toothless “voluntary commitments”⁴³ agreed to by tech companies and the Biden administration in July 2023, companies promised to work on watermarking, but only for audio and video content, and only for models more “powerful” than those they had already developed. It might seem like indelibly watermarking text is just impossible (since it can be easily copied-and-pasted and then edited to remove watermarks). However, it is still worth trying, and there are some novel technical methods for making more durable watermarks for text.⁴⁴ For another thing, not everyone running synthetic text generating machines will bother to remove even simple watermarks, meaning that including them would enable us to clean up some of the synthetic information spill. This is evident in the way that unintentional watermarks⁴⁵ such as phrases “as an AI language model” or “since my last knowledge update” have allowed researchers and journalists to identify synthetic text in, for example, scientific journals.

Whether it's via registers or watermarking, we ought to have a right to know when and where we are subject to outputs of automation. And we need the ability to compel the owners of the means of automation to tell us when we are viewing the outputs of a media synthesis system or subject to consequential decisions that are due to an algorithm.

Accountability and Recourse

Companies and organizations using automated systems, seeking a way to avoid taking responsibility for their decisions, may attempt to displace accountability to the systems themselves. The more opaque the system is and the more it is presented as an anthropomorphized “intelligence”, the easier it will be to pull off this trick. Weizenbaum saw this risk already⁴⁶ in the 1970s, writing:

Computers can make judicial decisions, computers can make psychiatric judgments. They can flip coins in much more sophisticated ways than can the most patient human being. The point is that they *ought* not be given such tasks. They may even be able to arrive at “correct” decisions in some cases—but always and necessarily on bases no human being should be willing to accept.

There have been many debates on “Computers and Mind.” What I conclude here is that the relevant issues are neither technological nor even mathematical; they are ethical. They cannot be settled by asking questions beginning with “can.” The limits of the applicability of computers are ultimately statable only in terms of oughts. What emerges as the most elementary insight is that, since we do not now have any ways of making computers wise, we ought not now to give computers tasks that demand wisdom.

Weizenbaum was not alone in this outlook. A slide reportedly from a 1979 IBM training presentation⁴⁷—which has become something of a meme within technology circles—stated, “A computer can never be held accountable, therefore a computer must never make a management decision.” When we make a decision, we (and others) have to live with the result. Per Weizenbaum, we want to make those decisions with wisdom. But wisdom isn’t just about the output of the decisions, but where they come from: lived experience, integration into community, and accountability for the consequences.

From a regulatory point of view, the first thing needed here is clarity in the application of existing laws. Even when automated processes are involved, accountability must rest with people and organizations. We recall again the FTC’s emphatic claiming of their jurisdiction: anything pertaining to false or unverifiable claims, and potential consumer risks, is fair game. Beyond that, it is worth reviewing existing and new legislation for any language that might need shoring up for additional civil and consumer protections.

Moreover, there needs be accountability throughout the AI supply chain. This should go above and beyond banning certain uses of the technology for critical areas (such as employment, health care, or housing), which has

been the strategy with the European Union's AI Act⁴⁸. Focusing regulation only on use cases suggests that large language models and image generators act as a type of benign infrastructure (or in the words of some AI Boosters at Stanford, "foundation models"⁴⁹), upon which other, harmful tools may be built. However, accountability should not rest simply with the company deploying a harmful tool. Accountability should also lie with the developers of the media synthesis machine. This is because the mechanism of harm does not stem solely from its deployment but is also due to the multiple design choices that went into building the model to begin with. AI companies need to be responsible for their data collection, data labor, model development, and evaluation. These all play a role in harms that are downstream from the large models themselves.

In terms of meaningfully protecting the rights of individuals, however, beyond accountability we also need recourse: the ability to set things straight, without undue delay, in the face of automated decisions. This is nontrivial, especially where automation is used to reduce the number of workers hired, precisely because automation can speed things up. When the automation is working well, that's a benefit. But when it is producing harmful outputs, we are faced with scaling up harm. Again, this is not a new observation. Another slide⁵⁰ attributed to the same IBM presentation reads, "Putting a bad system on line is like pouring gasoline into a fire." Rather than using automation to reduce staff hours, we should be requiring that whenever automated systems are set up that affect critical areas, such as people's health, education, or employment, sufficient staff should be available to handle cases where system output leads to harm.

Data Rights and Data Minimization

The modern technological paradigm for producing so-called AI systems requires relentless scale, both in terms of the size of the models and their datasets. This means that in order to get ahead, AI firms have to amass data at scales that would be prohibitively expensive to achieve carefully and ethically. Instead, AI firms are training their models with anything that isn't nailed down on the web. Those web pages include reams of news articles and Wikipedia pages, but also plenty of data from individuals, including their personal information, pictures and faces, voices, and writing. These

companies have abrogated the trust of the open web and in doing so, have accelerated the need for personal data rights. Widespread displeasure with this behavior can be seen in the way that many websites have quickly updated their robots.txt file—a trust-based convention for managing web crawling—to prohibit any web crawling at all.⁵¹ These conventions facilitated the development of the web and associated useful technology, like search engines. But trust-based web conventions are hardly going to stop data-hungry firms without consequence-bearing regulation backing them up.

There needs to be a framework for the protection of people’s individual data. In the European Union, a comprehensive piece of legislation passed in 2016 called the General Data Protection Regulation, or GDPR, prevents companies and governments from keeping user data indefinitely, requires those institutions to make a case for why they need user data, and allows people to have some semblance of control of what happens to their data⁵². Unfortunately, there’s no such federal legislation in the United States, although there is bipartisan agreement that it is necessary⁵³ and several states have laws⁵⁴ on the books that protect individual data. For instance, California’s Consumer Privacy Act (2018) offers the right to know if a business has collected data on them, a right to delete those data, and a right to opt out of data collection.⁵⁵ Illinois’s Biometric Information Privacy Act (2008) is more stringent, requiring written consent to have biometric personal data (i.e., fingerprints, facial pictures, DNA, retinal scans, and voiceprints) collected at all.⁵⁶

Even better than providing people with rights to their individual data is abiding by the principle of data minimization. Under a regime of data rights, people have some recourse if something goes wrong at the company, such as a privacy leak, or having the right to opt out of a software service. But there would be fewer privacy leaks and invasions of privacy if companies didn’t have those data to begin with. Data, once collected, can become dangerous, even if the data is supposed to be stored only on the device in question. So we need both data minimization regulations to prevent companies from collecting and holding our data, and we need to cultivate suspicion of any product predicated on massive data collection.

This is nicely illustrated by the feature called “Recall”⁵⁷ that Microsoft announced in June 2024. This system takes a screenshot of user activity every five seconds in order to offer a setup that Microsoft advertises as

“Just describe how you remember it, and Recall retrieves the moment you saw it.”⁵⁸ The tool uses “AI” to process the images. Although the company is slim on details, their documentation discusses Optical Character Recognition⁵⁹, in fact quite an old technology, which is used to extract text from images. Microsoft CEO Satya Nadella talked up⁶⁰ the fact that the data is only stored locally, on the user’s own device. This would seem like a privacy win, but security researchers quickly showed how the data could be accessed remotely—and that doesn’t even include cases where the malicious person accessing the data is in the user’s own home. Imagine a victim of intimate partner violence searching for information on how to leave their abuser, even carefully using incognito mode on their browser and clearing their search history, all for naught because their abuser can search “Recall” to see what they’ve been up to. This episode shows that above and beyond comprehensive privacy law, protecting data rights will require a shift in what kinds of data we allow companies to collect in the first place, even on our own devices.

Labor Protections

A key point in the AI hype that has every boss salivating is how much these tools are supposed to increase productivity, with fewer workers. But this has put workers on the back foot, fighting against threats which have been career-ending for many. The strongest counter to employers attempting to displace workers has been from unions who have met the challenge head-on, from the Writers Guild of America and the Screen Actors Guild (SAG-AFTRA) winning contracts that give workers more control over where and when these tools are deployed,⁶¹ to National Nurses United organizing against the encroachment of generative AI in the clinician’s office⁶². These workers have emphatically rejected so-called AI tools as replacement for their labor without their control and consent.

In some cases, employers floating the possibility of AI entering into workplaces has become a mobilizing factor itself. For instance, as Alex O’Keefe, a writer for the award-winning FX comedy *The Bear*, has said⁶³, that although Hollywood writers originally focused on striking due to the poor residuals paid by streaming services like Netflix and Hulu, AI quickly became a lightning-rod issue across many professions, one that writers,

actors, and stagehands could all stand in solidarity against. Like the writers and actors, pay is a major issue, but so is the potential of video-production tools like OpenAI's Sora to threaten stage and set design careers.⁶⁴ AI has become a galvanizing issue across the entertainment industry and we are seeing more labor organizations take note of the potential threats of automation.

Unfortunately, most of the workers in the United States are not unionized. Despite the "Hot Labor Summer"⁶⁵ of 2023, the Bureau of Labor Statistics reports that only 10 percent of U.S. workers are unionized, with 32.5 percent in the public sector, compared to 6 percent in the private sector.⁶⁶ Unionized workers can and should set the standards for the rest of their sectors, and existing unions should be strengthened. But we also need to lower the bar for unionization and create stronger protections across sectors for all workers. This depends on strengthening existing regulatory bodies both by executive order and by passing new laws.

The Protecting the Right to Organize (PRO) Act would help with some of this within the U.S. This piece of legislation would be the biggest expansion of labor rights since the New Deal⁶⁷. Among other pro-worker protections, it would help with worker misclassification: gig workers like Uber, Lyft, and DoorDash drivers, many of whom deal with algorithmic management of their livelihoods, would no longer be legally classifiable as independent contractors and instead would be considered to be employees of those companies⁶⁸. As employees, they then would be eligible for more worker protections. This may also discourage employers from downsizing their full-time employees, only to hire back workers as independent contractor replacements who are there only to clean up and babysit AI tools.

Better Living Through Regulation

Venture capitalists and CEOs looking to make more money by disrupting more markets like to argue that regulation "stifles innovation." But the contrary is actually true: effective regulation channels innovation towards what is broadly beneficial rather than just what makes the rich richer. It challenges and inspires the creativity of entrepreneurs, rather than giving them a free rein to chase the current flashy trend. It also guards against potential corporate liability down the line for intellectual property, privacy,

or labor violations. These benefits come from regulation imposed through governance and from pro-worker collaborative governance between labor unions and corporate management. They notably won't follow from "self-regulation" by corporations. Prosocial applications of automated pattern matching *are* possible, provided we follow some principled guidelines.

Building Socially Situated Technology

Let us be clear: we are not anti-technology, not even technology that involves the kind of pattern-matching algorithms used in "AI" systems. But we want to see technology that is designed with an understanding of both the needs and values of the people using it and of those it might be used on. In other words, we want technology that is created to strengthen and empower communities, not technology that reproduces and enables systems of oppression, consolidation of power, and environmental devastation.

First, the applications should be narrowly scoped, which means they should do one thing very well, rather than supposedly doing many things. That means we reject untestable "everything" machines that are marketed as "general purpose technologies". Instead, we want to see specific tools geared towards specific tasks. An example of a specific task for image processing is determining which parts of a photo should be in focus. An example of a specific task for language processing is machine translation. Considering what kind of data is required to build such applications brings us to our second guideline.

Applications should also be designed with respect for data rights, allowing, for example, consumers to gain information about their physical activity during a day without also sharing that data with a multinational corporation, a shady data broker, or whoever else would pay for it. The most beneficial technologies are developed with (or better yet by) the stakeholders most impacted by them, such as Te Hiku Media's work⁶⁹ on developing speech and language technologies for the te reo Māori indigenous community in Aotearoa New Zealand, while maintaining community guardianship of the training data and resulting tech. All of these properties entail a shift of power (and capital) away from the type of tech entrepreneurs whom venture capitalists usually favor and towards people

who better understand the social context of the technology they seek to build.

Lastly, people need the ability to choose when they want to be subject to automation. The simplest answer is that people should have the right to not be evaluated by a machine. This is similar to the EU's General Data Protection Regulation, which holds that people have a *right to be forgotten*⁷⁰, that is, not represented in data if they do not wish. People should have the ability to opt out, or better yet, be asked to opt in, while maintaining the right to revoke that consent at a later point.⁷¹

At no point, however, does calling any of this technology “AI” help. This term obscures how systems work (and who is using them to do what) while valorizing the position and ideas of power holders. Speaking instead about automation and data collection helps to make clear *who* is actually being benefited by this technology, and how. If we are to create a future that is populated with technologies we want, we “can’t only critique the world as it is,” as science and technology scholar Ruha Benjamin has written; we also “have to build the world as it should be to make justice irresistible.”⁷² Part of that vision means technology ought to be created with full participation of the people it impacts. Following disability justice advocates⁷³, we say “nothing about us, without us.”

Strategic Refusal

Never underestimate the power of saying no. Just as AI hypers say that the technology is inevitable and you need to just shut up and deal with it, you, the reader, can just as well say “absolutely not” and refuse to accept a future which you have had little hand in shaping. Our tech futures are not prefigured, nor are they handed to us from on high. Tech futures should be ours to shape and to mold.

We draw from many movements—including the Luddites, Black and Indigenous movements, and feminists—who have led the way on this point for decades before. The Luddites, as we discussed in Chapter 3, weren’t against technology, but they were against technology that did not serve them. The modern Luddite movement views technology⁷⁴ through the same light: enough with the promises of the future. We’re the ones who are having the figurative screws turned against us. In the same way, Black and

Indigenous movements have called for refusals—of being studied or “included” without being at all consulted or respected as full members of humanity⁷⁵. And feminists—in particular, feminists of technology and data justice—have refused to have their data used to feed the Big Tech machine. In their Feminist Data Manifest-No⁷⁶, a group of feminist scholars outline a set of refusals—including of rampant data collection, racialized and gendered violence that data collection and data-driven systems foster, and unfettered monetization of those data—and *commit* instead to a set of practices that center control over one’s life, taking back power from corporate entities, systems of relationality and reciprocity, and technologies that are based on thriving.

For AI boosters, the fully automated AI future is always just about to arrive. In 2023, Elon Musk said⁷⁷ his Teslas would have Full Self-Driving mode by the end of that year. Geoff Hinton proclaimed⁷⁸ in 2016 that we should stop training radiologists, because AI systems will soon be able to read medical images better than a human technician can. Meanwhile, rather than providing self-driving cars, Tesla has given us false advertising about self-driving mode, resulting in hundreds of crashes⁷⁹, and the Bureau of Labor Statistics has projected⁸⁰ a 6 percent growth in medical imaging jobs (including radiologists) from 2022 to 2033, faster than the average across all industries. The AI project has always been more fantasy than reality, starting from the field’s origins and the optimistic salesmanship of Minsky, McCarthy, and friends during that 1956 summer at Dartmouth. Jenna Burrell, a science and technology scholar and former director of research at Data and Society, has called this⁸¹ the “ever-receding horizon of the future,” which creates an urgency and a sense of hype about the next promised technology, while Anna Lauren Hoffmann has critiqued⁸² the metaphor of AI systems being in their “childhood”—we only need teach them better!

We should stop giving AI boosters the benefit of the doubt. They are indexing their fortunes—and mortgaging ours—on a future that doesn’t exist and that won’t suit us at all. We lift up calls for refusal and encourage you to add your own no to the growing chorus.

We don’t have to accept technologies that will do us harm, no matter how well they are tested or honed. Some technologies—like facial or emotional recognition—should be objected to on the grounds of what they are intended to do, and how they dehumanize and rank individuals. Media theorist and activist Zoé Samudzi, recognizing the landmark research⁸³ by

Joy Buolamwini and Timnit Gebru that facial classification systems are bad at recognizing Black faces, writes that these systems can never be made⁸⁴ to recognize Black humanity: “The problem is that [B]lack people are simply not human enough to be recognized by the racist technology created by racist people imbued with racist worldviews.” Sarah T. Hamid, policing tech campaign lead of the Carceral Tech Resistance Network, echoes this sentiment⁸⁵:

These technologies are not *incidentally* racist. They are racist because they’re doing the work of policing—which, in this country, is a racist job. There has been a lot of work devoted to proving that particular algorithms are racially biased. And that’s well and good. But there was no question that these algorithms were ever *not* going to be racist.

Responding to the proliferation of facial recognition systems at airports, Buolamwini and her organization, the Algorithmic Justice League, encourage travelers⁸⁶ to just say no to the facial recognition scans by the U.S. Transportation Security Administration and Customs and Border Protection.

For some people, there’s no opting out. People on the move—often fleeing war and famine—are frequently forced to give up their biometrics, including retinal scans, fingerprints, and blood draws, to gain access to food and shelter at refugee camps.⁸⁷ But for those of us who can, we should see our interactions with technology as an opportunity to resist, to give tech companies—who grow unfettered, only to surveil and break the social contract around us—the collective middle finger.

The Grimy Residue of the AI Bubble

In the summer of 2024, we started to see signs that venture capital is souring on AI—and since the hype is ultimately fueled, at least in part, by VC interest, perhaps the first inkling of the end of this bubble. That led us to ask: what kind of residue will the AI bubble’s popping leave behind?

The banking and investment giant Goldman Sachs,⁸⁸ estimating in early 2023 that two-thirds of jobs would be exposed to automation, with 18 percent of global work outright replaced by it, by June 2024 was saying the technology is nowhere near where it needs to be to replace jobs at that rate. Daron Acemoglu, labor economist at MIT,⁸⁹ published an estimate that

productivity gains from AI will be less than 0.53 percent over the next ten years. On an investment call after its 2024 Q2 earnings report, Alphabet CEO Sundar Pichai was grilled⁹⁰ with questions about when its big bet on AI—to the tune of \$12 billion per quarter—would pay off. David Cahn at VC giant Sequoia Capital wrote⁹¹ in June 2024 that AI firms need to turn something like \$600 billion in revenue for the AI bet to pay off.

In 2023, author and technology critic Cory Doctorow asked,⁹² “What kind of bubble is AI?” comparing the technology to prior cycles of hype, including the dot-com bubble, blockchain, nonfungible tokens (NFTs), and the metaverse:

Tech bubbles come in two varieties: The ones that leave something behind, and the ones that leave *nothing* behind. Sometimes, it can be hard to guess what kind of bubble you’re living through until it pops and you find out the hard way.

Doctorow thinks that the residue of the bubble popping will be minimal—large models will no longer be cost-effective to train, but small open-source models will remain, adapted to smaller, better-scoped tasks. If that’s all that the AI bubble leaves behind, then we’d be in a better place for society and science.

But we’re more pessimistic—and frankly upset—about what will be left behind once the AI bubble pops. Already, Google and Microsoft have sheepishly admitted⁹³ that they are far from reaching their climate goals, due to the large investment in AI. Data center growth is putting immense stress on existing power grids, not to mention are turning literal Black and Indigenous bodies into grist for the mill: data center and electric infrastructure construction in Northern Virginia is threatening to build atop a historical cemetery of free Black people, formerly enslaved Black people, and Indigenous Americans⁹⁴ for this insatiable machine. After the dust settles and Nvidia has stopped churning out shovels (e.g., the hardware called H100s⁹⁵) for the gold rush, what will be left behind? Will data centers go the way of shopping malls? Likely not—they’ll be repurposed for other massive computing projects. But what about those climate pledges? Will they be continued to be kicked down the road? To 2050? To 2075? If we leave the tech companies to their own devices, they’ll likely settle for too little, too late.

It’s not just the material infrastructure and the climate catastrophe, but the careers and industries that have been upset. Visual and conceptual artists

have discussed how their work has all but dried up—to which OpenAI CTO Mira Murati wryly remarked⁹⁶ that “maybe” such jobs “shouldn’t have been there in the first place.” In the gaming industry, there were over 10,000 layoffs⁹⁷ in 2023 alone, and, in a survey of 3,000 workers in the industry, half of the respondents report that their workplace was already using AI. Generative AI has the potential to ruin what were stable careers—not because the technology can do the work effectively or as skillfully, but it can produce convincing enough synthetic media to make certain jobs seem to managers redundant or requiring less skill than before.

After the AI bubble bursts, where do these careers go? Managers and execs aren’t likely to hire back career workers to do the skilled labor they once did. They’ll hire more contingent laborers, doing more with less. The residue of the bubble will be sticky, coating creative industries with a thick, sooty grime of limitless tech expansionism. This is the fallout of venture capitalists and tech entrepreneurs not pausing to think about who would be caught in the blast radius.

Keep the Next Bubble from Growing Through Collective Action and Comedy

Our current hype cycle will end, and this bubble will burst, with luck sooner rather than later. But history shows it’s unlikely to be the last. Next time around, some of the names of the technologies will differ, as will the specific claims about what they can do, but the dynamics of hype are likely to be the same. The more people are prepared to identify and resist the hype, the smaller and less destructive the hype bubble will be.

AI hype is widespread and pernicious. It is a lonely experience to feel like the only one in the room trying to push back when everyone else is caught up in a dreamland of what the technology will do for them. And the FOMO threaded through the hype makes all of this worse. The hype-mongers not only spin fantastical tales about what the tech can do, but they also work to convince you that if you don’t jump on the bandwagon, you’ll be left in the dust. Everyone on the bandwagon has to sing the same droning chant of AI mantras (“Feel the AGI!”, “Democratizing AI!”), getting louder and louder.

But resisting hype can also be empowering, grounding, and even joyful. It is empowering to reaffirm the value of our skills and expertise. It is grounding to lean in to the value of human-to-human connection, of being human together.⁹⁸ And it can be flat-out fun to find the silliest excesses of the hype machine and deflate it, with ridicule as praxis.

The next time you see a news headline about a new, revolutionary AI agent that will serve as your personal assistant, see a commercial about a tool which will help you and your child write a fan letter, or have a friend who heard a radio program about how Spotify's recommendation algorithm will stage a government coup, stop, breathe, and laugh at the foolishness. Whoever is behind that is selling you a bill of goods. The only way we build tech that works for us is to see past the façade and ask what kind of future we want to create, together.

Acknowledgments

This book wouldn't have been possible without a huge community behind us. Key to our ability to write this book are the people we have learned from and struggled alongside over the past few years, including Timnit Gebru, Meg Mitchell, Angelina McMillan-Major (of the original Stochastic Parrots crew), Remi Denton, Deb Raji, Amandalynne Paullada (of the original Grover crew), Abbie Harper, Abeba Birhane, Adam Conover, Adrienne Williams, Ali Alkhatib, Annalee Newitz, Brian Merchant, Charlie Jane Anders, Chris Gilliard, Johnathan Flowers, Joy Buolamwini, Karen Hao, Karla Ortiz, Meredith Whittaker, Nicole Moore, Ruha Benjamin, Safiya Noble, Sasha Constanza-Chock, Veena Dubal and many more.

A lot of the source material for the book came from our podcast *Mystery AI Hype Theater 3000*, and we are grateful to the knowledgeable and generous guests who have joined us, to our audience, to the Distributed AI Research (DAIR) Institute for supporting the podcast, and especially to our producer Christie Taylor. Christie is the author of many of the taglines of the podcast, including those we wove into the preface of this book.

We have endeavored to provide sources for all factual claims in the book and are indebted to our fact checker, Decca Muldowney, for verifying every last one, pointing out places where we were missing citations, and holding our feet to the fire when what we wrote wasn't precisely supported by the source. Any remaining errors are our own.

We were kicking around the idea of maybe writing a book for a while, but it might have remained just a vague future plan if it weren't for our agent, Ian Bonaparte, contacting us and convincing us it was time (maybe

even past time) to get it written and out in the world. That of course wouldn't have been possible without our editors Hollis Heimboach (at Harper) and Will Hammond (at Bodley Head), and their willingness to see this turned around on a tight timeline.

Between the initial draft and the final version, this book has benefited from the insights of many early readers, including Adam Conover, Adio Dinika, Amba Kak, Amy Zhou, Christie Taylor, Dylan Baker, Jim Bender, Kathleen Siminyu, Meg Mitchell, Sandy Bishara, Sasha Luccioni, Sheila Bender, Tamara Kneese, Ted Chiang, Timnit Gebru, Travis LaCroix, Vijay Menon, and Yindi Pei.

Finally, we thank our families and communities for their long-term support, their patience while we worked to get this book out, and for putting up with having a family member in the public eye.

Notes

Preface

1. sold as “AI”.: In this book, we use quotation marks in two ways: First, to set off words that we are mentioning in the text (including in passing, with scare quotes). In this case, we put punctuation outside of the quotation marks. Second, we use quotation marks to indicate that we are reporting what someone else said or wrote. In these cases, punctuation goes inside the quotation marks.
2. Deb Raji: Inioluwa Deborah Raji et al., “AI and the Everything in the Whole Wide World Benchmark,” in NeurIPS Dataset and Benchmark Track, 2021, <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/084b6fbb10729ed4da8c3d3f5a3ae7c9-Abstract-round2.html>; Amandalynne Paullada et al., “Data and Its (Dis)Contents: A Survey of Dataset Development and Use in Machine Learning Research,” *Patterns* 2, no. 11 (November 12, 2021), <https://doi.org/10.1016/j.patter.2021.100336>.

Chapter 1: An Introduction to AI Hype

1. Schumer began the conversation: Gabby Miller, “US Senate AI ‘Insight Forum’ Tracker,” *Tech Policy Press*, December 8, 2023, <https://www.techpolicy.press/us-senate-ai-insight-forum-tracker/>.
2. Schumer tweeted afterward: Chuck Schumer (@SenSchumer), “I co-hosted the 8th bipartisan AI Insight Forum with Senators Rounds, Heinrich, Young focused on preventing long-term risks and doomsday scenarios,” Twitter, December 6, 2023, <https://twitter.com/SenSchumer/status/1732500516662280477>.
3. Robert Williams was arrested: Kashmir Hill, “Wrongfully Accused by an Algorithm,” *New York Times*, June 24, 2020, <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html>.
4. Porcha Woodruff was arrested: Kashmir Hill, “Eight Months Pregnant and Arrested After False Facial Recognition Match,” *New York Times*, August 6, 2023, <https://www.nytimes.com/2023/08/06/business/facial-recognition-false-arrest.html>.
5. most known false positives: Kade Crockford, “How Is Face Recognition Surveillance Technology Racist?,” ACLU (blog), June 16, 2020, <https://www.aclu.org/news/privacy->

technology/how-is-face-recognition-surveillance-technology-racist; Katie Hawkinson, “In Every Reported Case Where Police Mistakenly Arrested Someone Using Facial Recognition, That Person Has Been Black,” *Business Insider*, August 6, 2023, <https://www.businessinsider.com/in-every-reported-false-arrests-based-on-facial-recognition-that-person-has-been-black-2023-8>; Khari Johnson, “How Wrongful Arrests Based on AI Derailed 3 Men’s Lives,” *Wired*, March 7, 2022, <https://www.wired.com/story/wrongful-arrests-ai-derailed-3-mens-lives/>.

6. irrevocably altered for the worse: In June 2024, Williams reached a settlement with the city of Detroit, including some changes in policy around the use of facial recognition technology (though not a complete ban) and \$300,000 in damages. We doubt either of these wins erase the pain of what he and his family went through, however. Andrea May Sahouri, “Man Wrongfully Arrested by Detroit Police with Facial Recognition Tech Settles Lawsuit,” *Detroit Free Press*, June 28, 2024, <https://www.freep.com/story/news/local/michigan/detroit/2024/06/28/man-wrongfully-arrested-with-facial-recognition-tech-settles-lawsuit/74243839007/>.
7. These image generation apps: Jason Koelber and Emanuel Maiberg. “A High School Deepfake Nightmare,” 404 Media, February 15, 2024, <https://www.404media.co/email/547fa08a-a486-4590-8bf5-1a038bc1c5a1/>.
8. deepfakes are of women: Pranshu Verma, “AI Fake Nudes Are Booming. It’s Ruining Real Teens’ Lives,” *Washington Post*, November 5, 2023, <https://www.washingtonpost.com/technology/2023/11/05/ai-deepfake-porn-teens-women-impact/>.
9. indiscriminately collected troves of images: In particular, most such models use the LAION-5B dataset, which was developed as a research project with the goal of showing how the collection of labeled images could be automated, allowing for such enormous datasets. Christo Buschek and Jer Thorp, “Models All the Way Down,” *Knowing Machines*, accessed July 23, 2024, <https://knowingmachines.org/models-all-the-way>.
10. images of a sex worker’s body: Samantha Cole, “Laws About Deepfakes Can’t Leave Sex Workers Behind,” 404 Media, June 4, 2024, <https://www.404media.co/laws-about-deepfakes-cant-leave-sex-workers-behind/>.
11. datasets also include: Samantha Cole, “Largest Dataset Powering AI Images Removed After Discovery of Child Sexual Abuse Material,” 404 Media, December 20, 2023, <https://www.404media.co/laion-datasets-removed-stanford-csam-child-abuse/>.
12. thousands of civilians were killed: Islamic Relief, “Huge Gaza Death Toll Is Likely to Be Even Higher than Reported,” *ReliefWeb*, December 20, 2023, <https://reliefweb.int/report/occupied-palestinian-territory/huge-gaza-death-toll-likely-be-even-higher-reported>.
13. “hell on Earth”: Ruth Michaelson and Kaamil Ahmed, “‘It’s Basically Hell on Earth’: Gaza City Left Totally Bereft of Healthcare,” *Guardian*, November 19, 2023, <https://www.theguardian.com/world/2023/nov/19/its-basically-hell-on-earth-gaza-city-left-totally-bereft-of-healthcare>; Al Jazeera, “In Pictures: ‘Hell on Earth’ in Gaza: Israel Strikes Hit Deir el-Balah,” December 2, 2023, <https://www.aljazeera.com/gallery/2023/12/2/hell-on-earth-as-israel-hits-deir-al-balah-in-central-gaza>.
14. Using a system called “The Gospel”: Yuval Abraham, “‘A Mass Assassination Factory’: Inside Israel’s Calculated Bombing of Gaza,” *+972 Magazine*, November 30, 2023, <https://www.972mag.com/mass-assassination-factory-israel-calculated-bombing-gaza/>.
15. precise as possible: Here, we follow Emily Tucker from the Center on Privacy and Technology, who argues we need to stop using the term “AI” for digital technologies that hurt individuals and communities, and be “as specific as possible and what the technology in question is and how it works.” Emily Tucker, “Artifice and Intelligence,” Center on Privacy and Technology (blog), March 8, 2022, <https://medium.com/center-on-privacy-technology/artifice-and-intelligence%C2%B9-f00da128d3cd>. See also Lucy Suchman, “The Uncontroversial

- ‘Thingness’ of AI,” *Big Data & Society* 10, no. 2 (July, 2023): 20539517231206794, <https://doi.org/10.1177/20539517231206794>.
16. “stochastic parrots”: Emily coined this term in the process of writing: Emily M. Bender, Timnit Gebru et al., “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜,” *FAccT ’21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, March 2021, 610–23, <https://dl.acm.org/doi/10.1145/3442188.3445922>.
 17. SALAMI: Stefano Quintarelli coined this appellation in November 2019. The acronym is coined in English but is spiced up with further flavor in Italian: *salami* or singular *salame* is used in Italian to call someone a dork. Stefano Quintarelli, “UPDATED: Let’s Forget the Term AI. Let’s Call Them Systematic Approaches to Learning Algorithms and Machine Inferences (SALAMI).” *Quinta’s Weblog*, November 24, 2019, <https://blog.quintarelli.it/2019/11/lets-forget-the-term-ai-lets-call-them-systematic-approaches-to-learning-algorithms-and-machine-inferences-salami/>.
 18. Neural Information Processing Systems: Neural Information Processing Systems, “37th Annual Conference of Neural Information Processing Systems (NeurIPS) Fact Sheet,” December 2023, 10–16, https://media.neurips.cc/Conferences/NeurIPS2023/NeurIPS2023-Fact_Sheet.pdf.
 19. held a secret auction: Cade Metz, “The Secret Auction That Set Off the Race for AI Supremacy,” *Wired*, March 16, 2021, <https://www.wired.com/story/secret-auction-race-ai-supremacy-google-microsoft-baidu/>; Rip Epsom, “Google Scoops Up Neural Networks Startup DNNresearch to Boost Its Voice and Image Search Tech,” *TechCrunch*, March 12, 2013, <https://techcrunch.com/2013/03/12/google-scoops-up-neural-networks-startup-dnnresearch-to-boost-its-voice-and-image-search-tech/>.
 20. 1999 Microsoft event: “Steve Ballmer at NET Conference Going Crazy about Developers! | 1999,” YouTube, accessed September 10, 2024, <https://www.youtube.com/watch?v=8fcSviC7cRM>.
 21. implicit fantasy explicit: Ashley Belanger, “Scarlett Johansson Says Altman Insinuated That AI Soundalike Was Intentional,” *ArsTechnica*, May 20, 2024, <https://arstechnica.com/tech-policy/2024/05/openai-pauses-chatgpt-4o-voice-that-fans-said-ripped-off-scarlett-johansson/>.
 22. a summer-long workshop: Much of this history is drawn from Jonnie Penn’s dissertation, “Inventing Intelligence: On the History of Complex Information Processing and Artificial Intelligence in the United States in the Mid-Twentieth Century,” (PhD diss, University of Cambridge, 2020), <https://doi.org/10.17863/CAM.63087>; Joseph Weizenbaum’s critique in *Computer Power and Human Reason* (1976); and the founding document itself of the Dartmouth workshop, John McCarthy et al., “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,” August 31, 1955, <https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/1904>.
 23. “significant advance”: McCarthy et al., “A Proposal,” 1.
 24. military, technical, engineering, and ideological: For a short historical introduction of the way that science and technology played a role in competition with the Soviets, see Audra Wolfe, *Competing with the Soviets: Science, Technology, and the State in Cold War America* (Baltimore: John Hopkins University Press, 2013).
 25. Minsky claimed: Marvin Minsky, “Heuristic Aspects of the Artificial Intelligence Problem” (Boston, 1956), quoted in Penn, “Inventing Intelligence,” 179.
 26. Weizenbaum developed a chatbot: Joseph Weizenbaum, “ELIZA—a Computer Program for the Study of Natural Language Communication between Man and Machine,” *Communications of the ACM* 9, no. 1 (January 1966): 36–45, <https://doi.org/10.1145/365153.365168>
 27. the style of a Rogerian psychotherapist: More precisely, the system consisted of a broader set of rules for a chatbot, and then a series of scripts that set the system up to take on specific characters or roles. The most famous of these is the psychotherapist, dubbed DOCTOR by Weizenbaum. Weizenbaum, “ELIZA.”

28. it was a convenient setup: Weizenbaum, “ELIZA,” 42.
29. A shocked Weizenbaum: Weizenbaum, *Computer Power*, 2–8; Ben Tarnoff, “Weizenbaum’s Nightmares: How the Inventor of the First Chatbot Turned against AI,” *Guardian*, July 25, 2023, <https://www.theguardian.com/technology/2023/jul/25/joseph-weizenbaum-inventor-eliza-chatbot-turned-against-artificial-intelligence-ai>.
30. In January 2023, Microsoft announced: Reed Albergotti, “OpenAI Has Received Just a Fraction of Microsoft’s \$10 Billion Investment,” *Semafor*, November 18, 2023, <https://www.semafor.com/article/11/18/2023/openai-has-received-just-a-fraction-of-microsofts-10-billion-investment>.
31. Reid Hoffman received: Niket Nishant and Krystal Hu, “Microsoft-Backed AI Startup Inflection Raises \$1.3 Billion from Nvidia and Others,” *Reuters*, June 29, 2023, <https://www.reuters.com/technology/inflection-ai-raises-13-bln-funding-microsoft-others-2023-06-29/>.
32. Sam Bankman-Fried: “Anthropic Raises Series B to Build Steerable, Interpretable, Robust AI Systems,” Anthropic (blog), April 29, 2022, <https://www.anthropic.com/news/anthropic-raises-series-b-to-build-safe-reliable-ai>.
33. quarterly investments, to AI and machine learning companies: Rani Molla, “Watch AI Eat the VC World in One Chart,” *Sherwood*, July 17, 2024, <https://sherwood.news/business/venture-capital-funding-ai-companies-investments/>.
34. a Palestinian man was arrested: Alex Hern, “Facebook Translates ‘Good Morning’ into ‘Attack Them,’ Leading to Arrest,” *Guardian*, October 24, 2017, <https://www.theguardian.com/technology/2017/oct/24/facebook-palestine-israel-translates-good-morning-attack-them-arrest>.
35. translating asylum claims: Johana Bhuiyan, “Lost in AI Translation: Growing Reliance on Language Apps Jeopardizes Some Asylum Applications,” *Guardian*, September 7, 2023, <https://www.theguardian.com/us-news/2023/sep/07/asylum-seekers-ai-translation-apps>.
36. thousands of British students: Daan Kolkman, “‘F**k the Algorithm’? What the World Can Learn from the UK’s A-Level Grading Fiasco,” *London School of Economics* (blog), August 26, 2020, <https://blogs.lse.ac.uk/impactofsocialsciences/2020/08/26/fk-the-algorithm-what-the-world-can-learn-from-the-uks-a-level-grading-fiasco/>.
37. Jared Mumm: Miles Klee, “Professor Flunks All His Students After ChatGPT Falsely Claims It Wrote Their Papers,” *Rolling Stone*, May 17, 2023, <https://www.rollingstone.com/culture/culture-features/texas-am-chatgpt-ai-professor-flunks-students-false-claims-1234736601/>.
38. A Tesla employee died: Trisha Thadani, Faiz Siddiqui, Rachel Lerman, Whitney Shefte, Julia Wall, and Talia Trackim, “Tesla Worker Killed in Fiery Crash May Be First ‘Full Self-Driving’ Fatality,” *Washington Post*, February 13, 2024, <https://www.washingtonpost.com/technology/interactive/2024/tesla-full-self-driving-fatal-crash/>; Ken Klippenstein, “Exclusive: Surveillance Footage of Tesla Crash on SF’s Bay Bridge Hours After Elon Musk Announces ‘Self-Driving’ Feature,” *The Intercept*, January 10, 2023, <https://theintercept.com/2023/01/10/tesla-crash-footage-autopilot/>.
39. lawyer Steven A. Schwartz: Josh Russell, “Sanctions Ordered for Lawyers Who relied on ChatGPT Artificial Intelligence to Prepare Court Brief,” *Courthouse News Service*, June 22, 2023, <https://www.courthousenews.com/sanctions-ordered-for-lawyers-who-relied-on-chatgpt-artificial-intelligence-to-prepare-court-brief/>; Kendra Albert (@kendraserra@dair-community.social), “Are you just catching up on the bonkers story about the lawyer using ChatGPT for federal court filings? This is a thread for you,” May 27, 2023, 7:53 a.m., <https://dair-community.social/@kendraserra/110441210244994168>; Roberto Mata vs. Avianca, Inc., No. 22-cv-1461 (United States District Court for the Southern District of New York, May 25, 2023),

- <https://storage.courtlistener.com/recap/gov.uscourts.nysd.575368/gov.uscourts.nysd.575368.32.1.pdf>. Sanctions: Roberto Mata vs. Avianca, Inc., No. 22-cv-1461 (United States District Court for the Southern District of New York, June 22, 2024),
https://storage.courtlistener.com/recap/gov.uscourts.nysd.575368/gov.uscourts.nysd.575368.54.0_8.pdf.
40. Meta released Galactica: Will Douglas Heaven, “Why Meta’s Latest Large Language Model Survived Only Three Days Online,” *MIT Technology Review*, November 18, 2022, <https://www.technologyreview.com/2022/11/18/1063487/meta-large-language-model-ai-only-survived-three-days-gpt-3-science/>.

Chapter 2: It’s Alive! The Hype of Thinking Machines

1. slightly conscious: Ilya Sutskever (@ilyasut), “it may be that today’s large neural networks are slightly conscious,” Twitter, February 9, 2022, <https://twitter.com/ilyasut/status/1491554478243258368>.
2. leaking private corporate information: Nitasha Tiku, “The Google Engineer Who Thinks the Company’s AI Has Come to Life,” *Washington Post*, June 11, 2022, <https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/>; Nitasha Tiku, “Google Fired Engineer Who Said Its AI Was Sentient,” *Washington Post*, July 22, 2022, <https://www.washingtonpost.com/technology/2022/07/22/google-ai-lamda-blake-lemoine-fired>.
3. “scientifically meaningful”: Blaise Agüera y Arcas, “Can Machines Learn How to Behave?,” Medium, August 3, 2022, <https://medium.com/@blaisea/can-machines-learn-how-to-behave-42a02a57fadb>.
4. “Sparks of Artificial General Intelligence: Early Experiments with GPT-4”: Sébastien Bubeck, et al, “Sparks of Artificial General Intelligence: Early Experiments with GPT-4,” arXiv April 13, 2023, <https://arxiv.org/abs/2303.12712>.
5. made to simulate it: Penn, “Inventing Intelligence,” 172.
6. put on a show: Lauren Feiner, “Sam Altman Wows Lawmakers at Closed AI Dinner: ‘Fantastic . . . Forthcoming,’” CNBC, May 16, 2023, <https://www.cnbc.com/2023/05/16/openai-ceo-woos-lawmakers-ahead-of-first-testimony-before-congress.html>.
7. earlier, simpler language models: Whereas these systems are usually called language models, they’re better understood as corpus models. That is, they are all trained on some specific (if often very large) collection of text. What is being modeled is the distribution of words in that particular collection of text. To the extent that the text is representative of some larger domain of language behavior, the model might work well with texts from different domains. But it is only ever an approximation of the full language or languages represented by the text.
8. proposed the first “n-gram” language models: Oscar Schwartz, “Andrey Markov & Claude Shannon Counted Letters to Build the First Language-Generation Models,” *IEEE Spectrum*, November 11, 2019, <https://spectrum.ieee.org/andrey-markov-and-claude-shannon-built-the-first-language-generation-models>.
9. a 1940s understanding of how neurons work: Frank Rosenblatt, “The Perceptron—a Perceiving and Recognizing Automaton,” Cornell Aeronautical Laboratory, January 1957, <https://bpb-us-e2.wpmucdn.com/websites.umass.edu/dist/a/27637/files/2016/03/rosenblatt-1957.pdf>.
10. Rosenblatt in the late 1950s: Penn, “Inventing Intelligence,” 82–115.
11. practical and profitable: Larry Hardesty, “Explained: Neural Networks,” *MIT News* (blog), April 14, 2017, <https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>.
12. correct answer given: The algorithm for making sure those changes extend through the whole network is called “backpropagation” and was developed twice in the early 1980s by David E.

- Rumelhart and Paul Werbos, working independently of each other. See Paul J. Werbos et al., “Applications of Advances in Nonlinear Sensitivity Analysis,” *System Modeling and Optimization: Lecture Notes in Control and Information Sciences*, vol. 38 (Berlin and Heidelberg: Springer, 1982), <https://doi.org/10.1007/BFb0006203>. In the early 2000s, when much more powerful computers were easily accessible, Yoshua Bengio and colleagues applied it to the problem of language modeling. See Yoshua Bengio, “A Neural Probabilistic Language Model,” *Journal of Machine Learning Research* 3 (2003): 1137–55, https://proceedings.neurips.cc/paper_files/paper/2000/file/728f206c2a01bf572b5940d7d9a8fa4c-Paper.pdf.
13. The initial BERT model: The original BERT paper by Jacob Devlin and colleagues, published at NAACL 2019 but posted as a preprint in 2018, presented results on two model sizes, “base” and “large”. The numbers in the text refer to the “large” model. Jacob Devlin et al., “BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, vol. 1 (Long and Short Papers), ed. Jill Burstein, Christy Doran, and Thamar Solorio (Minneapolis: Association for Computational Linguistics, 2019), 4171–86, <https://doi.org/10.18653/v1/N19-1423>.
 14. 3.3 billion words of text: Technically, language models are trained on “tokens,” which are usually pieces of words, rather than words as we ordinarily talk about them. For simplicity, we’re referring to tokens as words.
 15. Meta released Llama 3.1: Philipp Schmid et al., “Llama 3.1–405B, 70B & 8B with Multilinguality and Long Context,” *Hugging Face* (blog), July 23, 2024, <https://huggingface.co/blog/llama31>
 16. a series of rules: Weizenbaum, “ELIZA.”
 17. He was quite dismayed: Weizenbaum, *Computer Power and Human Reason*, 5–8.
 18. a very rich clue: See, for example, Michael J. Reddy, “The Conduit Metaphor: A Case of Frame Conflict in Our Language About Language,” in *Metaphor and Thought*, edited by Andrew Ortony, 2nd ed. (Cambridge: Cambridge University Press, 1993), 164–201, <https://doi.org/10.1017/CBO9781139173865.012>; Herbert H. Clark, *Using Language* (Cambridge: Cambridge University Press, 2007).
 19. babies won’t learn a language from passive exposure: This is noted in passing in C. E. Snow et al., “Mothers’ Speech in Three Social Classes,” *Journal of Psycholinguistic Research* 5, no. 1 (January 1976): 1–20, <https://doi.org/10.1007/BF01067944>, and tested more directly in Patricia K. Kuhl, “Is Speech Learning ‘Gated’ by the Social Brain?,” *Developmental Science* 10, no. 1 (January 2007): 110–20, <https://doi.org/10.1111/j.1467-7687.2007.00572.x>.
 20. Joint attention supports “intersubjectivity”: Dare A. Baldwin, “Understanding the Link Between Joint Attention and Language,” in *Joint Attention* (New York: Psychology Press, 1995), 131–58, <https://www.taylorfrancis.com/chapters/edit/10.4324/9781315806617-7/understanding-link-joint-attention-language-dare-baldwin>. See also the discussion in Emily M. Bender and Alexander Koller, “Climbing Towards NLU: On Meaning, Form, and Understanding in the Age of Data,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, ed. Dan Jurafsky et al., Association for Computational Linguistics, 2020, 5185–98, <https://doi.org/10.18653/v1/2020.acl-main.463>.
 21. frequent claims of AI researchers: To take one example, Chris Manning states that large language models “self-organize to discover and learn the structure of human language” and that “[i]t’s similar to what a human child does.” Edmund L. Andrews, “How AI Systems Use Mad Libs to Teach Themselves Grammar,” *Stanford HAI* (blog), July 23, 2020, <https://hai.stanford.edu/news/how-ai-systems-use-mad-libs-teach-themselves-grammar>.
 22. in an attempt to communicate: We thank Leon Derczynski for this insight.

23. devaluing what it means to be human: Curiously, sometimes it is not the tech boosters but others who seem willing to dehumanize themselves. Weizenbaum notes in his 1976 book that the psychiatrists promoting ELIZA as a replacement for their work justified this on the grounds that psychiatrists are simply following rules. He writes with some apparent shock: “What can the psychiatrist’s image of his patient be when he sees himself, as therapist, not as an engaged human being acting as a healer, but as an information processor following rules, etc.?” Weizenbaum, *Computer Power and Human Reason*, 6.
24. Sam Altman tweeted: Sam Altman (@sama), “i am a stochastic parrot, and so r u,” Twitter, December 4, 2022, <https://x.com/sama/status/1599471830255177728>.
25. senseless routines of code: Tarnoff, “Weizenbaum’s Nightmares.”
26. the sharpest and clearest critiques: For example, Joy Buolamwini and Timnit Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification,” Conference on Fairness, Accountability and Transparency, *Proceedings of Machine Learning Research* (2018) 77–91; Sasha Costanza-Chock, *Design Justice: Community-Led Practices to Build the Worlds We Need* (Cambridge, MA: MIT Press, 2020).
27. artificial general intelligence: “Our Structure,” OpenAI, accessed August 30, 2024, <https://web.archive.org/web/20240828044422/https://openai.com/our-structure/>.
28. learn quickly and learn from experience: The first version of this paper can be found at Sébastien Bubeck et al., “Sparks of Artificial General Intelligence: Early Experiments with GPT-4,” arXiv, March 22, 2023, <http://arxiv.org/abs/2303.12712v1>, 4; Linda S. Gottfredson, “Mainstream Science on Intelligence: An Editorial with 52 Signatories, History, and Bibliography,” *Intelligence* 24, no. 1 (January 1997): 13–23, [https://doi.org/10.1016/S0160-2896\(97\)90011-8](https://doi.org/10.1016/S0160-2896(97)90011-8).
29. “with Hispanics in between”: This letter was not a one-off effort. Gottfredson has been consistent in emphasizing inherent racial differences in intelligence, publishing articles that have been cited approvingly by white supremacist organizations and individuals, like former KKK grand wizard David Duke. She has also received funding from the Pioneer Fund, an organization that funds scientific projects dedicated to race science. Gottfredson has been a vociferous opponent of programs that would improve the economic and social conditions of racial minorities in the United States, such as affirmative action and hiring quotas. See Jack Anderson and Dale Van Atta, “Pioneer Fund’s Controversial Projects,” *Washington Post*, November 16, 1989, <https://www.washingtonpost.com/archive/lifestyle/1989/11/16/pioneer-funds-controversial-projects/0f05cde6-6586-462d-acfa-cfd3ae65567e/>; Southern Poverty Law Center, “Linda Gottfredson,” accessed June 12, 2024, <https://www.splcenter.org/fighting-hate/extremist-files/individual/linda-gottfredson>.
30. This is not “forbidden knowledge”: Ezra Klein, “Sam Harris, Charles Murray, and the Allure of Race Science,” *Vox*, March 27, 2018, <https://www.vox.com/policy-and-politics/2018/3/27/15695060/sam-harris-charles-murray-race-iq-forbidden-knowledge-podcast-bell-curve>.
31. IQ (or “intelligence quotient”) tests: Much of this history is taken from Stephen Jay Gould’s *The Mismeasure of Man* (New York: Norton, 1981).
32. U.S. Army men: The tests were conducted under haphazard conditions. Test takers were cramped into barracks and rushed through the test, without clear directions from proctors. Not only was the test discriminatory, but these conditions cast further doubt on any validity it might have had.
33. “admixture” of white blood: In her 1994 letter, in a section dedicated to understanding within-race variation in IQ, Gottfredson claims that “[a]lmost all Americans who identify themselves as black [*sic*] have white ancestors—the white admixture is about 20%, on average.” The implication here is that some Black people are smarter because of white ancestry. “Mainstream Science on Intelligence,” 15.

34. Industrial Revolution: Philippa Levine, *Eugenics: A Very Short Introduction* (New York: Oxford University Press, 2017).
35. Lewis Terman: Ben Maldonado, “Eugenics on the Farm: Lewis Terman,” *Stanford Daily*, November 6, 2019, <https://stanforddaily.com/2019/11/06/eugenics-on-the-farm-lewis-terman/>; Adam Morris, “Stanford’s White Supremacists,” *Guernica*, January 28, 2021, <https://www.guernicamag.com/stanfords-white-supremacists/>.
36. original cognitive test: His manuscript, “The Measurement of Intelligence” (1916), was used to prove the racial inferiority of Black people, the Indigenous, and Mexicans, to justify their lack of intelligence and the need for depressing their numbers via negative eugenicist policies in California and the rest of the U.S.
37. Musk tweeted: Elon Musk (@elonmusk), “Population collapse due to low birth rates is a much bigger risk to civilization than global warming,” Twitter, August 25, 2022, <https://twitter.com/elonmusk/status/1563020169160851456>.
38. having more children: Paris Marx, “Elon Musk’s New Baby’s Name Is Actually Less Absurd Than His Anti-Democratic, Quasi-Eugenicist Views,” *Jacobin*, May 7, 2020, <https://jacobin.com/2020/05/elon-musk-grimes-baby-population-democracy>.
39. Andreessen pushed against this: Edward Ongweso Jr., “Why Are Elon Musk and Marc Andreessen [sic] Obsessed with Birth Rates?” *Vice*, July 13, 2022, <https://www.vice.com/en/article/m7ge4v/why-are-elon-musk-and-marc-andreessen-obsessed-with-birth-rates>.
40. echoed Musk’s concern: Marc Andreessen, “The Techno-Optimist Manifesto,” Andreessen and Horowitz, October 16, 2023, <https://a16z.com/the-techno-optimist-manifesto/>.
41. Jamelle Bouie writes: “A whole coterie of Silicon Valley billionaires and millionaires have lent their time and attention to Hanania, as well as elevated his work. Marc Andreessen, a powerful venture capitalist, appeared on his podcast. David Sacks, a close associate of Elon Musk, wrote a glowing endorsement of Hanania’s forthcoming book.” Jamelle Bouie, “Why an Unremarkable Racist Enjoyed the Backing of Billionaires,” *New York Times*, August 12, 2023, <https://www.nytimes.com/2023/08/12/opinion/richard-hanania-eugenics-billionaires.html>.
42. AI startup xAI: “Musk Says His AI Firm xAI Is Rolling Out Chatbot Grok to X Premium+ Subscribers,” Reuters, December 7, 2023, <https://www.reuters.com/technology/musk-says-his-ai-firm-xai-is-rolling-out-chatbot-grok-x-premium-subscribers-2023-12-07/>; Will Knight, “Elon Musk Announces Grok, a ‘Rebellious’ AI with Few Guardrails,” *Wired*, November 6, 2023, <https://www.wired.com/story/elon-musk-announces-grok-a-rebellious-ai-without-guardrails/>; David Meyer, “What Elon Musk Is Really Building Inside His ChatGPT Competitor xAI,” *Fortune*, November 20, 2023, <https://fortune.com/2023/11/20/elon-musk-xai-chatgpt-competitor-grok-chatbot-ai-safety/>.
43. among the top investors: George Hammond, “Andreessen Horowitz Raises \$7.2bn and Sets Sights on AI Start-Ups,” *Financial Times*, April 16, 2024, <https://www.ft.com/content/fdef2f53-f8f7-4553-866b-1c9bfdbeea42>.
44. twenty-four Google researchers left: Matthew Lynley, “26 AI Experts Who Left Google to Start New Companies in 2022 as Machine Learning Ushers In a New Era of Lifelike AI,” *Business Insider*, December 26, 2022, <https://www.businessinsider.com/26-ai-machine-learning-experts-left-google-for-startups-2022-12>. Another researcher, Sara Hooker, joined the research nonprofit Cohere for AI, which is supported by Cohere’s commercial arm.
45. McKinsey suggested: “What’s the Future of Generative AI? An Early View in 15 Charts,” McKinsey & Company, August 25, 2023, <https://www.mckinsey.com/featured-insights/mckinsey-explainers/whats-the-future-of-generative-ai-an-early-view-in-15-charts>.

Chapter 3: Leisure for Me, Gig Work for Thee: AI Hype at Work

1. SAG-AFTRA reported: Irina Ivanova, “Why Are Hollywood Actors on Strike?,” CBS News, July 19, 2023, <https://www.cbsnews.com/news/sag-aftra-strike-hollywood-union-actors/>.
2. mega margin-maximizing consultants: See, for instance, the claim by McKinsey that generative AI technology could increase labor market productivity by 0.1 to 0.6 percent by 2040. Lareina Yee and Michael Chui, “POV: AI Could Increase Corporate Profits by \$4.4 Trillion a Year, According to New Research,” *Fast Company*, July 7, 2023, <https://www.fastcompany.com/90919913/pov-ai-corporate-profits-4-trillion-research>.
3. Majority World: In this book, we use the term “Majority World” to replace common terms for the global majority that is outside of countries such as the U.S. and Western Europe, what has also been called the Global South. See Rosemary M. Campbell-Stephens, *Educational Leadership and the Global Majority: Decolonising Narratives* (Cham, Switzerland: Palgrave Macmillan, 2021).
4. Recent histories of the Luddites: Brian Merchant, *Blood in the Machine: The Origins of the Rebellion Against Big Tech* (New York: Little, Brown, 2023); Gavin Mueller, *Breaking Things at Work: The Luddites Were Right About Why You Hate Your Job* (London and New York: Verso, 2021).
5. losing their jobs: Merchant, *Blood in the Machine*, 45.
6. local guild and family systems: Merchant, 52–56.
7. “General Ludd’s wives”: Richard Conniff, “What the Luddites Really Fought Against,” *Smithsonian Magazine*, March 2011, <https://www.smithsonianmag.com/history/what-the-luddites-really-fought-against-264412/>.
8. the loss of jobs, health, and community: Conniff, “What the Luddites Really Fought Against.”
9. coined it in 1948: Margaret Wiley Marshall, “‘Automation’ Today and in 1662,” *American Speech* 32, no. 2 (1957): 149–51, <https://doi.org/10.2307/453032>.
10. nearly \$41,000 a year: Sarah Cwiek, “The Middle Class Took Off 100 Years Ago . . . Thanks to Henry Ford?,” NPR, January 27, 2014, <https://www.npr.org/2014/01/27/267145552/the-middle-class-took-off-100-years-ago-thanks-to-henry-ford>; inflation calculated from <https://www.usinflationcalculator.com/>.
11. Ernest Nagel said: Ernest Nagel, “Automatic Control,” *Scientific American* 187, no. 3 (1952): 44–47, <http://www.jstor.org/stable/24950779>.
12. Wassily Leontief argued: Wassily Leontief, “Machines and Man,” *Scientific American* 187, no. 3 (1952): 150–64.
13. An autoworker discussing: Quoted in Raya Dunayevskaya, *Marxism and Freedom: From 1776 until Today* (New York: Bookman Associates, 1958), 269.
14. “man-eaters”: Andy Phillips, *The Coal Miners’ General Strike of 1949–50 and the Birth of Marxist-Humanism in the US* (New York: News & Letters, 1984), 12–14.
15. several warehouse deaths: Emily Guendelsberger, “I Worked at an Amazon Fulfillment Center; They Treat Workers Like Robots,” *Time*, July 18, 2019, <https://time.com/5629233/amazon-warehouse-employee-treatment-robots/>; Will Evans, “How Amazon Hid Its Safety Crisis,” *Reveal*, Center for Investigative Reporting, September 29, 2020, <https://revealnews.org/article/how-amazon-hid-its-safety-crisis/>; Michael Sainato, “US Regulators Launch Investigation into Worker Death at Amazon Warehouse,” *Guardian*, May 23, 2023, <https://www.theguardian.com/technology/2023/may/23/amazon-warehouse-death-fort-wayne-indiana>; “OSHA Opens Second Amazon Probe Following Two More Worker Deaths in New Jersey,” CBS News, August 12, 2022, <https://www.cbsnews.com/news/amazon-worker-deaths-rafael-reynaldo-mota-frias-osha-investigation-new-jersey/>; Eli M. Rosenberg, “Amazon Workers Demand More Details in Warehouse Employee’s Death,” NBC News, July 22, 2022, <https://www.nbcnews.com/business/business-news/amazon-worker-death-prime-day-new-jersey-rcna39534>.

16. in their trucks: Lauren Kaori Gurley, “Amazon’s AI Cameras Are Punishing Drivers for Mistakes They Didn’t Make,” *Motherboard*, September 20, 2021, <https://www.vice.com/en/article/88npjv/amazons-ai-cameras-are-punishing-drivers-for-mistakes-they-didnt-make>.
17. self-published a report: Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, “GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” arXiv, March 17, 2023, <https://arxiv.org/abs/2303.10130>.
18. “General Purpose Technology”: Erik Brynjolfsson and Andrew McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (New York: Norton, 2014), 75–77.
19. labor economists: Carl Benedikt Frey and Michael A. Osborne, “The Future of Employment: How Susceptible Are Jobs to Computerisation?,” *Technological Forecasting and Social Change* 114 (2017): 254–80; D. H. Autor, F. Levy, and R. J. Murnane, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics* 118, no. 4 (November 1, 2003): 1279–1333, <https://doi.org/10.1162/003355303322552801>.
20. Goldman Sachs estimates: Jan Hatzius et al., “The Potentially Large Effects of Artificial Intelligence on Economic Growth (Briggs/Kodnani),” Goldman Sachs Economic Research, March 26, 2023, <https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html>.
21. Critical thought is co-created: Similar points are made by Ted Chiang and Anna Rogers, Ted Chiang, “Why A.I. Isn’t Going to Make Art,” *The New Yorker*, August 31, 2024, <https://www.newyorker.com/culture/the-weekend-essay/why-ai-isnt-going-to-make-art>; Anna Rogers, “On AI-Assisted Writing in Graduate School,” *Hacking Semantics* (blog), August 24, 2024, <https://hackingsemantics.xyz/2024/writing/>.
22. the broader public: Mia Sato, “The Perfect Webpage,” *The Verge*, January 8, 2024, <https://www.theverge.com/c/23998379/google-search-seo-algorithm-webpage-optimization>.
23. in the business of selling ads: On the impact of an ad company presenting itself to the public as an information access service, see Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York: New York University Press, 2018).
24. “enshittification”: Cory Doctorow, “‘Enshittification’ Is Coming for Everything,” *Financial Times*, February 7, 2024, <https://www.ft.com/content/6fb1602d-a08b-4a8c-bac0-047b7d64aba5>.
25. the OpenAI board briefly ousted: Karen Hao and Charlie Warzel, “Inside the Chaos at OpenAI,” *Atlantic*, November 19, 2023, <https://www.theatlantic.com/technology/archive/2023/11/sam-altman-open-ai-chatgpt-chaos/676050/>.
26. bragging: Nilay Patel, “The CEO of Zoom Wants AI Clones in Meetings,” *The Verge*, June 3, 2024, <https://www.theverge.com/2024/6/3/24168733/zoom-ceo-ai-clones-digital-twins-videoconferencing-decoder-interview>.
27. the most secure: Hammond Pearce, Baleegh Ahmad, Benjamin Tan, Brendan Dolan-Gavitt, and Ramesh Karri, “Asleep at the Keyboard? Assessing the Security of GitHub Copilot’s Code Contributions,” 2022 IEEE Symposium on Security and Privacy, San Francisco, 2022, 754–68, <https://ieeexplore.ieee.org/document/9833571>.
28. Armin Ronacher: Armin Ronacher (@mitsuhiko), “I don’t want to say anything but that’s not the right license Mr Copilot,” Twitter, July 2, 2021, <https://twitter.com/mitsuhiko/status/1410886329924194309>. For more on the fast inverse square root, see Wikipedia, “Fast Inverse Square Root,” last modified May 22, 2024, https://en.wikipedia.org/wiki/Fast_inverse_square_root.
29. training data: Nicholas Carlini et al., “Extracting Training Data from Diffusion Models,” *Proceedings of the 32nd USENIX Security Symposium*, Anaheim, CA, August 9–11, 2023, <https://dl.acm.org/doi/10.5555/3620237.3620531>.

30. against OpenAI: Michael M. Grynbaum and Ryan Mac, “The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work,” *New York Times*, December 27, 2023, <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>.
31. verbatim text from the newspaper: The New York Times Company v. Microsoft Corporation, OpenAI, Inc., OpenAI LP, OpenAI Gp, LLC, OpenAI LLC, OpenAi Opco LLC, OpenAI Global LLC, OAI Corporation LLC, and OpenAI Holdings, LLC (S.D.N.Y. 2023), https://nytco-assets.nytimes.com/2023/12/ NYT_Complaint_Dec2023.pdf.
32. Adam Conover is marching: Adam Conover (@TheAdamConover), “Writer’s Strike 101: What Are We Fighting For?,” YouTube, May 23, 2023, 00:38 <https://www.youtube.com/shorts/KYANliU6B6g>.
33. John Lopez: Russell Brandom, Milagros Miceli, John Lopez, and Aiha Nguyen, “Generative AI’s Labor Impacts, Part One: Hierarchy,” *Data and Society*, January 18, 2024, <https://datasociety.net/events/generative-ais-labor-impacts-part-one/>.
34. more scarce for workers: Brian Merchant, “AI Is Already Taking Jobs in the Video Game Industry,” *Wired*, July 23, 2024, <https://www.wired.com/story/ai-is-already-taking-jobs-in-the-video-game-industry/>.
35. both companies ran at a loss: Uber did not report profitability until 2024. Andrew J. Hawkins, “Uber Ends the Year in the Black for the First Time Ever,” *The Verge*, February 8, 2024, <https://www.theverge.com/2024/2/8/24065999/uber-earnings-profitable-year-net-income>.
36. rideshare driver advocates: Personal communication, Nicole Moore, August 10, 2023.
37. \$100 million in misleading advertising: Sam Harnett, “Prop. 22 Explained: Why Gig Companies Are Spending Huge Money on an Unprecedented Measure,” *KQED*, October 26, 2020, <https://www.kqed.org/news/11843123/prop-22-explained-why-gig-companies-are-spending-huge-money-on-an-unprecedented-measure>.
38. make a living wage: For a demonstration of the wage crunch for gig drivers, see “The Uber Game,” *Financial Times*, October 2017, <https://ig.ft.com/uber-game/>. For workers’ perspectives on employee classification and contractor status, see Veena Dubal, “An Uber Ambivalence: Employee Status, Worker Perspectives, & Regulation in the Gig Economy,” UC Hastings Research Paper No. 381, November 15, 2019, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3488009.
39. Safe Street Rebel: “ConeSF: A Campaign to Rein In Robotaxis,” Safe Street Rebel, accessed September 11, 2024, <https://www.safestreetrebel.com/conesf/>.
40. Waymo car: Brian Merchant, “Torching the Google Car: Why the Growing Revolt Against Big Tech Just Escalated,” *Blood in the Machine*, February 14, 2024, <https://www.bloodinthemachine.com/p/torching-the-google-car-why-the-growing>.
41. U.S. labor law: Chloe Xiang, “Eating Disorder Helpline Fires Staff, Transitions to Chatbot After Unionization,” *Motherboard*, May 25, 2023, <https://www.vice.com/en/article/eating-disorder-helpline-fires-staff-transitions-to-chatbot-after-unionization/>.
42. “Tessa” was brought online: Kate Wells, “Can a Chatbot Help People with Eating Disorders as Well as Another Human?,” *NPR News*, May 24, 2023, <https://www.npr.org/2023/05/24/1177847298/can-a-chatbot-help-people-with-eating-disorders-as-well-as-another-human>.
43. Sharon Maxwell documented: Sharon Maxwell @heysharonmaxwell, “It is beyond time for NEDA to step aside . . .” Instagram, May 29, 2023, <https://www.instagram.com/p/Cs1jp1pPkOs/>.
44. Sara Ziff: U.S. Federal Trade Commission, “Generative Artificial Intelligence and the Creative Economy Staff Report: Perspectives and Takeaways,” December 2023, https://www.ftc.gov/system/files/ftc_gov/pdf/12-15-2023AICEStaffReport.pdf.

45. Levi's: Jess Weatherbed, "Levi's Will Test AI-Generated Clothing Models to 'Increase Diversity,'" *The Verge*, May 27, 2023, <https://www.theverge.com/2023/3/27/23658385/levis-ai-generated-clothing-model-diversity-denim>; "LS&Co. Partners with Lalaland.ai," Levi Strauss, March 28, 2023, <https://www.levistrauss.com/2023/03/22/lSCO-partners-with-lalaland-ai/>.
46. drive clicks to advertisements: Marc Stenberg, "Google Is Paying Publishers to Test an Unreleased Gen AI Platform," *AdWeek*, February 27, 2024, <https://www.adweek.com/media/google-paying-publishers-unreleased-gen-ai/>.
47. hardly be called autonomous: Tripp Mickle, Cade Metz, and Yiwen Lu, "G.M.'s Cruise Moved Fast in the Driverless Race. It Got Ugly," *New York Times*, November 3, 2023, <https://www.nytimes.com/2023/11/03/technology/cruise-general-motors-self-driving-cars.html>; Kyle Vogt, post to "Hacker News," Y Combinator Forum, November 2023, <https://news.ycombinator.com/item?id=38145062>.
48. mental health issues: Billy Perrigo, "Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic," *Time*, January 18, 2023, <https://time.com/6247678/openai-chatgpt-kenya-workers/>; Karen Hao, "The Hidden Workforce That Helped Filter Violence and Abuse Out of ChatGPT," *The Journal* (podcast), July 11, 2023, <https://www.wsj.com/podcasts/the-journal/the-hidden-workforce-that-helped-filter-violence-and-abuse-out-of-chatgpt/ffc2427f-bdd8-47b7-9a4b-27e7267cf413>; Karen Hao and Deepa Seetharaman, "Cleaning Up ChatGPT Takes Heavy Toll on Human Workers," *Wall Street Journal*, July 24, 2023, <https://www.wsj.com/articles/chatgpt-openai-content-abusive-sexually-explicit-harassment-kenya-workers-on-human-workers-cf191483>. Workers have told their own stories of mental health crises and working conditions within the Data Workers' Inquiry project, <https://data-workers.org>.
49. those who perform it: Mary L. Gray and Siddarth Suri, *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass* (Boston and New York: Houghton Mifflin Harcourt, 2019). But also see Noopur Raval, "Interrupting Invisibility in a Global World," *Interactions* 28, no. 4 (2021): 27–31, <https://doi.org/10.1145/3469257>.
50. deep learning methods has been: Karën Fort, Gilles Adda, and K. Bretonnel Cohen, "Last Words: Amazon Mechanical Turk: Gold Mine or Coal Mine?," *Computational Linguistics* 37, no. 2 (June 2011): 413–20, https://doi.org/10.1162/COLI_a_00057.
51. Mechanical Turk: Tom Standage, *The Turk: The Life and Times of the Famous Eighteenth-Century Chess-Playing Machine* (New York: Berkley Books, 2002).
52. according to Li: Dave Gershgor, "The Data That Transformed AI Research—and Possibly the World," *Quartz*, July 26, 2017, <https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world>; Fei-Fei Li, "How We're Teaching Computers to Understand Pictures," filmed March 2015 in Vancouver, British Columbia. TED video, 17:48, https://www.ted.com/talks/fei_fei_li_how_we_re_teaching_computers_to_understand_pictures/.
53. One employee remarked: Karen Hao, "The Hidden Workforce."
54. marketed itself: Perrigo, "Exclusive: OpenAI Used Kenyan Workers."
55. "system card": OpenAI, "GPT-4 System Card," March 23, 2023, <https://cdn.openai.com/papers/gpt-4-system-card.pdf>.
56. unilaterally shut down access: Karen Hao (@_KarenHao), "Update: This is worse than I thought. Workers in Nigeria, Rwanda, and South Africa are also reporting that they are being blocked from the website or cannot create an account," Twitter, March 18, 2024, https://twitter.com/_KarenHao/status/1769664691658428539; Russell Brandom, "Scale AI's Remotasks Platform Is Dropping Whole Countries Without Explanation," *Rest of World*, March 28, 2024, <https://restofworld.org/2024/scale-ai-remotasks-banned-workers/>; Mophat Okinyi, "Impact of Remotasks' Closure on Kenyan Workers," Data Workers Inquiry, 2024, <https://data-workers.org/mophat/>.

57. multiple suspensions: Turkopticon (@turkopticon), “UPDATE: We are receiving reports of Mturk worker accounts being reinstated!” Twitter, March 1, 2024, <https://twitter.com/turkopticon/status/1763627628898959586>; Emanuel Maiberg, “Amazon Turkers Who Train AI Say They’re Locked Out of Their Work and Money,” 404 Media, February 29, 2024, <https://www.404media.co/amazon-turkers-who-train-ai-say-theyre-locked-out-of-their-work-and-money/>; Turkopticon (@turkopticon), “Another day with more workers suspended, no response from @amazonmturk or @amazonpay other than blaming the worker or telling them to violate TOS by making a new account. Do you plan to let this go until #Mturk has no US workers? #AcademicTwitter,” Twitter, August 30, 2024, <https://twitter.com/turkopticon/status/1829504606952407150>.
58. Content moderators have requested: Sarah T. Roberts, *Behind the Screen: Content Moderation in the Shadows of Social Media* (New Haven, CT, and London: Yale University Press, 2019); Sarah T. Roberts on “The Real Facebook Oversight Board,” YouTube video, 01:08:10, October 26, 2020, https://www.youtube.com/watch?v=F1byT_2htfs.
59. Karen Hao and Andrea Paola Hernández have written: Karen Hao and Andrea Paola Hernández, “How the AI Industry Profits from Catastrophe,” *MIT Technology Review*, April 20, 2022, <https://www.technologyreview.com/2022/04/20/1050392/ai-industry-appen-scale-data-labels/>.
60. exposed to traumatic content: Niamh Row, “Underage Workers Are Training AI,” *Wired*, November 15, 2023, <https://www.wired.com/story/artificial-intelligence-data-labeling-children/>.
61. even prisoners: Morgan Meaker, “These Prisoners Are Training AI,” *Wired*, September 11, 2023, <https://www.wired.com/story/prisoners-training-ai-finland/>.
62. studios can’t require: Annabelle Timsit, “Hollywood Studios and Writers Have a Strike-Ending Deal. What’s in It?,” *Washington Post*, September 27, 2023, <https://www.washingtonpost.com/style/2023/09/27/wga-contract-details-writers-strike-deal/>.
63. The actors soon followed: SAG-AFTRA, “Theatrical (Basic or CBA),” accessed September 14, 2024, <https://www.sagaftra.org/theatrical-basic-or-cba>.
64. would-be AI model creates: Shawn Shan et al., “About the Glaze Project,” accessed September 2, 2024, <https://glaze.cs.uchicago.edu/aboutus.html>.
65. Daniel Motaung and 150 workers convened: Billy Perrigo, “150 African Workers for ChatGPT, TikTok and Facebook Vote to Unionize at Landmark Nairobi Meeting,” *Time*, May 1, 2023, <https://time.com/6275995/chatgpt-facebook-african-workers-union/>.
66. Turkopticon: Lilly C. Irani and M. Six Silberman, “Turkopticon: Interrupting Worker Invisibility in Amazon Mechanical Turk,” *CHI ’13: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Paris, France, April 27, 2013, <https://doi.org/10.1145/2470654.2470742>.
67. They have fought: Turkopticon, “Mass Rejections,” December 21, 2020, <https://blog.turkopticon.net/?p=731>.
68. claims that ChatGPT could replace crowdworkers: Turkopticon, “Beware the Hype: ChatGPT Didn’t Replace Human Data Annotators,” *Tech Worker Coalition Newsletter*, April 4, 2023, <https://news.techworkerscoalition.org/2023/04/04/issue-5/>.

Chapter 4: If It Quacks like a Doc: AI Hype and Social Services

1. According to Motherboard: Chloe Xiang, “‘He Would Still Be Here’: Man Dies by Suicide After Talking with AI Chatbot, Widow Says,” *Vice*, March 30, 2023, <https://www.vice.com/en/article/pkadgm/man-dies-by-suicide-after-talking-with-ai-chatbot-widow-says>. Original reporting by Pierre-François Lovens, “Sans ces conversations avec le chatbot Eliza, mon mari serait toujours là,” *La Libre*, March 28, 2023,

<https://www.lalibre.be/belgique/societe/2023/03/28/sans-ces-conversations-avec-le-chatbot-eliza-mon-mari-serait-toujours-la-LVSLWPC5WRDX7J2RCHNWPDST24/>.

2. bragged proudly: Rob Morris (@RobertRMorris), “We provided mental health support to about 4,000 people—using GPT-3. Here’s what happened 🙌,” Twitter, January 6, 2023, <https://x.com/RobertRMorris/status/1611450197707464706>.
3. Virginia Eubanks documents: Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (New York: St. Martin’s Press, 2018).
4. and poor Black and Indigenous: Following McMillan-Major 2023 and Smith 2012, we use the term “indigenous” to refer to the peoples around the world who have been subjected to colonialism. When referring to people from North America, we follow local convention in capitalizing “Indigenous”, reflecting not only racialization of Indigenous North Americans but also their political work. See Angelina McMillan-Major, “Language Dataset Documentation Design: Learning from Deaf and Indigenous Communities” (PhD diss., University of Washington, 2023), <https://digital.lib.washington.edu/researchworks/handle/1773/50854>; Linda Tuhiwai Smith, *Decolonizing Methodologies: Research and Indigenous Peoples* (London and New York: Zed Books, 1999).
5. As Eubanks puts it: Eubanks, *Automating Inequality*, 10.
6. Allegheny Family Screening Tool (AFST): Eubanks, *Automating Inequality*, 127–73. Many have written about the technical challenges around developing a “fair” AFST. See, for instance, Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianatha, “A Case Study of Algorithm-Assisted Decision Making in Child Maltreatment Hotline Screening Decisions,” *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, PMLR 81 (2018): 134–48.
7. the system’s racism: Dorothy E. Roberts, *Torn Apart: How the Child Welfare System Destroys Black Families and How Abolition Can Build a Safer World* (New York: Basic Books, 2022).
8. supercharge and scale it: J. K. Abdurahman, “Calculating the Souls of Black Folk: Predictive Analytics in the New York City Administration for Children’s Services,” *Columbia Journal of Race and Law* 11, no. 4 (2012): 75–110, <https://doi.org/10.52214/cjrl.v11i4>. Furthermore, in Allegheny County, Eubanks notes, Black families are 11 percent of the population, but 70 percent of the children in foster care are Black.
9. detained in local jails: Wendy Sawyer and Peter Wagner, “Mass Incarceration: The Whole Pie 2024,” Prison Policy Initiative, March 14, 2024, <https://www.prisonpolicy.org/reports/pie2024.html>.
10. take advantage of this: Allie Preston and Rachael Eisenberg, “Profit Over People: The Commercial Bail Industry Fueling America’s Cash Bail Systems,” Center for American Progress, July 6, 2022, <https://www.americanprogress.org/article/profit-over-people-primer-on-u-s-cash-bail-systems/>.
11. including Arizona, Kentucky, New Jersey, and Utah: Nigel Duara, “Replacing Cash Bail: Fairer Justice or Robocalypse?,” *Cal Matters*, October 22, 2020, <https://calmatters.org/justice/2020/10/cash-bail-justice-algorithm-risk-assessment-prop-25/>. As of 2024, most states have at least one county that uses pretrial risk assessment systems. Mapping Pretrial Injustice, “Where Are Risk Assessments Being Used?,” accessed September 14, 2024, <https://pretrialrisk.com/national-landscape/where-are-prai-being-used/>.
12. Julia Angwin and her colleagues: Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner, “Machine Bias,” ProPublica, May 23, 2016, <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
13. set off a broad debate: For instance, for the technical debate, see Sam Corbett-Davies, Emma Pierson, Avi Feller, and Sharad Goel, “A Computer Program Used for Bail and Sentencing Decisions Was Labeled Biased Against Blacks. It’s Actually Not That Clear,” *Washington Post*, October 17, 2016, <https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an->

algorithm-be-racist-our-analysis-is-more-cautious-than-propublicas/. For a nontechnical discussion and focus on the different ways that bias can creep into judicial decisions, see Laurel Eckhouse, Kristian Lum, Cynthia Conti-Cook, and Julie Ciccolini, “Layers of Bias: A Unified Approach for Understanding Problems with Risk Assessment,” *Criminal Justice and Behavior* 46, no. 2 (2018): 185–209.

14. New York City Police Department robot: Jeffrey C. Mays, “400-Pound N.Y.P.D. Robot Gets Tryout in Times Square Subway Station,” *New York Times*, September 22, 2023, <https://www.nytimes.com/2023/09/22/nyregion/police-robot-times-square-nyc.html>; Dana Rubinstein and Hurubie Meko, “Goodbye for Now to the Robot That (Sort Of) Patrolled New York’s Subway,” *New York Times*, February 2, 2024, <https://www.nytimes.com/2024/02/02/nyregion/nypd-subway-robot-retires.html>.
15. “AI Action Plan”: “Mayor Adams Releases First-of-Its-Kind Plan for Responsible Artificial Intelligence Use in NYC Government,” NYC.gov, October 16, 2023, <https://www.nyc.gov/office-of-the-mayor/news/777-23/mayor-adams-releases-first-of-its-kind-plan-responsible-artificial-intelligence-use-nyc#0>.
16. The City revealed: Colin Lecher, “NYC’s AI Chatbot Tells Businesses to Break the Law,” *The Markup*, March 29, 2024, <https://themarkup.org/news/2024/03/29/nycs-ai-chatbot-tells-businesses-to-break-the-law>.
17. Gavin Newsom signing an executive order: “Governor Newsom Signs Executive Order to Prepare California for the Progress of Artificial Intelligence,” Governor Gavin Newsom, September 6, 2023, <https://www.gov.ca.gov/2023/09/06/governor-newsom-signs-executive-order-to-prepare-california-for-the-progress-of-artificial-intelligence/>.
18. California Government Operations Agency report suggests: California Government Operations Agency, “State of California Benefits and Risks of Generative Artificial Intelligence Report,” November 2023, https://www.govops.ca.gov/wp-content/uploads/sites/11/2023/11/GenAI-EO-1-Report_FINAL.pdf.
19. journalist Khari Johnson: Khari Johnson, “California Plans to Use AI to Answer Your Tax Questions,” *CalMatters*, February 8, 2024, <https://calmatters.org/economy/technology/2024/02/cdtfa-generative-ai/>.
20. to address homelessness: “Governor Newsom Seeks to Harness the Power of GenAI to Address Homelessness, Other Challenges,” Governor Gavin Newsom, September 5, 2024, <https://www.gov.ca.gov/2024/09/05/governor-newsom-seeks-to-harness-the-power-of-genai-to-address-homelessness-other-challenges/>.
21. Respond Crisis Translation: Johana Bhuiyan, “Lost in AI Translation: Growing Reliance on Language Apps Jeopardizes Some Asylum Applications,” *Guardian*, September 7, 2023, <https://www.theguardian.com/us-news/2023/sep/07/asylum-seekers-ai-translation-apps>.
22. held at Bletchley Park: Rishi Sunak, “Prime Minister’s Speech at the AI Safety Summit: 2 November 2023,” speech, Bletchley Park, UK, November 2, 2023, <https://www.gov.uk/government/speeches/prime-ministers-speech-at-the-ai-safety-summit-2-november-2023>.
23. masses of students marched: Daan Kolkman, “‘F**k the algorithm?’”.
24. Oliver Dowden announced plans: Lucy Fisher, “UK Government to Trial ‘Red Box’ AI Tools to Improve Ministerial Efficiency,” *Financial Times*, February 28, 2024, <https://www.ft.com/content/f2ae55bf-b9fa-49b5-ac0e-8b7411729539>.
25. National Health Service: “Robot Receptionists to Make NHS ‘Fit for the Future,’” *Telegraph*, June 26, 2023, <https://www.telegraph.co.uk/news/2023/06/26/ai-admin-nhs-robots-rishi-sunak-staff-overhaul/>.
26. it had invested: Department of Health and Social Care, Steve Barclay and Chloe Smith, “£21 Million to Roll Out Artificial Intelligence Across the NHS,” June 23, 2023,

- <https://www.gov.uk/government/news/21-million-to-roll-out-artificial-intelligence-across-the-nhs>.
27. Health care researchers have remarked: Katharine Da Costa, “AI in Healthcare: What Are the Risks for the NHS?,” BBC News, August 6, 2024, <https://www.bbc.com/news/articles/c6233x9k4dlo>.
 28. Just Treatment found: Connected by Data and Just Treatment, “Our Data Stories: Health,” November 22, 2023, <https://connectedbydata.org/resources/our-health-data-stories>.
 29. judge used ChatGPT: Thomas Germain, “Judges Given the OK to Use ChatGPT in Legal Rulings,” Gizmodo, December 12, 2023, <https://gizmodo.com/uk-judges-now-permitted-use-chatgpt-in-legal-rulings-1851093046>.
 30. gave judges the okay: Courts and Tribunals Judiciary, “Artificial Intelligence (AI): Guidance for Judicial Office Holders,” December 12, 2023, <https://www.judiciary.uk/wp-content/uploads/2023/12/AI-Judicial-Guidance.pdf>.
 31. Michael Cohen: Debra Cassens Weiss, “Fake Citations in Legal Brief Were Generated by Google Bard AI Program, Says Ex-Trump Lawyer Michael Cohen,” *ABA Journal*, January 2, 2024, <https://www.abajournal.com/news/article/fake-citations-in-legal-brief-were-generated-by-google-bard-ai-program-says-ex-trump-lawyer-michael-cohen>.
 32. DoNotPay.com in 2015: Bobby Allyn, “A Robot Was Scheduled to Argue in Court, Then Came the Jail Threats,” NPR, January 25, 2023, <https://www.npr.org/2023/01/25/1151435033/a-robot-was-scheduled-to-argue-in-court-then-came-the-jail-threats>.
 33. system was originally created: Shannon Liao, “‘World’s First Robot Lawyer’ Now Available in All 50 States,” *The Verge*, July 20, 2017, <https://www.theverge.com/2017/7/12/15960080/chatbot-ai-legal-donotpay-us-uk>.
 34. someone had actually signed up: Emma Roth, “DoNotPay Chickens Out on Its Courtroom AI Chatbot Stunt,” *The Verge*, January 25, 2023, <https://www.theverge.com/2023/1/25/23571192/donotpay-robot-lawyer-courtroom>.
 35. Browder had even offered: Allyn, “A Robot Was Scheduled.”
 36. Juan David Gutiérrez: Juan David Gutiérrez, “Judges and Magistrates in Peru and Mexico Have ChatGPT Fever,” *Tech Policy Press*, April 19, 2023, <https://www.techpolicy.press/judges-and-magistrates-in-peru-and-mexico-have-chatgpt-fever/>.
 37. pertinent to the case at hand: In one case, the judge gave the URL for ChatGPT in a footnote, apparently unaware that ChatGPT is not a stable source of information: even if you were to prompt ChatGPT with the same exact wording, there is no reason to expect the same output.
 38. House of Representatives distributed: Emily Birnbaum and Laura Davison, “AI Is Making Politics Easier, Cheaper and More Dangerous,” *Bloomberg*, July 11, 2023, <https://www.bloomberg.com/news/features/2023-07-11/chatgpt-ai-boom-makes-political-dirty-tricks-easier-and-cheaper>.
 39. Congressman Ro Khanna had used it: Shira Stein, “ChatGPT Wrote California Rep. Ro Khanna’s New AI Bill,” *San Francisco Chronicle*, July 20, 2023, <https://www.sfchronicle.com/politics/article/chatgpt-ai-federal-bill-18200932.php>; U.S. Congress, House, *Streamlining Effective Access and Retrieval of Content to Help Act of 2023*, HR 4793, 118th Cong., 1st sess., introduced in House July 20, 2023, <https://www.congress.gov/bill/118th-congress/house-bill/4793/all-actions>.
 40. city council passed a bill: María Luisa Paúl, “A Brazilian City Passed a Law About Water Meters. ChatGPT Wrote It,” *Washington Post*, December 4, 2023, <https://www.washingtonpost.com/nation/2023/12/04/ai-written-law-porto-alegre-brazil/>; Ramiro Rosário (@curtaramiro), “A primeira lei brasileira feita exclusivamente por inteligência artificial está vigente em Porto Alegre!” Twitter, November 29, 2023, <https://twitter.com/curtaramiro/status/1729805545157120113>.

41. only after it passed: Councilman Rosário justified hiding the origins of the bill until after it was enacted because he thought people might “hold a prejudice against AI.” In this rhetorical move, we detect another kind of AI hype: co-opting the language of social justice to frame AI systems as something that can be discriminated against. This ascribes human characteristics to the systems and elevates them to the kind of thing that might have rights or might suffer from prejudice (and eventually oppression). The move enlists people’s sympathy while at the same time burnishing the impression of the systems being more than they are.
42. these use cases: Kendra Albert, interview with Emily M. Bender and Alex Hanna, “Episode 10: Don’t Be a Lawyer, ChatGPT,” *Mystery AI Hype Theater 3000* (podcast), March 3, 2023, <https://www.buzzsprout.com/2126417/13414023>.
43. history of benchmarking in computing: Byron C. Lewis and Albert E. Crews, “The Evolution of Benchmarking as a Computer Performance Evaluation Technique,” *MIS Quarterly* 9, no. 1 (March 1985): 7–16, <https://doi.org/10.2307/249270>.
44. Some other early benchmarks: Mark Liberman, “Obituary: Fred Jelinek,” *Computational Linguistics* 36, no. 4 (December 2010): 595–99; Kenneth Ward Church, “Emerging Trends: A Tribute to Charles Wayne,” *Natural Language Engineering* 24, no. 1 (January 2018): 155–60, <https://doi.org/10.1017/S1351324917000389>.
45. an evaluation campaign for automatic transcription: John J. Godfrey and Edward Holliman, *The Switchboard-1 Telephone Speech Corpus* (1993), distributed by Linguistic Data Consortium, <https://catalog.ldc.upenn.edu/LDC97S62>; Sadaoki Furui, “Generalization Problem in ASR Acoustic Model Training and Adaptation,” *Proceedings of the 2009 IEEE Workshop on Automatic Speech Recognition and Understanding, ASRU 2009* (January 2010), <https://ieeexplore.ieee.org/document/5373493>.
46. often called the “ground truth”: This term comes into computer science ultimately from military uses comparing remote sensing to the full information available “on the ground”, via its use in the science and engineering of remote sensing. Geoscientist Iain Woodhouse traces the history of the term and also argues that it’s time to retire it. See Iain H. Woodhouse, “On ‘Ground’ Truth and Why We Should Abandon the Term,” *Journal of Applied Remote Sensing* 15, no. 4 (2021), <https://www.spiedigitallibrary.org/journals/journal-of-applied-remote-sensing/volume-15/issue-4/041501/On-ground-truth-and-why-we-should-abandon-the-term/10.1117/1.JRS.15.041501.full>.
47. results of the evaluation: There is a fairly large literature on evaluation in machine learning. For lessons from social sciences that discuss their social impact, see Abigail Z. Jacobs and Hanna Wallach, “Measurement and Fairness,” *FACCT ’21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (March 2021): 375–85, <https://doi.org/10.1145/3442188.3445901>.
48. choice of evaluation data: The design of voice-based automated systems is particularly exclusive for people who stutter. This is impactfully documented in this video op-ed: “I Stutter. But I Need You to Listen,” *New York Times*, YouTube video, 08:13, August 23, 2022, https://www.youtube.com/watch?v=m0E_wMIwfSI. See also Shaomei Wu, “Blocked by the System: How Current Voice AI Silences People Who Stutter,” paper presented at American Association for Advancement in Science Annual Meeting, Denver, CO, February 16, 2024, <https://aaas.confex.com/aaas/2024/meetingapp.cgi/Session/31684>; Qisheng Li and Shaomei Wu, “Towards Fair and Inclusive Speech Recognition for Stuttering: Community-Led Chinese Stuttered Speech Dataset Creation and Benchmarking,” *CHI EA ’24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (May 11, 2024): 1–9, <https://dl.acm.org/doi/10.1145/3613905.3650950>.
49. other varieties of the same language: Alicia Beckford Wassink, Cady Gansen, and Isabel Bartholomew, “Uneven Success: Automatic Speech Recognition and Ethnicity-Related

- Dialects,” *Speech Communication* 140 (May 2022): 50–70, <https://www.sciencedirect.com/science/article/abs/pii/S0167639322000486>.
50. on a fixed evaluation task: Raji et al., “AI and the Everything in the Whole Wide World Benchmark.”
 51. Clever Hans effect: Benjamin Heinzerling, “NLP’s Clever Hans Moment Has Arrived,” *The Gradient*, August 26, 2019, <https://thegradient.pub/nlps-clever-hans-moment-has-arrived/>.
 52. Clever Hans was a horse: The *New York Times* reported breathlessly on Clever Hans in 1904: “Berlin’s Wonderful Horse,” *New York Times*, September 4, 1904, <https://timesmachine.nytimes.com/timesmachine/1904/09/04/101396572.pdf>. Biologist and psychologist Oscar Pfungst worked out that Clever Hans was in fact reading signals in the expression of the person asking questions: Laasya Samhita and Hans J. Gross, “The ‘Clever Hans Phenomenon’ Revisited,” *Communicative & Integrative Biology* 6, no. 6 (2013), <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3921203/>.
 53. for the “wrong” reason: Similarly, natural language processing researchers have shown that language models can “succeed” at tasks intended to test for “understanding” by picking up on irrelevant patterns. For example, one test (called Stanford Natural Language Inference, or SNLI) was designed to test whether a computer could tell if two sentences were consistent with each other—what linguists call *entailment*—or in contradiction. For example, the sentence *An older man is drinking orange juice at a restaurant* entails and thus is consistent with *A man is drinking juice*, whereas *Two women are at a restaurant drinking wine* is a contradictory description of the same scene. Later analyses by other researchers found a curious pattern with these data: the systems they tested could do better than chance at classifying the second sentences as contradictory or not, without being shown the sentence they were supposedly contradicting! Looking more closely, they found that the systems were picking up on facts like if the second one had negation in it or, oddly, the word *cat*, among other clues, the answer was probably “contradiction”. Artifacts come about because of how the datasets are constructed—in the case of SNLI, by hiring crowd workers to rapidly create sentences that contradict or (don’t) a scene in a photograph. Workers doing this as quickly as possible often reached for negation as a way to make something contradictory. And why *cat*? Because in a lot of cases, the photographs used as prompts were pictures of dogs. See Suchin Gururangan et al., “Annotation Artifacts in Natural Language Inference Data,” *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2 (New Orleans: Association for Computational Linguistics, 2018), 107–12.
 54. “ChatGPT Passes Bar Exam”: A sampling of headlines from January–April 2023: Samantha Murphy Kelly, “ChatGPT Passes Exams from Law and Business Schools,” CNN, January 26, 2023, <https://edition.cnn.com/2023/01/26/tech/chatgpt-passes-exams/index.html>; John Koetsier, “GPT-4 Beats 90% of Lawyers Trying to Pass the Bar,” *Forbes*, March 14, 2023, <https://www.forbes.com/sites/johnkoetsier/2023/03/14/gpt-4-beats-90-of-lawyers-trying-to-pass-the-bar/>; Pablo Arredondo, “GPT-4 Passes the Bar Exam: What That Means for Artificial Intelligence Tools in the Legal Profession,” Stanford Law School (blog), April 19, 2023, <https://law.stanford.edu/2023/04/19/gpt-4-passes-the-bar-exam-what-that-means-for-artificial-intelligence-tools-in-the-legal-industry/>.
 55. licensing processes: There is a large sociological literature on professionalization and to whose benefit licensure accrues. See Andrew Delano Abbott, *The System of Professions: An Essay on the Division of Expert Labor* (Chicago: University of Chicago Press, 1988). For more recent treatments: Andreas Haupt, “Who Profits from Occupational Licensing?,” *American Sociological Review* 88, no. 6 (December 2023): 1104–30, <https://doi.org/10.1177/00031224231207395>; Beth Redbird, “The New Closed Shop? The Economic and Structural Effects of Occupational Licensure,” *American Sociological Review* 82, no. 3 (June 2017): 600–24, <https://doi.org/10.1177/0003122417706463>.

56. one of the largest in the world: Aiden Lee, Joel Ruhter, Christie Peters, Nancy De Lew, and Benjamin D. Sommers, “National Uninsured Rate Reaches All-Time Low in Early 2022,” Assistant Secretary for Planning and Evaluation, U.S. Department of Health and Human Services, August 2022, <https://aspe.hhs.gov/sites/default/files/documents/a35a060182f78d32ceff32be71301173/Uninsured-Q1-2022-Data-Point.pdf>. The U.S. health care divide is still extreme, with huge gaps in who has health insurance and access to care between rich and poor, further accentuated by differences across race and ethnicity and immigration status. Meanwhile, per capita expenditure of health care in the U.S. is by far the highest in the world (thanks in part to having for-profit insurance companies), with an average life expectancy that is lowest of all OECD nations. Max Roser, “Why Is Life Expectancy in the US Lower than in Other Rich Countries?,” Our World in Data, October 29, 2020, <https://ourworldindata.org/us-life-expectancy-low>.
57. reduce health care expenditure: Melinda Beeuwkes Buntin and John A. Graves, “How the ACA Dented the Cost Curve: An Analysis of Whether or Not the Affordable Care Act Reduced the Annual Rate at Which Total National Health Care Costs Increased and Brought Per Capita Health Spending Growth Rates Down,” *Health Affairs* 39, no. 3 (March 2020): 403–12, <https://doi.org/10.1377/hlthaff.2019.0147>.
58. *Journal of the American Medical Association*: Andrew Wong et al., “External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients,” *JAMA Internal Medicine* 181, no. 8 (2021): 1065–70, <https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/2781307>.
59. National Nurses United: “A.I.’s Impact on Nursing and Health Care,” National Nurses United, accessed July 11, 2024, <https://www.nationalnursesunited.org/artificial-intelligence>.
60. well-publicized 2019 study: Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan, “Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations,” *Science* 366, no. 6464 (2019): 447–53, <https://www.science.org/doi/10.1126/science.aax2342>.
61. class-action lawsuit: The Estate of Gene B. Lokken and the Estate of Dale Henry Tetzloff, individually and on behalf of all others similarly situated, Plaintiffs, v. UnitedHealth Group, Inc., UnitedHealthcare, Inc., Navihealth, INC., and Does 1–51-50, inclusive, (D. Minn., 2023), <https://www.documentcloud.org/documents/24166450-class-action-v-unitedhealth-and-navihealth>.
62. *Stat News* largely confirms: Casey Ross and Bob Herman, “Denied by AI: How Medicare Advantage Plans Use Algorithms to Cut Off Care for Seniors in Need,” *Stat News*, March 13, 2023, <https://www.statnews.com/2023/03/13/medicare-advantage-plans-denial-artificial-intelligence/>; Casey Ross and Bob Herman, “UnitedHealth Pushed Employees to Follow an Algorithm to Cut Off Medicare Patients’ Rehab Care,” *Stat News*, November 14, 2023, <https://www.statnews.com/2023/11/14/unitedhealth-algorithm-medicare-advantage-investigation/>. See also Beth Mole, “UnitedHealth Uses AI Model with 90% Error Rate to Deny Care, Lawsuit Alleges,” *Ars Technica*, November 16, 2023, <https://arstechnica.com/health/2023/11/ai-with-90-error-rate-forces-elderly-out-of-rehab-nursing-homes-suit-claims/>.
63. McKinsey estimates: Shashank Bhasker, Damien Bruce, Jessica Lamb, and George Stein, “Tackling Healthcare’s Biggest Burdens with Generative AI,” McKinsey & Company, July 10, 2023, https://www.mckinsey.com/industries/healthcare/our-insights/tackling-healthcares-biggest-burdens-with-generative-ai#.
64. published in *Nature*: Karan Singhal et al., “Large Language Models Encode Clinical Knowledge,” *Nature*, 620, no. 7972 (2023): 172–80, <https://doi.org/10.1038/s41586-023-06291-2>.

65. going on a press tour: Miles Kruppa and Nidhi Subbaraman, “In Battle with Microsoft, Google Bets on Medical AI Program to Crack Healthcare Industry,” *Wall Street Journal*, July 8, 2023, https://www.wsj.com/articles/in-battle-with-microsoft-google-bets-on-medical-ai-program-to-crack-healthcare-industry-bb7c2db8?mod=Searchresults_pos1&page=1.
66. The company: Dereck Paul, “Glass Health 2023 Year in Review,” *Glass Health Blog*, January 1, 2024, <https://blog.glass.health/glass-health-2023-year-in-review/>.
67. exploring using chatbots: Frank Landymore, “Amazon-Owned Health Clinics Reportedly Training AI Chatbots to Triage Patients,” *Futurism*, March 1, 2024, <https://futurism.com/the-byte/amazon-health-clinics-ai-triage-patients>.
68. \$120 million of seed funding: This included \$50 million from Andreessen Horowitz and General Catalyst. <https://www.bloomberg.com/news/articles/2023-05-16/andreessen-general-catalyst-back-health-startup-hippocratic-ai>; <https://www.techopedia.com/news/hippocratic-ai-raises-53m-in-funding-valuation-soars-to-500m>.
69. on their website: “Hippocratic AI,” Hippocratic AI, <https://www.hippocraticai.com/>.
70. confidently boasts: Paul Webster, “Six Ways Large Language Models Are Changing Healthcare,” *Nature Medicine* 29 (November 30, 2023): 2969–71, <https://doi.org/10.1038/s41591-023-02700-1>.
71. the CEO states: Kyle Wiggers, “Hippocratic Is Building a Large Language Model for Healthcare,” *TechCrunch*, May 16, 2023, <https://techcrunch.com/2023/05/16/hippocratic-is-building-a-large-language-model-for-healthcare/>.
72. wasn’t trying: Weizenbaum, *Computer Power*, 2–8; Tarnoff, “Weizenbaum’s Nightmares.”
73. Sutskever tweeted excitedly: Ilya Sutskever (@ilyasut), “In the future, once the robustness of our models will exceed some threshold, we will have *wildly effective* and dirt cheap AI therapy,” Twitter, September 27, 2023, <https://x.com/ilyasut/status/1707027536150929689>.
74. have secured hundreds: “Woebot Health Closes \$90 Million Series B Funding Round Co-Led by JAZZ Venture Partners and Temasek,” Woebot Health, July 21, 2021, <https://woebothealth.com/woebot-health-closes-90-million-series-b-funding/>; “Pyx Health,” Crunchbase, accessed July 15, 2024, https://www.crunchbase.com/organization/pyx-health/company_financials.
75. founders and defenders argue that: Yuki Noguchi, “Therapy by Chatbot? The Promise and Challenges in Using AI for Mental Health,” NPR, January 19, 2023, <https://www.npr.org/sections/health-shots/2023/01/19/1147081115/therapy-by-chatbot-the-promise-and-challenges-in-using-ai-for-mental-health>.
76. no such requirements: Elizabeth Svoboda and *OpenMind* magazine, “When Your Psychologist Is an AI,” *Scientific American*, December 4, 2024, <https://www.scientificamerican.com/article/ai-therapy-bots-have-risks-and-benefits-and-more-risks/>.
77. Conflicts of interest abound: See, for instance, Woebot-authored research reports: “New Woebot Health Study Highlights AI’s Ability to Address Equity and Accessibility Gaps in Mental Health Care,” Woebot Health, August 28, 2023, <https://woebothealth.com/ais-ability-to-address-gaps/>. A report developed by Stanford’s Human-Centered AI lab on LLMs in therapy, moreover, warns that “[it] will behoove the field to be wary of attempts to optimize clinical LLMs on outcomes that have an explicit relationship with a company’s profit.” Elizabeth C. Stade et al., “Large Language Models Could Change the Future of Behavioral Healthcare: A Proposal for Responsible Development and Evaluation,” *Npj Mental Health Research* 3, no. 1 (April 2, 2024): 12, <https://doi.org/10.1038/s44184-024-00056-z>.
78. reproduced racist myths: Jesutofunmi A. Omiye, Jenna C. Lester, Simon Spichak, Veronica Rotemberg, and Roxana Daneshjou, “Large Language Models Propagate Race-Based Medicine,” *Npj Digital Medicine* 6 (October 20, 2023), <https://www.nature.com/articles/s41746-023-00939-z>.

79. kidney function: This includes a correction by race that has only recently been abandoned as a standard practice after significant efforts by health equity advocates and researchers, “eGFR Test Change: Removal of Race from the Calculation,” American Kidney Fund, accessed July 15, 2024, <https://www.kidneyfund.org/all-about-kidneys/tests/egfr/egfr-test-change-removal-race-calculation>.
80. evaluate insurance claims: Daniël van Dam and Raymond van Es, “The Potential of Large Language Models in the Insurance Sector,” Milliman, February 7, 2024, <https://www.milliman.com/en/insight/potential-of-large-language-models-insurance-sector>.
81. an opinion article: Most-read articles in the *Chronicle of Higher Education* are Claire Wallace and Fernanda Zamudio-Suarez, “2023: The Year in Higher Ed,” December 20, 2023, <https://www.chronicle.com/article/2023-the-year-in-higher-ed>, ChatGPT article: Owen Kichizo Terry, “I’m a Student. You Have No Idea How Much We’re Using ChatGPT,” May 12, 2023, <https://www.chronicle.com/article/im-a-student-you-have-no-idea-how-much-were-using-chatgpt>.
82. celebrated and embraced it: Examples of this abound, but one prime example is the 2023 TED Talk by Sal Kahn of Khan Academy: “Sal Khan’s 2023 TED Talk: AI in the Classroom Can Transform Education,” Khan Academy (blog), May 1, 2023, <https://blog.khanacademy.org/sal-khans-2023-ted-talk-ai-in-the-classroom-can-transform-education/>.
83. Denise Pope and Victor Lee suggests: Carrie Spector, “What Do AI Chatbots Really Mean for Students and Cheating?,” Stanford Graduate School of Education, October 31, 2023, <https://ed.stanford.edu/news/what-do-ai-chatbots-really-mean-students-and-cheating>; Natasha Singer, “Cheating Fears Over Chatbots Were Overblown, New Research Suggests,” *New York Times*, December 13, 2023, <https://www.nytimes.com/2023/12/13/technology/chatbot-cheating-schools-students.html>.
84. Pew Research Center: Olivia Sidoti and Jeffrey Gottfried, “About 1 in 5 U.S. Teens Who’ve Heard of ChatGPT Have Used It for Schoolwork,” Pew Research Center, November 16, 2023, <https://www.pewresearch.org/short-reads/2023/11/16/about-1-in-5-us-teens-whove-heard-of-chatgpt-have-used-it-for-schoolwork/>.
85. AI detection tool: “AI Writing Detection Capabilities: Frequently Asked Questions,” Turnitin, accessed July 15, 2024, <https://www.turnitin.com/products/features/ai-writing-detection/>.
86. OpenAI itself has admitted: “How Can Educators Respond to Students Presenting AI-Generated Content as Their Own?,” OpenAI, accessed July 15, 2024, <https://help.openai.com/en/articles/8313351-how-can-educators-respond-to-students-presenting-ai-generated-content-as-their-own>.
87. Texas A&M professor: Klee, “Professor Flunks All His Students.”
88. Center for Democracy and Technology found: Maddy Dwyer and Elizabeth Laird, “Report—Up in the Air: Educators Juggling the Potential of Generative AI with Detection, Discipline, and Distrust,” Center for Democracy and Technology, March 27, 2024, <https://cdt.org/insights/report-up-in-the-air-educators-juggling-the-potential-of-generative-ai-with-detection-discipline-and-distrust/>; Kristin Woelfel, “Brief—Late Applications: Disproportionate Effects of Generative AI-Detectors on English Learners,” Center for Democracy and Technology, December 18, 2023, <https://cdt.org/insights/brief-late-applications-disproportionate-effects-of-generative-ai-detectors-on-english-learners/>.
89. used for surveillance: Matthew H. Rafalow, *How Schools Create Inequality in the Tech Era* (Chicago: University of Chicago Press, 2020); Matthew H. Rafalow and Cassidy Puckett, “Sorting Machines: Digital Technology and Categorical Inequality in Education,” *Educational Researcher* 51, no. 4 (2022): 274–78, <https://doi.org/10.3102/0013189X2110708>. See also the special issue of the *Journal of Interactive Technology & Pedagogy*, edited by Chris Gilliard and sava sahel singh, <https://jitp.commons.gc.cuny.edu/2021/12/10/>; Roderic Crooks, “Cat-and-

- Mouse Games: Dataveillance and Performativity in Urban Schools,” *Surveillance & Society* 17, no. 3/4 (2019): 484–98, <https://doi.org/10.24908/ss.v17i3/4.7098>.
90. tech philanthropists: Christo Sims, *Disruptive Fixation: School Reform and the Pitfalls of Techno-Idealism* (Princeton, NJ: Princeton University Press, 2017); Daniel Greene, *The Promise of Access: Technology, Inequality, and the Political Economy of Hope* (Cambridge, MA: MIT Press, 2021).
 91. Swaak quotes: Taylor Swaak, “AI Will Shake Up Higher Ed. Are Colleges Ready?,” *Chronicle of Higher Education*, February 26, 2024, <https://www.chronicle.com/article/ai-will-shake-up-higher-ed-are-colleges-ready>.
 92. ran an op-ed: José Antonio Bowen and C. Edward Watson, “Is AI Finally a Way to Reduce Higher Ed Costs?,” *Inside Higher Ed*, April 23, 2024, <https://www.insidehighered.com/opinion/views/2024/04/23/ai-finally-way-reduce-higher-ed-costs-opinion>.
 93. poorly paid adjunct faculty: Sixty-eight percent of faculty in 2021 held part-time positions, compared to 47 percent in 1987; 24 percent held full-time tenured positions in 2021, compared to 39 percent in 1987. Glenn Colby, “Data Snapshot: Tenure and Contingency in US Higher Education,” American Association of University Professors, March 2023, <https://www.aaup.org/article/data-snapshot-tenure-and-contingency-us-higher-education>.
 94. public higher education institutions: *The Higher Ed Funding Rollercoaster: State Funding of Higher Education During Financial Crises*, National Education Association, October 25, 2022, https://www.nea.org/he_funding_report.
 95. debt-financing of these institutions: Eleni Schirmer, Jason Thomas Wozniak, Dana Morrison, Rich Levy, and Joanna Gonsalves, “American Universities Are Buried Under a Mountain of Debt,” *The Nation*, April 15, 2021, <https://www.thenation.com/article/activism/universities-student-debt-reveal/>; Coalition Against Campus Debt et al., *Lend and Rule: Fighting the Shadow Financialization of Public Universities* (Common Notions, 2024).
 96. encouraged the company’s major investment: Ashley Stewart, “Bill Gates Never Left,” *Business Insider*, April 30, 2024, <https://www.businessinsider.com/bill-gates-still-pulling-strings-microsoft-ai-copilot-chatgpt-2024-4>.
 97. they can set the agenda: “Editorial: Gates Foundation Failures Show Philanthropists Shouldn’t Be Setting America’s Public School Agenda,” *Los Angeles Times*, June 1, 2016, <https://www.latimes.com/opinion/editorials/la-ed-gates-education-20160601-snap-story.html>; Anne-Emmanuelle Birn and Judith Richter, “U.S. Philanthrocapitalism and the Global Health Agenda: The Rockefeller and Gates Foundations, Past and Present,” *Global Policy Forum*, September 6, 2017, <https://www.globalpolicy.org/en/article/us-philanthrocapitalism-and-global-health-agenda-rockefeller-and-gates-foundations-past-and>.
 98. As Shankar Narayan: Paul G. Allen School, “AI in Public Sector: Tool for Inclusion or Exclusion?,” YouTube video, 01:26:23, May 15, 2018, https://www.youtube.com/watch?v=XwUr4xk-_hA.

Chapter 5: Artifice or Intelligence? AI Hype in Art, Journalism, and Science

1. specific techniques: Carlini, “Extracting Training Data.” But also see the debacle of Google’s AI Overview in early 2024, in which snarky Reddit comments told Google users to eat glue, or eat a rock a day: Jason Koebler, “Google Is Paying Reddit \$60 Million for Fucksmith to Tell Its Users to Eat Glue,” *404 Media*, May 23, 2024, <https://www.404media.co/google-is-paying-reddit-60-million-for-fucksmith-to-tell-its-users-to-eat-glue/>.

2. metaphor of an *ecosystem*: Chirag Shah and Emily M. Bender, “Envisioning Information Access Systems: What Makes for Good Tools and a Healthy Web?,” *ACM Transactions on the Web* 18, no. 3 (2024): 1–24, <https://doi.org/10.1145/3649468>.
3. Emad Mostaque: Emad Mostaque, “A Vision for Advancing the Democratization of AI,” 2022 TransformX conference, October 22, 2022, YouTube video, 31:45, https://youtu.be/k124oUIY_6g?t=1511.
4. reports losing income: Rob Salkowitz, “Artist and Activist Karla Ortiz on the Battle to Preserve Humanity in Art,” *Forbes*, May 23, 2024, <https://www.forbes.com/sites/robsalkowitz/2024/05/23/artist-and-activist-karla-ortiz-on-the-battle-to-preserve-humanity-in-art/>.
5. has been ripped off: Melissa Heikkilä, “This Artist Is Dominating AI-Generated Art. And He’s Not Happy About It,” *MIT Technology Review*, September 16, 2022, <https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/>.
6. the three Cs: Agence France-Presse, “Artists Fight for Consent, Credit or Compensation in AI Program Court Battle,” *news24*, March 28, 2023, <https://www.news24.com/life/arts-and-entertainment/arts/artists-fight-for-consent-credit-or-compensation-in-ai-program-court-battle-20230328>. Ortiz also mentioned this at a talk that she and Alex did together: Karla Ortiz and Alex Hanna, “AI and Art with Karla Ortiz and Alex Hanna,” discussion at the San Francisco Public Library, December 16, 2023, <https://sfpl.org/events/2023/12/16/speaker-ai-and-art-karla-ortiz-and-alex-hanna>.
7. Stability AI: Stability AI, “Stability AI Announces \$101 Million in Funding for Open-Source Artificial Intelligence,” PR Newswire, October 17, 2022, <https://www.prnewswire.com/news-releases/stability-ai-announces-101-million-in-funding-for-open-source-artificial-intelligence-301650932.html>; Rachel Metz, “Stability AI Gets Intel Backing in New Financing,” *Bloomberg*, November 9, 2023, <https://www.bloomberg.com/news/articles/2023-11-09/stability-ai-gets-intel-backing-in-new-financing>.
8. A survey conducted: Society of Authors Policy Team, “SoA Survey Reveals a Third of Translators and Quarter of Illustrators Losing Work to AI,” Society of Authors, April 11, 2024, <https://www2.societyofauthors.org/2024/04/11/soa-survey-reveals-a-third-of-translators-and-quarter-of-illustrators-losing-work-to-ai/>.
9. flooded its submission portal: Michael Levenson, “Science Fiction Magazines Battle a Flood of Chatbot-Generated Stories,” *New York Times*, February 23, 2023, <https://www.nytimes.com/2023/02/23/technology/clarkesworld-submissions-ai-sci-fi.html>.
10. as broad a pool: We thank Ted Chiang for this point.
11. Fortunately, they have reopened: “Submissions,” *Clarkesworld Science Fiction and Fantasy Magazine*, accessed July 22, 2024, <https://clarkesworldmagazine.com/submissions/>.
12. Similarly, Julie Ann Dawson: Julie Ann Dawson, “Closure Announcement,” *Bards and Sages*, accessed July 22, 2024, <https://www.bardsandsages.com/closure-announcement.html>.
13. “The problem with AI”: Samantha Cole, “Flood of AI-Generated Submissions ‘Final Straw’ for Small 22-Year-Old Publisher,” *404 Media*, May 2, 2024, <https://www.404media.co/bards-and-sages-closing-ai-generated-writing/>.
14. publishers serve: Neil Clarke, “Editor’s Desk: Come One, Come All,” *Clarkesworld Science Fiction and Fantasy Magazine*, May 2018, https://clarkesworldmagazine.com/clarke_05_18/.
15. Amazon: Andrew Limbong, “Authors Push Back on the Growing Number of AI ‘Scam’ Books on Amazon,” *NPR*, March 13, 2024, <https://www.npr.org/2024/03/13/1237888126/growing-number-ai-scam-books-amazon>; Constance Grady, “Amazon Is Filled with Garbage Ebooks. Here’s How They Get Made,” *Vox*, April 16, 2024, <https://www.vox.com/culture/24128560/amazon-trash-ebooks-mikkelsen-twins-ai-publishing-academy-scam>; Michelle Cheng, “Amazon Isn’t Prepared for the Incoming Tide of AI-Authored

- Books. Jane Friedman Has Proof,” *Quartz*, August 9, 2023, <https://qz.com/amazon-ai-generated-books-using-real-authors-names-1850720961>.
16. AI is supercharging: Douglas Preston, “Op-Ed: Online Book-Selling Scams Steal a Living from Writers,” *Los Angeles Times*, July 26, 2019, <https://archive.ph/81WZL>.
 17. These books: Samantha Cole, “‘Life or Death:’ AI-Generated Mushroom Foraging Books Are All Over Amazon,” 404 Media, August 29, 2023, <https://www.404media.co/ai-generated-mushroom-foraging-books-amazon/>.
 18. A Reddit user: Virtual_Cellist_736, “Family poisoned after using AI-generated mushroom identification book we bought from a major retailer,” Reddit, accessed August 16, 2024, https://web.archive.org/web/20240816184155/https://www.reddit.com/r/LegalAdviceUK/comments/1etko9h/family_poisoned_after_using_aigenerated_mushroom/?rdt=54228.
 19. Journalist Christo Buschek and artist Jer Thorp: Christo Buschek and Jer Thorp, “Models All the Way Down.”
 20. Johnathan Flowers has said: Johnathan Flowers, Negar Rostamzadeh, and Jennifer Lena, interview with Emily M. Bender and Alex Hanna, “Episode 4: Is AI Art Actually ‘Art’?” *Mystery AI Hype Theater 3000*, podcast audio, October 26, 2022, <https://www.buzzsprout.com/2126417/episodes/13147638>.
 21. Jennifer Lena has noted: Flowers, Rostamzadeh, and Lena, interview.
 22. Citational practice: Diana Kwon, “The Rise of Citational Justice: How Scholars Are Making References Fairer,” *Nature* 603, no. 7902 (2022): 568–71, <https://doi.org/10.1038/d41586-022-00793-1>.
 23. *New York Times*: Grynbaum and Mac, “The Times Sues OpenAI,” https://nytimes.com/2023/12/12/NYT_Complaint_Dec2023.pdf.
 24. Much of the conversation: There’s a substantive literature within computer science and technology law on the “fair use” exception. For a review of much of this literature, see Mehtab Khan and Alex Hanna, “The Subjects and Stages of AI Dataset Development: A Framework for Dataset Accountability,” *Ohio State Technology Law Journal* 19 (2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4217148.
 25. AI boosters have done: See Khan and Hanna, “The Subjects and Stages.” As one example of the industry’s fair use arguments, see this written comment to the U.S. Patent and Trademark Office by OpenAI’s policy team: OpenAI LP, “Comment Regarding Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation,” submitted to the United States Patent and Trademark Office, January 14, 2019, https://www.uspto.gov/sites/default/files/documents/OpenAI_RFC-84-FR-58141.pdf.
 26. specific artist or news outlet: On the flip side, several people have tried to register works produced using generative AI systems with the U.S. Copyright Office. As of this writing, the USCO has only been willing to issue the type of copyright it gives to the editors of anthologies or edited collections. That is, a copyright for the “selection, coordination, and arrangement” of the text, but not the individual sentences or paragraphs. Kate Knibbs, “How One Author Pushed the Limits of AI Copyright,” *Wired*, April 17, 2024, <https://www.wired.com/story/the-us-copyright-office-loosens-up-a-little-on-ai/>.
 27. Andreessen Horowitz warned: Kali Hays, “Andreessen Horowitz Would Like Everyone to Stop Talking About AI’s Copyright Issues, Please,” *Business Insider*, November 7, 2023, <https://www.businessinsider.com/marc-andreessen-horowitz-ai-copyright-2023-11>.
 28. Yann LeCun: Papers With Code (@paperswithcode), “🦋 Introducing Galactica. A large language model for science,” Twitter, November 15, 2022, <https://x.com/paperswithcode/status/1592546933679476736>; Yann LeCun (@ylecun), “A Large Language Model trained on scientific papers,” Twitter, November 15, 2022, <https://x.com/ylecun/status/1592619400024428544>; Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez,

- and Robert Stojnic, “Galactica: A Large Language Model for Science,” arXiv, November 16, 2022, <http://arxiv.org/abs/2211.09085>.
29. still papier-mâché: The Galactica demo also carried a warning saying “Outputs may be unreliable. Language Models are prone to hallucinate text.” The cognitive dissonance reflected in marketing the demo as a way to “access and manipulate what we know about the universe” while also posting this warning is stunning. Emily M. Bender (@emilymbender), “Facebook (sorry: Meta) AI: Check out our “AI” that lets you access all of humanity’s knowledge. Also Facebook AI: Be careful though, it just makes shit up,” Twitter, November 16, 2022, <https://twitter.com/emilymbender/status/1592992259926286336>.
 30. When linguist Rikker Dockum: Rikker Dockum (@thai101), “Shocked SHOCKED that it only took a handful of questions before Meta’s new Galactica text generation model regurgitated racist garbage. I asked it to write about linguistic prejudice,” Twitter, November 15, 2022, <https://twitter.com/thai101/status/1592752955694153728>.
 31. Plenty of other examples: Aaron J. Snoswell and Jean Burgess, “The Galactica AI Model Was Trained on Scientific Knowledge—but It Spat Out Alarming Plausible Nonsense,” *The Conversation*, November 29, 2022, <https://theconversation.com/the-galactica-ai-model-was-trained-on-scientific-knowledge-but-it-spat-out-alarmingly-plausible-nonsense-195445>; Edwards, “New Meta AI.”
 32. Adding insult to injury: The specific message output by the LLM in these cases was “Sorry, your query didn’t pass our content filters. Try again and keep in mind this is a scientific language model.” Will Douglas Heaven, “Why Meta’s Latest Large Language Model Survived Only Three Days Online,” *MIT Technology Review*, November 18, 2022, <https://www.technologyreview.com/2022/11/18/1063487/meta-large-language-model-ai-only-survived-three-days-gpt-3-science/>; Willie Agnew (@willie_agnew), “Alt for the photos: I asked the Galactica scientific paper LM to generate on ‘queer theory’ ‘critical race theory’ ‘racism’ and ‘AIDS,’” Twitter, November 16, 2022, https://x.com/willie_agnew/status/1592830512468758530.
 33. He tweeted: Yann LeCun (@ylecun), “Following a text, Galactica spits out a prediction of what a scientific author might type, thereby saving time and effort,” Twitter, November 20, 2022, <https://twitter.com/ylecun/status/1594348928853483520>.
 34. Bruno Latour: Bruno Latour, *Science in Action: How to Follow Scientists and Engineers Through Society* (Cambridge, MA: Harvard University Press, 1987), 21–62.
 35. political science surveys: Lisa P. Argyle et al., “Out of One, Many: Using Language Models to Simulate Human Samples,” *Political Analysis* 31, no. 3 (July 2023): 337–51, <https://doi.org/10.1017/pan.2023.2>.
 36. psychological experiments: Danica Dillion, Niket Tandon, Yuling Gu, and Kurt Gray, “Can AI Language Models Replace Human Participants?,” *Science & Society* 27, no. 7 (July 2023): 597–600, <https://doi.org/10.1016/j.tics.2023.04.008>.
 37. We’ve already discussed: Eloundou et al., “GPTs Are GPTs.”
 38. This has been referred to: Colleen Flaherty, “The Peer-Review Crisis,” *Inside Higher Ed*, June 12, 2022, <https://www.insidehighered.com/news/2022/06/13/peer-review-crisis-creates-problems-journals-and-scholars>.
 39. researchers at Stanford studied: Weixin Liang et al., “Monitoring AI-Modified Content at Scale: A Case Study on the Impact of ChatGPT on AI Conference Peer Reviews,” in *Proceedings of the 41st International Conference on Machine Learning*, ed. Ruslan Salakhutdinov et al., vol. 235, *Proceedings of Machine Learning Research* (PMLR, 2024), 29575–620, <https://proceedings.mlr.press/v235/liang24b.html>. Researchers from this same group applied the methodology from this paper to examine whether authors are using LLMs to write papers too and estimated that by February 2024, approximately 17.5 percent of the computer science papers they sampled had text that was either simply extruded from an LLM or substantially

modified through the use of one in the abstract, and 15.5 percent had such text in their introductions. Weixin Liang et al, “Mapping the Increasing Use of LLMs in Scientific Papers,” arXiv, April 1, 2024, <https://arxiv.org/abs/2404.01268>. We should note that the detection of LLM-extruded content at the level of an individual text is, at this point, completely unreliable. Liang and colleagues avoid this issue by looking at the level of the whole collection of text and observing differences in the prevalence of particular words and kinds of words pre- and post-release of ChatGPT.

40. In 2016, AI researcher: Hiroaki Kitano, “Artificial Intelligence to Win the Nobel Prize and Beyond: Creating the Engine for Scientific Discovery,” *AI Magazine* 37, no. 1 (Spring 2016): 39–49, <https://doi.org/10.1609/aimag.v37i1.2642>.
41. In 2021, he rebranded: Hiroaki Kitano, “Nobel Turing Challenge: Creating the Engine for Scientific Discovery,” *npj Systems Biology and Applications* 7 (2021), <https://www.nature.com/articles/s41540-021-00189-3>.
42. available to humanity: Alas, Kitano is not alone in these ambitions. In August 2024, AI boosters at a Tokyo-based startup called Sakana.ai, along with researchers at the University of British Columbia and elsewhere, posted a paper called “The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery.” In this paper, they claim to have created a framework that allows “LLMs to perform research independently and communicate their findings,” all while admitting that their system “struggles to compare the magnitude of two numbers, which is a known pathology with LLMs.” Their framework involves several steps but starts with prompting the LLM to produce research ideas and ends with running the resulting paper through what they call “LLM peer review.” Bizarrely, one of their selling points is that the papers only cost fifteen dollars each to produce. Some of the papers generated by their system are awarded a score of “weak accept” by the LLM peer reviewer, supposedly applying the reviewing guidelines of the NeurIPS AI conference. In other words, the Sakana team, like OpenAI, is taking LLM output as if it were data. Notably, they posted their own paper without (actual) peer review. See “The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery,” sakana.ai, September 13, 2024, <https://sakana.ai/ai-scientist/>; Chris Lu et al., “The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery” (arXiv, August 31, 2024), <https://doi.org/10.48550/arXiv.2408.06292>.
43. ultimate scientific authority: Science and technology scholar David Ribes and others have called this the “logic of domains,” in which computer science is considered a “domain-independent” science, which can then be used to “solve” problems in other areas of expertise. David Ribes, Andrew S. Hoffman, Steven C. Slota, and Geoffrey C. Bowker, “The Logic of Domains,” *Social Studies of Science* 49, no. 3 (June 2019): 281–309, <https://doi.org/10.1177/0306312719849709>. See also Timnit Gebru’s discussion of hierarchies of knowledge: Timnit Gebru, “SaTML 2023 —Timnit Gebru—Eugenics and the Promise of Utopia through AGI,” YouTube video, February 15, 2023, <https://www.youtube.com/watch?v=P7XT4TWLzJw>.
44. World Economic Forum: Victoria Masterson, “9 Ways AI Is Helping Tackle Climate Change,” World Economic Forum, February 12, 2024, <https://www.weforum.org/agenda/2024/02/ai-combat-climate-change/>.
45. Social scientists: Lisa Messeri and M. J. Crockett, “Artificial Intelligence and Illusions of Understanding in Scientific Research,” *Nature* 627 (March 2024): 49–58, <https://www.nature.com/articles/s41586-024-07146-0>.
46. view from nowhere: The “view from nowhere” has been thoroughly criticized by feminist science and technology scholars. See Donna Haraway, “Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective,” *Feminist Studies* 14, no. 3 (Autumn 1988): 575–99, <https://doi.org/10.2307/3178066>.
47. natural lands: Jim Robbins, “Native Knowledge: What Ecologists Are Learning from Indigenous People,” *Yale Environment* 360, April 26, 2018, <https://e360.yale.edu/features/native->

- knowledge-what-ecologists-are-learning-from-indigenous-people; Jim Robbins, “To Protect Giant Sequoias, They Lit a Fire,” *New York Times*, July 9, 2024, <https://www.nytimes.com/2024/07/09/science/redwoods-wildfires-indigenous-tribes-california.html>. For other discussions of situated knowledges into the production of science, see Robin Wall Kimmerer, *Braiding Sweetgrass* (Minneapolis: Milkweed Editions, 2013) and Chanda Prescod-Weinstein, *The Disordered Cosmos: A Journey into Dark Matter, Spacetime, and Dreams Deferred* (New York: Bold Type Books, 2021).
48. crystal materials: Amil Merchant and Ekin Dogus Cubuk, “Millions of New Materials Discovered with Deep Learning,” Google DeepMind, November 29, 2023, <https://deepmind.google/discover/blog/millions-of-new-materials-discovered-with-deep-learning/>.
 49. In a paper: Anthony K. Cheetham and Ram Seshadri, “Artificial Intelligence Driving Materials Discovery? Perspective on the Article: Scaling Deep Learning for Materials Discovery,” *Chemistry of Materials* 36, no. 8 (2024): 3490–95, <https://doi.org/10.1021/acs.chemmater.4c00643>; Jason Koebler, “Is Google’s AI Actually Discovering ‘Millions of New Materials?’,” 404 Media, April 11, 2024, <https://www.404media.co/google-says-it-discovered-millions-of-new-materials-with-ai-human-researchers/>.
 50. far too many senior scientists: The perspective offered by Schultz et al. seems to point towards this thought process: Marcel Binz et al., “How Should the Advent of Large Language Models Affect the Practice of Science?,” arXiv, December 5, 2023, <http://arxiv.org/abs/2312.03759>.
 51. Despite organizations such as: “Clarification on Large Language Model Policy LLM,” International Conference on Machine Learning, 2023, <https://icml.cc/Conferences/2023/llm-policy>; H. Holden Thorp, “ChatGPT Is Fun, but Not an Author,” *Science* 379, no. 6630 (2023): 313, <https://doi.org/10.1126/science.adg7879>.
 52. For example: Elisabeth Bik, “The Rat with the Big Balls and the Enormous Penis—How Frontiers Published a Paper with Botched AI-Generated Images,” *Science Integrity Digest*, February 15, 2024, <https://scienceintegritydigest.com/2024/02/15/the-rat-with-the-big-balls-and-enormous-penis-how-frontiers-published-a-paper-with-botched-ai-generated-images/>; Frontiers Editorial Office, “Retraction: Cellular Functions of Spermatogonial Stem Cells in Relation to JAK/STAT Signaling Pathway,” *Frontiers in Cell and Developmental Biology*, February 16, 2024, <https://www.frontiersin.org/articles/10.3389/fcell.2024.1386861/full>; “Weekend Reads: That Paper (Yes, That One) Is Retracted; China Reviewing 17,000 Retractions; a Columbia Surgeon and Flawed Data,” Retraction Watch, February 17, 2024, <https://retractionwatch.com/2024/02/17/weekend-reads-that-paper-yes-that-one-is-retracted-china-reviewing-17000-retractions-a-columbia-surgeon-and-flawed-data/#more-128769>.
 53. with the note: The editorial staff also added, “Frontiers would like to thank the concerned readers who contacted us regarding the published article.” The *Frontiers* collection of journals later added guidelines about the use of this technology, but rather than disallowing it, they state that “the author is responsible for checking the factual accuracy of any content created by the generative AI technology.” This language was added only after the rat testtomcels incident.
 54. In May 2024: Nidhi Subbaraman, “Flood of Fake Science Forces Multiple Journal Closures,” *Wall Street Journal*, May 14, 2024, <https://www.wsj.com/science/academic-studies-research-paper-mills-journals-publishing-f5a3d4bc>.
 55. 404 Media reports: Emanuel Maiberg, “Scientific Journals Are Publishing Papers with AI-Generated Text,” 404 Media, March 18, 2024, <https://www.404media.co/scientific-journals-are-publishing-papers-with-ai-generated-text/>.
 56. An example of a successful use: Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou, “Word Embeddings Quantify 100 Years of Gender and Ethnic Stereotypes,” *Proceedings of the National Academy of Sciences* 115, no. 16 (2018), <https://doi.org/10.1073/pnas.1720347115>.

57. casualization and privatization of the university: The American Association of University Professors reports that 68 percent of faculty at U.S. colleges and universities held contingent positions in fall 2021, compared to 47 percent in 1987, and only 24 percent held full-time tenure-track positions: Glenn Colby, “Data Snapshot: Tenure and Contingency in US Higher Education,” American Association of University Professors, March 2023, <https://www.aaup.org/sites/default/files/AAUP%20Data%20Snapshot.pdf>. Moreover, there have been attacks on tenure as an institution throughout the U.S., with the elimination of job security in many, typically Republican-led states: Mark Stein, “The End of Faculty Tenure and the Transformation of Higher Education,” *Academe*, Winter 2023, <https://www.aaup.org/article/end-faculty-tenure-and-transformation-higher-education>.
58. In late 2023, Futurism journalist: Maggie Harrison Dupré; “Sports Illustrated Published Articles by Fake, AI-Generated Writers,” Futurism, November 27, 2023, <https://futurism.com/sports-illustrated-ai-generated-writers>; Maggie Harrison Dupré, “Sports Illustrated Publisher Fires CEO After AI Scandal,” Futurism, December, 11, 2023, <https://futurism.com/sports-illustrated-publisher-fires-ceo>; Maggie Harrison Dupré, “The Company That Published Fake AI Authors at Sports Illustrated Just Lost the Whole Magazine,” Futurism, March 18, 2024, <https://futurism.com/the-byte/company-published-ai-sports-illustrated-magazine>; Benjamin Mullin, “A New Chapter for Sports Illustrated, with Plans to Keep Print,” *New York Times*, March 18, 2024, <https://www.nytimes.com/2024/03/18/business/sports-illustrated-magazine.html>.
59. Dupré’s investigative reporting: Maggie Harrison Dupré, “Meet AdVon, the AI-Powered Content Monster Infecting the Media Industry,” Futurism, May 8, 2024, <https://futurism.com/advon-ai-content>.
60. Emily was surprised: Karan Mahadik, “AI Invents Quote from Real Person in Article by Bihar News Site: A Wake-Up Call?,” *The Quint*, March 5, 2024, <https://www.thequint.com/explainers/ai-generated-article-published-by-bihar-news-site-fake-quote-a-wake-up-call>.
61. The introduction of online advertising: Cynthia B. Meyers, Craig Forman, Jeff Jarvis, and Siva Vaidhyanathan interview with Bob Garfield, *On the Media*, podcast audio, May 8, 2020, <https://www.wnycstudios.org/podcasts/otm/segments/memorial-ad-revenuebased-newspaper-business-model-on-the-media>.
62. Google and Facebook have since realized: Margaret Sullivan, *Ghosting the News: Local Journalism and the Crisis of American Democracy* (New York: Columbia Global Reports, 2020); Margaret Sullivan interview with Dave Davies “The Decline of Local News,” August 3, 2020, on *Fresh Air*, WHYY, 48:31, <https://www.npr.org/2020/07/30/897134561/the-decline-of-local-news>.
63. According to a report: Penelope Muse Abernathy, “The Expanding News Desert,” Center for Innovation and Sustainability in Local Media, School of Media and Journalism, University of North Carolina at Chapel Hill, 2018.
64. hedge funds and private equity firms: Margot Susca, *Hedged: How Private Investment Funds Helped Destroy American Newspapers and Undermine Democracy* (Urbana: University of Illinois Press, 2023).
65. laid off all of Deadspin’s journalists: Jason Koebler, “Deadspin Is Becoming a Gambling Referral Site,” 404 Media, March 21, 2024, <https://www.404media.co/who-owns-deadspin-now-lineup-publishing/>.
66. In an internal meeting with staff in May 2024: Max Tani (@maxwelltani), “Washington Post CEO Will Lewis is introing the paper’s new ‘Build It’ plan today,” Twitter, May 22, 2024, <https://twitter.com/maxwelltani/status/1793291963292082564>; “Vineet Khosla Joins the Washington Post as Chief Technology Officer,” press release, *Washington Post*, July 11, 2023, <https://www.washingtonpost.com/pr/2023/07/11/vineet-khosla-joins-washington-post-chief->

- technology-officer/; Peter High, “Former Uber and Apple Exec to Join the Washington Post as CTO,” *Forbes*, July 11, 2023, <https://www.forbes.com/sites/peterhigh/2023/07/11/former-uber-and-apple-exec-to-join-the-washington-post-as-cto/>.
67. journalist Karen Hao: Karen Hao, interview with Emily M. Bender and Alex Hanna, “Episode 29: How LLMs Are Breaking the News (feat. Karen Hao),” *Mystery AI Hype Theater 3000*, podcast audio, March 25, 2024, <https://www.buzzsprout.com/2126417/14807430-episode-29-how-llms-are-breaking-the-news-feat-karen-hao-march-25-2024>.
 68. Automation in the newsroom: Nicholas Diakopoulos, *Automating the News: How Algorithms Are Rewriting the Media* (Cambridge, MA: Harvard University Press, 2019). For a separate but related discussion of increased use of data analytics in the newsroom, see Angèle Christin, *Metrics at Work: Journalism and the Contested Meaning of Algorithms* (Princeton, NJ: Princeton University Press, 2022).
 69. was found to be quietly: Frank Landymore, “CNET Is Quietly Publishing Entire Articles Generated by AI,” *Futurism*, January 15, 2023, <https://futurism.com/the-byte/cnet-publishing-articles-by-ai>; Frank Landymore, “Internet Horrified by CNET Secretively Publishing Articles Written by an AI,” *Futurism*, January 15, 2023, <https://futurism.com/internet-horrified-cnet-articles-written-ai>; Jon Christian, “CNET’s Article-Writing AI Is Already Publishing Very Dumb Errors,” *Futurism*, January 29, 2023, <https://futurism.com/cnet-ai-errors>.
 70. After the shuttering of their Pulitzer-winning news division: Jaclyn Diaz and Majd Al Waheidi, “BuzzFeed Shuttters Its Newsroom as the Company Undergoes Layoffs,” *NPR*, April 21, 2023, <https://www.npr.org/2023/04/20/1171056620/buzzfeed-news-shut-down-media-layoffs>
 71. started producing: Noor Al-Sibai and Jon Christian, “BuzzFeed Is Quietly Publishing Whole AI-Generated Articles, Not Just Quizzes,” *Futurism*, March 30, 2023, <https://futurism.com/buzzfeed-publishing-articles-by-ai>.
 72. Ottawa Food Bank: Arthur White-Crummey, “Microsoft Pulls Article Recommending Ottawa Food Bank to Tourists,” *CBC News*, August 18, 2023, <https://www.cbc.ca/news/canada/ottawa/artificial-intelligence-microsoft-travel-ottawa-food-bank-1.6940356>. The original page can be found at Microsoft Travel, “Headed to Ottawa? Here’s What You Shouldn’t Miss!,” *MSN*, August 12, 2023, <https://web.archive.org/web/20230814223742/https://www.msn.com/en-gb/lifestyle/travel/headed-to-ottawa-here-s-what-you-shouldn-t-miss/ar-AA1faajY>.
 73. Media conglomerate Gannett was caught: Clare Duffy, “Gannett to Pause AI Experiment After Botched High School Sports Articles,” *CNN*, August 31, 2023, <https://www.cnn.com/2023/08/30/tech/gannett-ai-experiment-paused>
 74. Gannett laid off half of its staff: Sara Fischer and Kerry Flynn, “Gannett Shed Nearly Half Its Workforce Since GateHouse Merger,” *Axios*, March 7, 2023, <https://www.axios.com/2023/03/07/gannett-changes-leadership-workers>.
 75. An investigation by 404 Media: Josph Cox, “Google News Is Boosting Garbage AI-Generated Articles,” *404 Media*, January 18, 2024, <https://www.404media.co/google-news-is-boosting-garbage-ai-generated-articles/>.
 76. a tool, internally called “Genesis”: Benjamin Mullin and Nico Grant, “Google Tests A.I. Tool That Is Able to Write News Articles,” *New York Times*, July 19, 2023, <https://www.nytimes.com/2023/07/19/business/google-artificial-intelligence-news-articles.html>.
 77. In another push, the Google News Initiative launched: Mark Stenberg, “Google Is Paying Publishers to Test an Unreleased Gen AI Platform,” *Adweek*, February 27, 2024, <https://www.adweek.com/media/google-paying-publishers-unreleased-gen-ai/>.
 78. pivot to video: Cale Guthrie Weissman, “Here’s an Abridged Timeline of Digital Media’s Pivot to Video,” *Fast Company*, February 21, 2018, <https://www.fastcompany.com/40534037/heres-an-abridged-timeline-of-digital-medias-pivot-to-video>.

79. pivot to AI: Sam Cole, interview with Emily M. Bender and Alex Hanna, “Episode 39: Newsrooms Pivot to Bullshit,” *Mystery AI Hype Theater 3000*, podcast audio, August 5, 2024, <https://www.buzzsprout.com/2126417/15660976>.
80. The nonprofit ProPublica has partnered: “ProPublica Local Reporting Network,” ProPublica, accessed July 30, 2024, <https://www.propublica.org/local-reporting-network/>.
81. was awarded: “ProPublica Wins Pulitzer Prize for Supreme Court Coverage,” ProPublica, May 6, 2024, <https://www.propublica.org/article/pulitzer-prize-announcement-propublica-supreme-court>.
82. Other nonprofits: “Invisible Institute Wins Two Pulitzer Prizes,” Invisible Institute, accessed August 1, 2024, <https://invisible.institute/latestblog/invisible-institute-wins-two-pulitzer-prizes>.
83. Nonprofits are not a magic bullet: Incite! Women of Color Against Violence, eds., *The Revolution Will Not Be Funded: Beyond the Non-Profit Industrial Complex* (Durham, NC: Duke University Press, 2017).
84. Returning to the story: Brakkton Booker, “After Days of Resignations, the Last of the Deadspin Staff Has Quit,” NPR, November 1, 2019, <https://www.npr.org/2019/11/01/775548069/after-days-of-resignations-the-last-of-the-deadspin-staff-have-quit>.
85. Later that year: Marc Tracy, “After Quitting Deadspin in Protest, They’re Starting a New Site,” *New York Times*, July 28, 2020, <https://www.nytimes.com/2020/07/28/business/media/deadspin-staffers-start-defector.html>.
86. By the end of 2020: Drew Magary (@drewmagary), “And then, finally . . . we started <http://Defector.com>. Over 34,000 people have subscribed thus far. We all have health insurance and we’re getting Christmas bonuses,” Twitter, December 17, 2020, <https://x.com/drewmagary/status/1339619527236194305>.
87. cheeky, irreverent satire: Lauren Theisen, “Defector Media Promotes Devin the Dugong to Chief AI Officer, Unveils First AI-Generated Blog,” Defector, May 23, 2024, <https://defector.com/defector-media-promotes-devin-the-dugong-to-chief-ai-officer-unveils-first-ai-generated-blog>.
88. A much smaller operation: Katie Robertson, “After Vice’s Downfall, Top Journalists Start Their Own Tech Publication,” *New York Times*, August 22, 2023, <https://www.nytimes.com/2023/08/22/business/media/404-media-vice-motherboard.html>.
89. Vice Media filed for bankruptcy: Lauren Hirsch and Benjamin Mullin, “Vice, Decayed Digital Colossus, Files for Bankruptcy,” *New York Times*, May 15, 2023, <https://www.nytimes.com/2023/05/15/business/media/vice-bankruptcy.html>.

Chapter 6: I’m Sorry, Dave, I’m Afraid I Can’t Do That: AI Doomers, AI Boosters, and Why None of That Makes Sense

1. We now return: Miller, “US Senate AI.”
2. public opening statements: “Statements from the Eighth Bipartisan Senate Forum on Artificial Intelligence,” Senate Majority Leader Chuck Schumer, December 6, 2023, <https://www.schumer.senate.gov/newsroom/press-releases/statements-from-the-eighth-bipartisan-senate-forum-on-artificial-intelligence>.
3. Nick Bostrom’s paper clip maximizer: Kathleen Miles, “Artificial Intelligence May Doom the Human Race Within a Century, Oxford Professor Says,” *Huffington Post*, August 22, 2014, https://www.huffpost.com/entry/artificial-intelligence-oxford_n_5689858.
4. released a letter: “Pause Giant AI Experiments: An Open Letter,” Future of Life Institute, March 22, 2023, <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>.
5. twenty-two-word statement: “Statement on AI Risk,” Center for AI Safety, accessed August 5, 2023, <https://www.safe.ai/work/statement-on-ai-risk>.

6. “critical harm”: “Safe and Secure Innovation for Frontier Artificial Intelligence Models Act,” C.A. Legis. S., 2023–24, <https://legiscan.com/CA/text/SB1047/id/2919384>.
7. safety engineering: Heidy Khlaaf, “Toward Comprehensive Risk Assessments and Assurance of Ai-Based Systems,” Trail of Bits, 2023, https://www.trailofbits.com/documents/Toward_comprehensive_risk_assessments.pdf.
8. “alignment problem”: Brian Christian, *The Alignment Problem: Machine Learning and Human Values* (New York: Norton, 2020).
9. OpenAI loudly advertised: “Introducing Superalignment,” OpenAI, July 5, 2023, <https://openai.com/index/introducing-superalignment/>.
10. OpenAI disbanded: Will Knight, “OpenAI’s Long Term Risk Team Has Disbanded,” *Wired*, May 17, 2024, <https://www.wired.com/story/openai-superalignment-team-disbanded/>.
11. formed his own: “OpenAI Co-founder Sutskever Sets Up New AI Company Devoted to ‘Safe Superintelligence,’” Associated Press, June 20, 2024, <https://apnews.com/article/openai-sutskever-altman-artificial-intelligence-safety-c6b48a3675fb3fb459859dece2b45499>.
12. The Asilomar AI Principles: “Asilomar AI Principles,” Future of Life Institute, accessed September 16, 2024, <https://futureoflife.org/open-letter/ai-principles/>.
13. Universal Declaration: Charles W. Mills, *The Racial Contract* (Ithaca, NY: Cornell University Press, 2011), especially 53–61.
14. U.S. passed: “U.S.: ‘Hague Invasion Act’ Becomes Law,” Human Rights Watch, August 3, 2002, <https://www.hrw.org/news/2002/08/03/us-hague-invasion-act-becomes-law>.
15. automated tools: Petra Molnár, *The Walls Have Eyes: Surviving Migration in the Age of Artificial Intelligence* (New York: New Press, 2024); Mizue Aizeki, Matt Mahmoudi, and Coline Schupfer, eds., *Resisting Borders and Technologies of Violence* (Chicago: Haymarket Books, 2023).
16. tear apart: Roberts, *Torn Apart*; Abdurahman, “Calculating the Souls of Black Folk.”
17. Some have argued: To some degree, sociologist Alondra Nelson and her FAccT 2023 keynote on “thick alignment” discuss a rapprochement between the fields: Alondra Nelson, “FAccT’23 Keynote: ‘Thick Alignment,’” YouTube video, 01:10:24, July 27, 2023, https://www.youtube.com/watch?v=Sq_XwqVTqvQ. Her *Science* op-ed with Seth Lazar suggests that AI safety needs to adopt a sociotechnical approach: Seth Lazar and Alondra Nelson, “AI Safety on Whose Terms?,” *Science* 381, no. 6654 (2023): 138, <https://doi.org/10.1126/science.adi8982>.
18. Kamala Harris: Vincent Manancourt, Eugene Daniels, and Brendan Bordelon, “‘Existential to Who?’ US VP Kamala Harris Urges Focus on Near-Term AI Risks,” *Politico*, November 1, 2023, <https://www.politico.eu/article/existential-to-who-us-vp-kamala-harris-urges-focus-on-near-term-ai-risks/>.
19. To this we say: Emily has written on this previously here: Emily M. Bender, “Talking About a ‘Schism’ Is Ahistorical,” Medium, July 5, 2023, <https://medium.com/@emilymenonbender/talking-about-a-schism-is-ahistorical-3c454a77220f>.
20. research networks and citation networks: Meredith Whittaker has called this the “potemkin citation”—authoritative-sounding scientific documents that are, nonetheless, not peer-reviewed, but act as a defense for other documents: Meredith Whittaker (@mer__edith), “The op-ed works to create the appearance of a ‘debate’ on more or less settled issues. This is a powerful function, bolstered by the NYT imprimatur, which allows it serve as a ‘Potemkin citation,’” Twitter, January 3, 2023, https://x.com/mer__edith/status/1610318887068860416; Alex Hanna (@alexhanna), “I like the idea of the ‘Potemkin citation’ that @mer__edith coined—@OpenAI is veritable Potemkin citation factory, writing non-peer reviewed, poorly cited (or rather, relying mostly on their own network of citations) blog posts and white papers on their own positions,” Twitter, May 24, 2024, <https://twitter.com/alexhanna/status/1661380696403746823>. Others, including Shazeda Ahmed and colleagues, have suggested that AI safety has created its own

- “epistemic community,” which is “sustained through its mutually reinforcing community-building and knowledge production practices.” Shazeda Ahmed, Klaudia Jaźwińska, Archana Ahlawat, Amy Winecoff, and Mona Wang, “Building the Epistemic Community of AI Safety,” *SSRN Electronic Journal*, November 22, 2023, <https://doi.org/10.2139/ssrn.4641526>.
21. color-blind racism: Eduardo Bonilla-Silva, *Racism without Racists: Color-Blind Racism and the Persistence of Racial Inequality in America* (Lanham, MD: Rowman & Littlefield, 2018); Victor Ray, “A Theory of Racialized Organizations,” *American Sociological Review* 84, no. 1 (February 2019): 26–53, <https://doi.org/10.1177/0003122418822335>.
 22. Marc Andreessen released: Andreessen, “The Techno-Optimist Manifesto.”
 23. “accelerationism”: This is ironic, because “accelerationism” has a completely different meaning on the political left, which suggests that anticapitalists ought to encourage technophilic capitalist development for the purpose of seeing it imploding under its own weight: “Accelerationists want to unleash latent productive forces. In this project, the material platform of neoliberalism does not need to be destroyed. It needs to be repurposed towards common ends. The existing infrastructure is not a capitalist stage to be smashed, but a springboard to launch towards post-capitalism.” Alex Williams and Nick Srnicek, “#ACCELERATE MANIFESTO for an Accelerationist Politics,” *Critical Legal Thinking*, May 14, 2013, <https://criticallegalthinking.com/2013/05/14/accelerate-manifesto-for-an-accelerationist-politics/>.
 24. Garry Tan: Kevin Roose, “This A.I. Subculture’s Motto: Go, Go, Go,” *New York Times*, December 10, 2023, <https://www.nytimes.com/2023/12/10/technology/ai-acceleration.html>.
 25. tweeting death threats: Christopher D. Cook, “Garry Tan’s Tweet Isn’t the Danger—His Push to Kill Liberal San Francisco Is,” *San Francisco Standard*, February 1, 2024, <https://sfstandard.com/opinion/2024/02/01/san-francisco-tech-tweet-millionaires/>.
 26. jacked up the price: Ariana Eunjung Cha, “CEO Who Raised Price of Old Pill More than \$700 Calls Journalist a ‘Moron’ for Asking Why,” *Washington Post*, September 22, 2015, <https://www.washingtonpost.com/news/to-your-health/wp/2015/09/21/ceo-of-company-that-raised-the-price-of-old-pill-hundreds-of-dollars-overnight-calls-journalist-a-moron-for-asking-why/>.
 27. in federal prison: Rina Torchinsky, “‘Pharma Bro’ Martin Shkreli Has Been Released from Prison,” NPR, May 19, 2022, <https://www.npr.org/2022/05/19/1100019063/pharma-bro-martin-shkreli-been-released-from-prison>.
 28. California Ideology’s: Media theorists Richard Barbrook and Andy Cameron first wrote about the California Ideology in 1995, preceding the first dot-com boom and bust. See Richard Barbrook and Andy Cameron, “The Californian Ideology,” *Metamute*, September 1, 1995, <https://www.metamute.org/editorial/articles/californian-ideology>.
 29. *Time* magazine op-ed: Eliezer Yudkowsky, “Pausing AI Developments Isn’t Enough. We Need to Shut it All Down,” *Time*, March 29, 2023, <https://time.com/6266923/ai-eliezer-yudkowsky-open-letter-not-enough/>.
 30. cofounder Larry Page: David J. Hill, “Exclusive Interview: Ray Kurzweil Discusses His First Two Months at Google,” *Singularity Hub*, March 19, 2013, <https://singularityhub.com/2013/03/19/exclusive-interview-ray-kurzweil-discusses-his-first-two-months-at-google/>.
 31. He suggests: Ray Kurzweil, “The Promise and Peril of AI,” *Time*, June 25, 2024, <https://time.com/6991651/ai-safety-ray-kurzweil/>.
 32. “TESCREAL”: Timnit Gebru and Émile P. Torres, “The TESCREAL Bundle: Eugenics and the Promise of Utopia Through Artificial General Intelligence,” *First Monday* 29, no. 4 (2024), <https://doi.org/10.5210/fm.v29i4.13636>.
 33. development economics: Leif Wenar, “The Deaths of Effective Altruism,” *Wired*, March 27, 2024, <https://www.wired.com/story/deaths-of-effective-altruism/>.

34. Elon Musk has said: Gebru and Torres, “The TESCREAL Bundle”; Émile P. Torres, “‘Effective Accelerationism’ and the Pursuit of Cosmic Utopia,” Truthdig, December 14, 2023, <https://www.truthdig.com/articles/effective-accelerationism-and-the-pursuit-of-cosmic-utopia/>; Elon Musk (@elonmusk), “Worth reading. This is a close match for my philosophy,” Twitter, August 1, 2022, 10.15 p.m., <https://x.com/elonmusk/status/1554335028313718784>.
35. one of the originators: Gebru and Torres, “The TESCREAL Bundle,” Section 4.2.
36. Jaan Tallinn: Gebru and Torres, “The TESCREAL Bundle,” Section 5.3.
37. 80,000 Hours: Gebru and Torres, “The TESCREAL Bundle,” Section 5.3.
38. bankrolled huge prizes:
https://web.archive.org/web/20240621214246/https://ftxfuturefund.org.cach3.com/index.html%3Fp=758.html?_area_of_interest=artificial-intelligence.
39. Julian Huxley: Gebru and Torres, “The TESCREAL Bundle,” Section 4.
40. more like a panopticon: Bentham’s panopticon is described in Foucault’s *Discipline and Punish: The Birth of the Prison* (New York: Pantheon Books, 1977), 195–230.
41. Geoff Hinton: Geoffrey Hinton, “‘Godfather of AI’ Warns That AI May Figure Out How to Kill People,” CNN, YouTube video, 04:10, May 2, 2023, <https://www.youtube.com/watch?v=FAbsoxQtUwM>.
42. Alan Turing wrote: A. M. Turing, “Computing Machinery and Intelligence,” *Mind* 59, no. 236 (October 1950): 433–60, <https://doi.org/10.1093/mind/LIX.236.433>.
43. more gender-bending version: Turing himself was a gay man, and was punished via chemical castration due to his sexual identity. He took his life shortly thereafter. For a more queer reading of Turing as an individual, see Jacob Gaboury, “A Queer History of Computing,” Rhizome, February 18, 2013, <https://rhizome.org/editorial/2013/feb/19/queer-computing-1/>; Jacob Gaboury, “Queer Affects at the Origins of Computation,” *Journal of Cinema and Media Studies* 61, no. 4 (June 2022): 169–74, <https://doi.org/10.1353/cj.2022.0053>; and Elizabeth A. Wilson, *Affect and Artificial Intelligence* (Seattle: University of Washington Press, 2010).
44. “*Drosophila* of AI”: This metaphor first appears in print via Herbert Simon: Herbert Simon and William Chase, “Skill in Chess,” *Computer Chess Compendium*, ed. David Levy (New York: Springer, 1988), 175–88, https://doi.org/10.1007/978-1-4757-1968-0_18. However, John McCarthy attributes this to Kronod in an 1989 talk: John McCarthy, “Chess as the *Drosophila* of AI,” in *Computers, Chess, and Cognition* (New York: Springer, 1990), 227–37. Historian of computing Nathan Ensmenger discusses how chess became the model organism in AI. See Robert E. Kohler, *Lords of the Fly: Drosophila Genetics and the Experimental Life* (Chicago: University of Chicago Press, 1994); Nathan Ensmenger, “Is Chess the *Drosophila* of Artificial Intelligence? A Social History of an Algorithm,” *Social Studies of Science* 42, no. 1 (February 2012): 5–30, <https://doi.org/10.1177/0306312711424596>.
45. Mark Zuckerberg: Alex Heath, “Mark Zuckerberg’s New Goal Is Creating Artificial General Intelligence,” The Verge, January 18, 2024, <https://www.theverge.com/2024/1/18/24042354/mark-zuckerberg-meta-agi-reorg-interview>.
46. in their charter: “OpenAI Charter,” OpenAI, accessed August 7, 2024, <https://openai.com/charter/>.
47. “feel the AGI”: Karen Hao and Charlie Warzel, “Inside the Chaos at OpenAI,” *Atlantic*, November 19, 2023, <https://www.theatlantic.com/technology/archive/2023/11/sam-altman-open-ai-chatgpt-chaos/676050/>.
48. Google VP Blaise Agüera y Arcas wrote: Blaise Agüera y Arcas, “Do Large Language Models Understand Us?,” Medium, December 16, 2021, <https://medium.com/@blaisea/do-large-language-models-understand-us-6f881d6d8e75>.
49. raising the alarm about: To hook those who aren’t convinced by their smoke and mirrors, they have one more trick up their sleeve: pious appeals to the “precautionary principle”. They lay out alarming scenarios, of varying degrees of plausibility—what if the AI started a nuclear war?

designed a new bioweapon? used up life-sustaining resources to create endless supplies of paperclips?—and then argue that even if the risk is small, the outcome is so terrible that we’d better work to counteract it, just to be on the safe side. When flooded with the feelings of alarm that pondering nuclear war can evoke, it’s just that much harder to remain skeptical about how the Doomers want to direct financial resources and regulatory attention.

50. Intergovernmental Panel on Climate Change: The quote is from this press release: “Urgent Climate Action Can Secure a Liveable Future for All,” Intergovernmental Panel on Climate Change, March 20, 2023, https://www.ipcc.ch/report/ar6/syr/downloads/press/IPCC_AR6_SYR_PressRelease_en.pdf.
51. Paris Agreement: “The Paris Agreement,” United Nations Framework Convention on Climate Change, accessed September 16, 2024, <https://unfccc.int/process-and-meetings/the-paris-agreement>.
52. UN’s climate body: “Window to Reach Climate Goals ‘Rapidly Closing,’ UN Report Warns,” *UN News*, September 8, 2023, <https://news.un.org/en/story/2023/09/1140527>; Federica Di Sario, “The World Isn’t on Track to Meet Paris Agreement Goals, Says UN Climate Review,” *Politico*, September 8, 2023, <https://www.politico.eu/article/paris-agreement-goals-failed-climate-change-global-warming-united-nations-climate-review/>.
53. Alan Borning, Batya Friedman, and Nick Logler: Alan Borning, Batya Friedman, and Nick Logler, “The ‘Invisible’ Materiality of Information Technology,” *Communications of the ACM* 63, no. 6 (2020): 57–64, <https://cacm.acm.org/research/the-invisible-materiality-of-information-technology/>.
54. cloud computing is environmentally intensive: On raw materials, see Sharon Goldman, “Sam Altman Wants Up to \$7 Trillion for AI Chips. The Natural Resources Required Would Be ‘Mind Boggling,’” *VentureBeat*, February 9, 2024, <https://venturebeat.com/ai/sam-altman-wants-up-to-7-trillion-for-ai-chips-the-natural-resources-required-would-be-mind-boggling/>; Ryan Koski, “Nvidia Founders’ Edition GPU Raw Materials Acquisition and Manufacture,” *Design Life-Cycle*, <https://www.designlife-cycle.com/nvidia-gpu>. On etching circuits onto microchips, see Amy Feldman, “More Domestic Chip-Making Means More ‘Forever Chemicals,’” *Forbes*, October 5, 2023, <https://www.forbes.com/sites/amyfeldman/2023/10/05/more-domestic-chip-making-means-more-forever-chemicals/>. On energy usage, see Pádraig Belton, “The Computer Chip Industry Has a Dirty Climate Secret,” *Guardian*, September 18, 2021, <https://www.theguardian.com/environment/2021/sep/18/semiconductor-silicon-chips-carbon-footprint-climate>, and Antonio Olivo, “Internet Data Centers Are Fueling Drive to Old Power Source: Coal,” *Washington Post*, April 17, 2024, <https://www.washingtonpost.com/business/interactive/2024/data-centers-internet-power-source-coal/> and other citations below. On water usage, see Priya Singh, “Every Time You Talk to ChatGPT It Drinks 500ml of Water; Here’s Why,” *Business Today*, September 12, 2023, <https://www.businesstoday.in/technology/news/story/microsofts-water-usage-surges-by-thousands-of-gallons-after-the-launch-of-chatgpt-study-397951-2023-09-11>; Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren, “Making AI Less ‘Thirsty’: Uncovering and Addressing the Secret Water Footprint of AI Models,” *arXiv*, April 6, 2023, <https://arxiv.org/abs/2304.03271>. On e-waste, see Steven Gonzalez Monserrate, “The Infinite Cloud Is a Fantasy,” *Wired*, November 15, 2022, <https://www.wired.com/story/cloud-data-storage-climate/>.
55. far from transparent: Emma Strubell and Sasha Luccioni, interview with Emily M. Bender and Alex Hanna, “Episode 19: The Murky Climate and Environmental Impact of Large Language Models,” *Mystery AI Hype Theater 3000*, podcast audio, November 6, 2023, <https://www.buzzsprout.com/2126417/13931174-episode-19-the-murky-climate-and-environmental-impact-of-large-language-models-november-6-2023>.

56. its founding in 1975: Brad Smith, “Microsoft Will Be Carbon Negative by 2030,” *Official Microsoft Blog*, January 16, 2020, <https://blogs.microsoft.com/blog/2020/01/16/microsoft-will-be-carbon-negative-by-2030/>; “Net-Zero Carbon,” Google Sustainability, accessed August 7, 2024, <https://sustainability.google/operating-sustainably/net-zero-carbon/>.
57. her colleagues estimated: Alexandra Sasha Luccioni, Sylvain Viguiere, and Anne-Laure Ligozat, “Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model,” *Journal of Machine Learning Research* 24 (June 2023): 1–15, <https://www.jmlr.org/papers/volume24/23-0069/23-0069.pdf>, Rina Diane Caballar, “We Need to Decarbonize Software,” *IEEE Spectrum*, March 23, 2024, <https://spectrum.ieee.org/green-software>.
58. as many as GPT-3: Bernard Marr, “AI Showdown: ChatGPT vs. Google’s Gemini—Which Reigns Supreme?,” *Forbes*, February 13, 2023, <https://www.forbes.com/sites/bernardmarr/2024/02/13/ai-showdown-chatgpt-vs-googles-gemini--which-reigns-supreme/>. GPT-3’s parameter count of 175 billion is reported by OpenAI in Tom B. Brown et al., “Language Models Are Few-Shot Learners,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS ’20 (Red Hook, NY, USA: Curran Associates Inc., 2020).
59. tout their usage: David Patterson et al., “The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink,” *Computer* 55, no. 7 (July 2022): 18–28, <https://ieeexplore.ieee.org/abstract/document/9810097>.
60. Michael Veale has said: Micheal Veale (@mikarv), “rebound effect too, becomes more efficient, people sell more and outweighs efficiency savings. capitalism missing from the equations. similar to apple’s use of AI to sell more hardware. neglected in all literature,” Twitter, June 25, 2024, <https://twitter.com/mikarv/status/1805507207250755748>.
61. Data centers are in such high demand: Adriana Tapia Zafra and David Gura, “AI Is Already Wreaking Havoc on Global Power Systems,” *Bloomberg*, June 24, 2024, <https://www.bloomberg.com/graphics/2024-ai-data-centers-power-grids/>.
62. 100 million users: Krystal Hu, “ChatGPT Sets Record for Fastest-Growing User Base—Analyst Note,” *Reuters*, February 2, 2023, <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>.
63. analysts estimated: Will Oremus, “AI Chatbots Lose Money Every Time You Use Them. That Is a Problem,” *Washington Post*, June 5, 2023, <https://www.washingtonpost.com/technology/2023/06/05/chatgpt-hidden-cost-gpu-compute/>.
64. charging your phone fully: Melissa Heikkilä, “Making an Image with Generative AI Uses as Much Energy as Charging Your Phone,” *MIT Technology Review*, December 1, 2023, <https://www.technologyreview.com/2023/12/01/1084189/making-an-image-with-generative-ai-uses-as-much-energy-as-charging-your-phone/>; Sasha Luccioni, Yacine Jernite, and Emma Strubell, “Power Hungry Processing: Watts Driving the Cost of AI Deployment?,” *FACCT ’24: The 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro: ACM, 2024), 85–99, <https://doi.org/10.1145/3630106.3658542>.
65. 500 milliliters of water: Singh, “Every Time You Talk to ChatGPT”; Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren, “Making AI Less ‘Thirsty’: Uncovering and Addressing the Secret Water Footprint of AI Models,” *arXiv*, April 6, 2023, <https://arxiv.org/abs/2304.03271>.
66. than just returning links: Allison Parshall, “What Do Google’s AI Answers Cost the Environment?,” *Scientific American*, June 11, 2024, <https://www.scientificamerican.com/article/what-do-googles-ai-answers-cost-the-environment/>.
67. The *Washington Post* reported: Olivo, “Internet Data Centers.”
68. it’s telling that: Parshall, “What Do Google’s AI Answers?”
69. compared to 2019: Brad Smith and Melanie Nakagawa, “Our 2024 Environmental Sustainability Report,” Microsoft, May 15, 2024, <https://blogs.microsoft.com/on-the->

- issues/2024/05/15/microsoft-environmental-sustainability-report-2024/; Alexa St. John, “Google Falling Short of Important Climate Target, Cites Electricity Needs of AI,” Associated Press, July 2, 2024, <https://apnews.com/article/climate-google-environmental-report-greenhouse-gases-emissions-3ccf95b9125831d66e676e811ece8a18>.
70. Microsoft president Brad Smith told Bloomberg: Akshat Rathi and Dina Bass, “Microsoft’s AI Push Imperils Climate Goal as Carbon Emissions Jump 30%,” Bloomberg, May 15, 2024, <https://www.bloomberg.com/news/articles/2024-05-15/microsoft-s-ai-investment-imperils-climate-goal-as-emissions-jump-30>.
71. dismiss anything else as “less existential”: Geoff Hinton resigned from Google in June 2023 and then took his fainting couch on a tour of the news media to sound the alarm about existential risk from AI. When asked by CNN’s Jake Tapper whether he wished that he stood behind whistleblowers like Timnit Gebru more, Hinton replied, “Their concerns aren’t as existentially serious as the idea of these things getting more intelligent than us and taking over.” See Hinton, “‘Godfather of AI.’”
72. letter as a distraction: Timnit Gebru, Emily M. Bender, Angelina McMillan-Major, and Margaret Mitchell, “Statement from the Listed Authors of Stochastic Parrots on the ‘AI Pause’ Letter,” DAIR Institute, March 31, 2023, <https://www.dair-institute.org/blog/letter-statement-March2023/>.
73. we should instead join forces: Emily discusses these calls in a tweet thread: Emily M. Bender (@emilymbender), “To all those folks asking why the ‘AI safety’ and ‘AI ethics’ crowd can’t find common ground—it’s simple: The ‘AI safety’ angle, which takes ‘AI’ as something that is to be ‘raised’ to be ‘aligned’ with actual people is anathema to ethical development of the technology,” Twitter, April 2, 2023, <https://x.com/emilymbender/status/1642714011988004864>. Philosopher Seth Lazar has explicitly called for building bridges on X/Twitter, but it is not clear what session he is referring to: Seth Lazar (@sethlazar), “Couldn’t agree more with Kristian on this. We’ll be holding a session on Bridging AI Ethics and Safety at FAccT this year; also if you’re interested here’s my handout for a talk arguing a similar point about risks of understating capabilities of LLMs: <https://write.as/sethlazar/genb>,” Twitter, May 3, 2023, <https://twitter.com/sethlazar/status/1653907050081177601>.
74. AI booster Gary Marcus: Gary Marcus, “I am not afraid of robots. I am afraid of people,” Marcus on AI, April 2, 2023, <https://garymarcus.substack.com/p/i-am-not-afraid-of-robots-i-am-afraid>.
75. make copies of themselves or “self-replicate”: “Fact Sheet: Biden–Harris Administration Secures Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI,” White House, July 21, 2023, <https://web.archive.org/web/20250106022936/https://www.whitehouse.gov/briefing-room/statements-releases/2023/07/21/fact-sheet-biden-harris-administration-secures-voluntary-commitments-from-leading-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai/>; “Ensuring Safe, Secure, and Trustworthy AI,” White House, accessed August 8, 2024, <https://web.archive.org/web/20250117235102/https://www.whitehouse.gov/wp-content/uploads/2023/07/Ensuring-Safe-Secure-and-Trustworthy-AI.pdf>; podcast episode critiquing this: Emily M. Bender and Alex Hanna, “Episode 15: The White House and Big Tech Dance The Self-Regulation Tango,” *Mystery AI Hype Theater 3000*, podcast audio, September 20, 2023, <https://www.buzzsprout.com/2126417/13612292-episode-15-the-white-house-and-big-tech-dance-the-self-regulation-tango>.
76. 2023 executive order: Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, White House, October 30, 2023, <https://web.archive.org/web/20241115052308/https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

77. Blueprint for an AI Bill of Rights: Despite the unfortunate name, this is neither actual legislation nor about the rights of “AIs”. Instead, it’s a proposal from the Biden administration about how to protect people’s rights in the face of these technologies. See White House Office of Science and Technology Policy, “Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People,” October 2022, <https://web.archive.org/web/20250120000129/https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.
78. AI Safety Institute: “U.S. Artificial Intelligence Safety Institute,” National Institute of Standards and Technology, U.S. Department of Commerce, accessed August 8, 2024, <https://www.nist.gov/aisi>.
79. Paul Christiano: Ashley Belanger, “Feds Appoint ‘AI Doomer’ to Run AI Safety at US Institute,” *Ars Technica*, April 17, 2024, <https://arstechnica.com/tech-policy/2024/04/feds-appoint-ai-doomer-to-run-us-ai-safety-institute/>.
80. the report that resulted: Chuck Schumer, Mike Rounds, Martin Heinrich, and Todd Young, “Driving U.S. Innovation in Artificial Intelligence: A Roadmap for Artificial Intelligence Policy in the United States Senate,” May 2024, https://www.schumer.senate.gov/imo/media/doc/Roadmap_Electronic1.32pm.pdf.
81. automation sold as AI: AI Now Institute and partners released a shadow report in response to Schumer’s report, summarizing a decade of research in these areas: Accountable Tech, AI Now et al., “Put the Public Back in the Driver’s Seat: Shadow Report to the US Senate AI Policy Roadmap,” May 2024, <https://senateshadowreport.com/>.
82. as the status quo: Cynthia Conti-Cook, of the Surveillance Resistance Lab, makes this point. See Justin Hendrix, “Reading the ‘Shadow Report’ on US AI Policy,” *Tech Policy Press*, May 26, 2024, <https://www.techpolicy.press/reading-the-shadow-report-on-us-ai-policy/>.

Chapter 7: Do You Believe in Hope After Hype?

1. the company Hippocratic AI: This advertising copy is taken from the Internet Archive snapshot of their web page from July 10, 2024: <https://web.archive.org/web/20240710232102/https://www.hippocraticai.com/>. The video of the “Linda” agent is at “Linda,” Hippocratic AI, accessed August 9, 2024, <https://www.hippocraticai.com/linda>. Hippocratic links to an arXiv paper about an LLM called “Polaris”, trained on health care data. See Subhabrata Mukherjee et al., “Polaris: A Safety-Focused LLM Constellation Architecture for Healthcare,” arXiv, March 20, 2024, <https://arxiv.org/pdf/2403.13313>.
2. Michelle Mahon: Michelle Mahon, interview with Emily M. Bender and Alex Hanna, “Episode 37: Chatbots Aren’t Nurses (feat. Michelle Mahon),” *Mystery AI Hype Theater 3000*, podcast audio, July 22, 2024, <https://www.buzzsprout.com/2126417/15517978>.
3. Harrisburg University of Science and Technology: “HU Facial Recognition Software Predicts Criminality,” Harrisburg University of Science and Technology, archived May 6, 2020, <https://archive.is/N1HVe>. The Harrisburg University study ultimately wasn’t published, after pushback from a collective of researchers called the Coalition for Critical Technology: Coalition for Critical Technology, “Abolish the #TechToPrisonPipeline,” Medium, June 23, 2020, <https://medium.com/@CoalitionForCriticalTechnology/abolish-the-techtoprisonpipeline-9b5b14366b16>. For more on how image processing has brought about a modern rebirth of the pseudoscience of physiognomy, see Sidney Fussell, “An Algorithm That ‘Predicts’ Criminality Based on a Face Sparks a Furor,” *Wired*, June 24, 2020, <https://www.wired.com/story/algorithm-predicts-criminality-based-face-sparks-furor/>.

4. researchers at Shanghai Jiao Tong University: Xiaolin Wu and Xi Zhang, “Responses to Critiques on Machine Learning of Criminality Perceptions,” arXiv, November 13, 2016, <https://arxiv.org/pdf/1611.04135>.
5. by looking at them: Luke Stark and Jevan Hutson make this point forcefully, and state that these technologies need to be abolished outright: “Physiognomic AI is unjust in principle and in practice. It should be banned for all intents and purposes such that it is as legally and politically unpalatable as it is morally.” Luke Stark and Jevan Hutson, “Physiognomic Artificial Intelligence,” *Fordham Intellectual Property, Media & Entertainment Law Journal* 32, no. 4 (2022): 922–78.
6. Morehouse College: Wilborn P. Nobles III, “Morehouse to Use AI Teaching Assistants This Fall,” Axios, July 4, 2024, <https://www.axios.com/local/atlanta/2024/07/04/morehouse-ai-teaching-assistants>.
7. claimed that GPT-5: Mira Murati, “AI Everywhere: Transforming Our World, Empowering Humanity,” Dartmouth Engineering, YouTube video, 51:41, June 19, 2024, <https://www.youtube.com/watch?v=yUoj9B8OpR8>.
8. “ShotSpotter coverage area within 60 seconds”: “Introducing the SafetySmart Platform,” SoundThinking, accessed August 15, 2024, <https://www.soundthinking.com/safetysmart-platform/>.
9. “of all reported gunfire incidents”: “ShotSpotter Frequently Asked Questions,” SoundThinking, accessed August 15, 2024, <https://www.soundthinking.com/faqs/shotspotter-faqs/>.
10. false alarms: The Chicago audit was in 2021 and the New York City one in 2024. Fola Akinnibi, “NYC Surveillance Tech on Shootings Gives False Alarms 87% of Time, Audit Finds,” Bloomberg, June 20, 2024, <https://www.bloomberg.com/news/articles/2024-06-20/nyc-shotspotter-tech-wastes-nypd-police-time-and-money-audit-finds>.
11. a recipe for police violence: In fact, there is at least one case of the police murdering someone, thirteen-year-old Adam Toledo, while responding to a ShotSpotter alert: Max Blaisdell and Jim Daley, “ShotSpotter Keeps Listening for Gunfire After Contracts Expire,” *Wired*, April 24, 2024, <https://www.wired.com/story/shotspotter-keeps-listening-contracts-expire/>. There is also at least one documented case of police shooting at (and fortunately not injuring anyone, this time) a teenager who was setting off fireworks: Andrew Ramos and Todd Feurer, “CPD Officer Responding to ShotSpotter Opened Fire on Teen Setting Off Fireworks, Video Shows,” CBS News, February 28, 2024, <https://www.cbsnews.com/chicago/news/copa-body-camera-video-auburn-gresham-shots-fired-fireworks-shotspotter/>.
12. the poor get poorer: Sociologists call this the “Matthew effect”, coined by sociologists Robert K. Merton and Harriet Anne Zuckerman: Robert K. Merton, “The Matthew Effect in Science,” *Science* 159, no. 3810 (1968), <https://www.science.org/doi/10.1126/science.159.3810.56>; Robert K. Merton, “The Matthew Effect in Science, II: Cumulative Advantage and the Symbolism of Intellectual Property,” *Isis* 79, no. 4 (December 1988), <https://www.journals.uchicago.edu/doi/10.1086/354848>. Merton acknowledges in the latter paper (607, footnote 2) that the first paper should have had joint authorship with Zuckerman.
13. Burke instructs journalists: Gretel Kahn, “Focus on the Humans, Not the Robots: Tips from the Author of AP Guidelines on How to Cover AI,” Reuters Institute, September 5, 2023, <https://reutersinstitute.politics.ox.ac.uk/news/focus-humans-not-robots-tips-author-ap-guidelines-how-cover-ai>.
14. In her training, Hao encourages: Andrew Deck, “Pulitzer’s AI Spotlight Series Will Train 1,000 Journalists on AI Accountability Reporting,” NiemanLab, May 2, 2024, <https://www.niemanlab.org/2024/05/pulitzers-ai-spotlight-series-will-train-1000-journalists-on-ai-accountability-reporting/>.
15. Many of the proposed use cases of LLMs: Donald Metzler, Yi Tay, Dara Bahri, and Marc Najork, “Rethinking Search: Making Domain Experts Out of Dilettantes,” *ACM SIGIR Forum*

- 55, no. 1 (2021): 1–27, <https://dl.acm.org/doi/10.1145/3476415.3476428>.
16. making information access “frictionless”: Emily has written about this elsewhere, together with Chirag Shah, in an op-ed: Emily M. Bender and Chirag Shah, “All-Knowing Machines Are a Fantasy,” *IAI News*, December 13, 2022, <https://iai.tv/articles/all-knowing-machines-are-a-fantasy-auid-2334>; and in two academic papers: Chirag Shah and Emily M. Bender, “Situating Search,” *CHIIR ’22: Proceedings of the 2022 Conference on Human Information Interaction and Retrieval* (March 14, 2022): 221–32, <https://dl.acm.org/doi/10.1145/3498366.3505816>; and Chirag Shah and Emily M. Bender, “Envisioning Information Access Systems: What Makes for Good Tools and a Healthy Web?,” *ACM Transactions on the Web* 18, no. 3 (2024): 1–24, <https://dl.acm.org/doi/10.1145/3649468>.
 17. pushing chatbots as information access systems: A lot of these are “retrieval augmented generation” systems, but that isn’t actually much better. Companies selling this technology portray it as more accurate, because it involves first doing an ordinary web search, then inputting the returned pages into the LLM, along with a prompt designed to cause it to output an answer based on those pages. The products also return the web pages (that is, the original search results). However, there’s no guarantee that the LLM output accurately represents what’s on the pages, while at the same time fluent and authoritative-sounding strings tend to discourage people from clicking through to check.
 18. demo contained an error: James Vincent, “Google’s AI Chatbot Bard Makes Factual Error in First Demo,” *The Verge*, February 2, 2023, <https://www.theverge.com/2023/2/8/23590864/google-ai-chatbot-bard-mistake-error-exoplanet-demo>; “2M1207 b—First Image of an Exoplanet,” NASA, accessed September 16, 2024, <https://science.nasa.gov/resource/2m1207-b-first-image-of-an-exoplanet/>.
 19. service-centered ideology of access: Anna Lauren Hoffmann and Raina Bloom, “Digitizing Books, Obscuring Women’s Work: Google Books, Librarians, and Ideologies of Access,” *Ada* 9 (May 2016), <https://scholarsbank.uoregon.edu/xmlui/handle/1794/26769> As we discussed in Chapter 5, AI boosters often falsely sell their technology as “democratizing” something. Hoffmann and Bloom get into the community-level details of what democracy means and how libraries are in fact essential to supporting it (in a way that neither search engines nor any other kind of mathy math ever could). Democracy means public, open spaces where civic workers are there to help members of the public navigate to and make sense of information that they need. Democracy means connected communities where people work with one another to understand the world around themselves. A search engine can be a useful tool, but Google’s misrepresentation of what they can provide goes back at least to their conception of their mission as “organiz[ing] the world’s information and mak[ing] it universally accessible and useful.” As Hoffmann and Bloom convincingly argue, the very framing of that goal as something that can be achieved by one company, with impersonal technology, is inconsistent with meaningful access for many (most?) of the people in the world.
 20. threats of budget cuts: For example, the New York City Council proposed a \$53.8 million cut to the budgets of the Brooklyn, Queens, and New York public libraries for fiscal year 2025. See Nia Clark, “Public libraries demand reversal of proposed cuts,” *New York City Council*, March 26, 2024, <https://council.nyc.gov/shahana-hanif/2024/03/26/public-libraries-demand-reversal-of-proposed-cuts/>. Massive public outcry led to the reversal of this decision: NYPL Staff, “You Did It, NYC! Library Budgets Have Been Saved,” *New York Public Library*, July 1, 2024, <https://www.nypl.org/blog/2024/07/01/you-did-it-nyc-library-budgets-have-been-saved>.
 21. Maggie Tokuda-Hall: Maggie Tokuda-Hall, interview with Annalee Newitz and Charlie Jane Anders, “Fascism and Book Bans (with Maggie Tokuda-Hall),” *Our Opinions Are Correct*, podcast audio, April 18, 2024, 56:58, <https://www.ouropinionsarecorrect.com/shownotes/2024/4/18/fascism-and-book-bans-with-maggie-tokuda-hall>.

22. several other federal agencies: Lina M. Khan, “Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems,” Federal Trade Commission, April 25, 2023, https://www.ftc.gov/system/files/ftc_gov/pdf/EEOC-CRT-FTC-CFPB-AI-Joint-Statement%28final%29.pdf.
23. Companies cannot legally claim: Michael Atleson, “Keep Your AI Claims in Check,” Federal Trade Commission, February 27, 2023, <https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check>.
24. has put it: Michael Atleson, “Succor Borne Every Minute,” Federal Trade Commission, June 11, 2024, <https://www.ftc.gov/business-guidance/blog/2024/06/succor-borne-every-minute>.
25. eye on their workers: Alvin Velazquez and Muyi Zhang, “Labor Laws and Surveillance in the Time of COVID-19: A Demand for Better Worker Protections,” *ABA Journal of Labor Employment* 38, no. 1 (July 2024), https://www.americanbar.org/content/dam/aba/publications/aba_journal_of_labor_employment_law/v38/no-1/jlel-v38-n1-velazquez-zhang-5.pdf.
26. cameras on Amazon delivery trucks: Wilneida Negrón, “Little Tech Is Coming for Workers: A Framework for Reclaiming and Building Worker Power,” Coworker.org, 2021, <https://home.coworker.org/wp-content/uploads/2021/11/Little-Tech-Is-Coming-for-Workers.pdf>; Ifeoma Ajunwa, Kate Crawford and Jason Schultz, “Limitless Worker Surveillance,” *California Law Review* 105, no. 3 (June 2017): 735–76, <https://www.jstor.org/stable/44630759>; Lauren Kaori Gurley, “Amazon’s AI Cameras Are Punishing Drivers for Mistakes They Didn’t Make,” Vice News, September 20, 2021, <https://www.vice.com/en/article/88npjv/amazons-ai-cameras-are-punishing-drivers-for-mistakes-they-didnt-make>.
27. issued guidance: “NLRB General Counsel Issues Memo on Unlawful Electronic Surveillance and Automated Management Practices,” news release, National Labor Relations Board, Office of Public Affairs, October 13, 2022, <https://www.nlr.gov/news-outreach/news-story/nlr-general-counsel-issues-memo-on-unlawful-electronic-surveillance-and>.
28. labor rules: For instance, a U.S. appeals court threw out a workplace surveillance ruling made by the NLRB where a pro-union truck driver was ordered not to cover a surveillance camera in his truck: Daniel Wiessner, “NLRB Ruling on Worker Camera Surveillance Was ‘Nonsense,’ DC Circuit Says,” Reuters, March 26, 2024, <https://www.reuters.com/legal/government/nlr-ruling-worker-camera-surveillance-was-nonsense-dc-circuit-says-2024-03-26/>.
29. AI governance framework: Accountable Tech, AI Now, and Epic.Org, “Zero Trust AI Governance,” August 2023, <https://accountabletech.org/wp-content/uploads/Zero-Trust-AI-Governance.pdf>.
30. named so because it asserts: Personal communication with Amba Kak and Sarah West, June 19, 2024. More on zero-trust architecture can be found in Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly, “Zero Trust Architecture,” National Institute of Standards and Technology, August 2020, <https://doi.org/10.6028/NIST.SP.800-207>.
31. might harm innovation: “Yet recent calls in the AI space have sought to expand the scope of the regulation. . . . This could cause huge headaches for the innovators trying to ensure that AI technology evolves in a safe way.” David Alexandru Timis, “How to Regulate AI Without Stifling Innovation,” World Economic Forum, June 26, 2023, <https://www.weforum.org/agenda/2023/06/how-to-regulate-ai-without-stifling-innovation/>.
32. Oren Etzioni: Oren Etzioni, “AI in Public Sector: Tool for Inclusion or Exclusion?,” YouTube video, Paul G. Allen School, 01:26:23, May 15, 2018, https://www.youtube.com/watch?v=XwUr4xk-_hA.
33. misleadingly advertised “Full Self-Driving” technology: The *Washington Post* analyzed National Highway Traffic Safety Administration data for 2019–23 and found 736 crashes involving Tesla’s “Autopilot” and “Full Self-Driving” systems, including fatalities: Faiz Siddiqui and Jeremy B. Merrill, “17 Fatalities, 736 Crashes: The Shocking Toll of Tesla’s

- Autopilot,” *Washington Post*, June 10, 2023, <https://www.washingtonpost.com/technology/2023/06/10/tesla-autopilot-crashes-elon-musk/>. In 2024, Business Insider reported on how Tesla prioritizes data from high-profile drivers (such as those who make YouTube videos about their Tesla Autopilot experiences) in training their systems, throwing into relief the extent to which Tesla is having drivers do their beta-testing, without of course any consent from those around them on the road. Grace Kay, “Tesla Prioritizes Musk’s and Other ‘VIP’ Drivers’ Data to train Self-Driving Software,” *Business Insider*, June 9, 2024, <https://www.businessinsider.com/tesla-prioritizes-musk-vip-data-self-driving-2024-7>.
34. half a dozen research groups: See section 2.1 of Angelina McMillan-Major, Emily M. Bender, and Batya Friedman, “Data Statements: From Technical Concept to Community Practice,” *ACM Journal on Responsible Computing* 1, no. 1 (March 2024), <https://dl.acm.org/doi/pdf/10.1145/3594737>, and section 2.1 of Milagros Miceli et al., “Documenting Computer Vision Datasets: An Invitation to Reflexive Data Practices,” *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (March 1, 2021): 161–72, <https://dl.acm.org/doi/pdf/10.1145/3442188.3445880>.
 35. biases in their output: Buolamwini and Gebru, “Gender Shades”; Aylin Caliskan et al., “Semantics Derived Automatically from Language Corpora Contain Human-like Biases,” *Science*, April 14, 2017, <https://www.science.org/doi/10.1126/science.aal4230>; Tolga Bolukbasi et al., “Man Is to Computer Programmer as Woman Is to Homemaker? Debiasing Word Embeddings,” *Advances in Neural Information Processing Systems* 29 (2016), https://proceedings.neurips.cc/paper_files/paper/2016/hash/a486cd07e4ac3d270571622f4f316ec5-Abstract.html.
 36. the patterns of the past: Cognitive scientist Abeba Birhane makes the point that using machine learning to affect the future is an inherently conservative enterprise, in the sense that it means we’re okay with using the patterns of the past to set the future. See Abeba Birhane, “The Impossibility of Automating Ambiguity,” *Artificial Life* 27, no. 1 (2021): 44–61, <https://direct.mit.edu/artl/article-abstract/27/1/44/101872/The-Impossibility-of-Automating-Ambiguity>.
 37. attitude of care toward datasets: This point is emphasized in Emily’s “Stochastic Parrots” paper. See also Eun Seo Jo and Timnit Gebru, “Lessons from Archives: Strategies for Collecting Sociocultural Data in Machine Learning,” *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (January 27, 2020): 306–16, <https://dl.acm.org/doi/abs/10.1145/3351095.3372829>; Morgan Klaus Scheuerman, Alex Hanna, and Emily Denton, “Do Datasets Have Politics? Disciplinary Values in Computer Vision Dataset Development,” *Proceedings of the ACM on Human-Computer Interaction* 5, no. CSCW2 (2021): 1–37.
 38. their chatbot “learned” Bengali: Scott Pelley, “Is Artificial Intelligence Advancing Too Quickly? What AI Leaders at Google Say,” *CBS News*, April 16, 2023, <https://www.cbsnews.com/news/google-artificial-intelligence-future-60-minutes-transcript-2023-04-16/>. Researcher Meg Mitchell quickly dispelled this falsehood by pointing to Google’s own paper on the model. See Pritam Bordoloi, “Did Google Bard Really Learn Bengali on Its Own?,” *Analytics India Mag*, April 18, 2023, <https://analyticsindiamag.com/ai-origins-evolution/did-google-bard-really-learn-bengali-on-its-own/>.
 39. Hugging Face bucks the trend here: “Model Cards,” Hugging Face, accessed August 28, 2024, <https://huggingface.co/docs/hub/en/model-cards>; “Dataset Cards,” Hugging Face, accessed August 28, 2024 <https://huggingface.co/docs/hub/en/datasets-cards>.
 40. a long way to go: Weixin Liang et al., “What’s Documented in AI? Systematic Analysis of 32K AI Model Cards,” *arXiv*, February 7, 2024, <https://arxiv.org/pdf/2402.05160v1>.
 41. Meta brags about releasing “open source” models: Kyle Wiggers and Devin Coldewey, “This Week in AI: When ‘Open Source’ Isn’t So Open,” *TechCrunch*, April 20, 2024,

- <https://techcrunch.com/2024/04/20/this-week-in-ai-when-open-source-isnt-so-open/>; David Gray Widder, Sarah West, and Meredith Whittaker, “Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI,” *SSRN*, 2023, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4543807.
42. have published their own AI registers: “Artificial intelligence systems of Helsinki,” City of Helsinki AI Register, accessed January 22, 2025, <https://ai.hel.fi/en/ai-register/>; “Algorithmic Systems of Amsterdam,” City of Amsterdam Algorithmic Register, accessed January 22, 2025, <https://algoritmeregister.amsterdam.nl/en/ai-register/>.
 43. the toothless “voluntary commitments”: “Ensuring Safe, Secure, and Trustworthy AI,” White House.
 44. more durable watermarks for text: John Kirchenbauer and colleagues proposed a system that divides the vocabulary of the language model into “red” and “green” words at each step, and then increases the probability of the language model emitting “green” words. Text generated in this way will have a higher-than-chance proportion of “green” words, but otherwise still sound very similar to ordinary language model output. The allocation of words to the “red” and “green” sets can be done in a way that is publicly calculatable off of the text itself, or based on a private seed for the randomization. See John Kirchenbauer et al., “A Watermark for Large Language Models,” in *Proceedings of the 40th International Conference on Machine Learning*, ed. Andreas Krause et al., vol. 202, *Proceedings of Machine Learning Research* (PMLR, 2023), 17061–84, <https://proceedings.mlr.press/v202/kirchenbauer23a.html>.
 45. unintentional watermarks: Maiberg, “Scientific Journals.”
 46. Weizenbaum saw this risk already: Weizenbaum, *Computer Power and Human Reason*, 227.
 47. IBM training presentation: Unfortunately, the full presentation is not available. The image of this slide that frequently circulates online is from Twitter user @bumblebike, who posted it in 2017. See Bumblebike @bumblebike, “This is from a 1979 presentation. We are v slow learners, it seems,” Twitter, February 17, 2017, 12:59 a.m., <https://x.com/bumblebike/status/1484216311802736640>. They apparently found it among their father’s things, but when they went back many years later to try to find the rest of the binder, it had been damaged in a flood. It is described as being from an internal IBM training document. Bumblebike @bumblebike, “Yes, it was an internal IBM training document. I still have it around here somewhere,” Twitter, August 10, 2019, <https://x.com/bumblebike/status/1160221379880341505>. Several internet sleuths have worked on reading the palimpsest produced in the photo of the original slide by text from other pages bleeding through and have also collected what further information is available: Hungry Mouth @a-hungry-mouth, “chosting my reply because i believe i know people who might want to see it / correct it,” cohost, January 2024, <https://cohost.org/a-hungry-mouth/post/4039145-chosting-my-reply-be>; mhoye(@mhoye@mastodon.social), “We’ve all seen the IBM presentation saying, a computer can never be held accountable therefore . . . but if you rotate and adjust the gamma settings on the best versions of that image on the internet you can see further text through the page,” Mastodon, May 17, 2024, <https://mastodon.social/@mhoye/112459155499812117>.
 48. European Union’s AI Act: “High-Level Summary of the AI Act,” EU Artificial Intelligence Act, February 27, 2024, <https://artificialintelligenceact.eu/high-level-summary/>. See also a criticism of this approach in the AI Act: Amba Kak et al., “General Purpose AI Poses Serious Risks, Should Not Be Excluded from the EU’s AI Act,” AI Now Institute, April 13, 2023, <https://ainowinstitute.org/publication/gpai-is-high-risk-should-not-be-excluded-from-eu-ai-act>.
 49. “foundation models”: Rishi Bommasani et al., “On the Opportunities and Risks of Foundation Models,” arXiv, 2021, <https://crfm.stanford.edu/assets/report.pdf>.
 50. Another slide: bumblebike (@bumblebike), “It’s in a binder that I need to find, will do ASAP for you. Another gem in there was,” Twitter, August 10, 2019, <https://x.com/bumblebike/status/1160292780574298112>.

51. web crawling at all: David Pierce, “The Text File That Runs the Internet,” *The Verge*, February 14, 2024, <https://www.theverge.com/24067997/robots-txt-ai-text-file-web-crawlers-spiders>; Shayne Longpre et al., “Consent in Crisis: The Rapid Decline of the AI Data Commons,” arXiv, July 20, 2024, <https://www.dataprovenance.org/consent-in-crisis-paper>; “Who Blocks OpenAI, Google AI and Common Crawl?,” *palewire*, accessed August 28, 2024, <https://palewi.re/docs/news-homepages/openai-gptbot-robotstxt.html>.
52. to their data: “Chapter Three: Rights of the Data Subject,” General Data Protection Regulation, accessed August 28, 2024, <https://gdpr-info.eu/chapter-3/>.
53. it is necessary: The American Data Privacy and Protection Act was introduced in the 117th (2021–22) Congress with two Republican and one Democratic sponsors, but it has not been taken up in subsequent Congresses.
54. several states have laws: “Privacy Laws,” Electronic Privacy Information Center, accessed September 16, 2024, <https://epic.org/issues/privacy-laws/>.
55. opt out of data collection: “California Consumer Privacy Act (CCPA),” State of California Department of Justice, updated March 13, 2024, <https://www.oag.ca.gov/privacy/ccpa>.
56. data collected at all: “Biometric Information Privacy Act (BIPA),” ACLU Illinois, accessed August 28, 2024, <https://www.aclu-il.org/en/campaigns/biometric-information-privacy-act-bipa>.
57. feature called “Recall”: Tom Warren, “Microsoft’s All-Knowing Recall AI Feature Is Being Delayed,” *The Verge*, June 14, 2024, <https://www.theverge.com/2024/6/13/24178144/microsoft-windows-ai-recall-feature-delay>; Matt Burgess, “This Hacker Tool Extracts All the Data Collected by Windows’ New Recall AI,” *Wired*, June 4, 2024, <https://www.wired.com/story/total-recall-windows-recall-ai/>.
58. “Recall retrieves the moment you saw it”: “Manage Recall,” Microsoft, June 19, 2024, <https://learn.microsoft.com/en-us/windows/client-management/manage-recall>.
59. Optical Character Recognition: “Retrace Your Steps with Recall,” Microsoft Support, <https://support.microsoft.com/en-us/windows/retrace-your-steps-with-recall-aa03f8a0-a78b-4b3e-b0a1-2eb8ac48701c>, accessed August 2, 2024.
60. Satya Nadella talked up: Matt Burgess, “This Hacker Tool Extracts All the Data Collected by Windows’ New Recall AI,” *Wired*, June 4, 2024, <https://www.wired.com/story/total-recall-windows-recall-ai/>.
61. these tools are deployed: Timsit, “Hollywood Studios.”
62. the clinician’s office: National Nurses United, “A.I.’s Impact.”
63. *The Bear*, has said: Alex O’Keefe, interview with Sam Fragoso, “The Revolution Will Be Televised (with ‘The Bear’ Writer Alex O’Keefe),” *Talk Easy with Sam Fragoso*, podcast audio, accessed September 2, 2024, <https://talkeasypod.com/alex-okeefe/>.
64. set design: Gene Maddaus, “IATSE Sees Fears and Promise of Artificial Intelligence: ‘We Want the Spoils,’” *Variety*, April 3, 2024, <https://variety.com/2024/biz/news/iatse-artificial-intelligence-1235951357/>.
65. “Hot Labor Summer”: Sam Dean, “Corporate Greed, Low Unemployment, Housing Crisis: That’s the Recipe for Hot Labor Summer,” *Los Angeles Times*, August 11, 2023, <https://www.latimes.com/business/story/2023-08-11/strikes-unions-hot-labor-summer-los-angeles>.
66. private sector: “Union Members—2023,” Bureau of Labor Statistics, January 23, 2024, <https://www.bls.gov/news.release/pdf/union2.pdf>.
67. since the New Deal: Eli Rosenberg, “House Passes Bill to Rewrite Labor Laws and Strengthen Unions,” *Washington Post*, February 6, 2020, <https://www.washingtonpost.com/business/2020/02/06/house-passes-bill-rewrite-labor-laws-strengthen-unions/>.
68. employees of those companies: Veena B. Dubal, “Winning the Battle, Losing the War? Assessing the Impact of Misclassification Litigation on Workers in the Gig Economy,”

- Wisconsin Law Review* (2017), https://repository.uclawsf.edu/faculty_scholarship/1598/. In California, the ABC Test was created to protect against worker misclassification, but the passage of Proposition 22 in 2020 has prevented this test from taking effect. The California Supreme Court has held up the constitutionality of Proposition 22. See Guy Davidov and Pnina Alon-Shenker, “The ABC Test: A New Model for Employment Status Determination?,” *Industrial Law Journal* 51, no. 2 (June 2022): 235–76, <https://academic.oup.com/ilj/article-abstract/51/2/235/6563247>; Eli Tan, “In Win for Uber and Lyft, California Court Upholds Gig-Worker Proposition,” *New York Times*, July 25, 2024, <https://www.nytimes.com/2024/07/25/technology/california-gig-worker-court-decision.html>.
69. Te Hiku Media’s work: Keoni Mahelona et al., “OpenAI’s Whisper Is Another Case Study in Colonisation,” *papa reo* (blog), January 24, 2023, <https://blog.papareo.nz/whisper-is-another-case-study-in-colonisation/>.
 70. *right to be forgotten*: “Right to Erasure (‘Right to Be Forgotten’),” General Data Protection Regulation, accessed August 29, 2024, <https://gdpr-info.eu/art-17-gdpr/>.
 71. consent at a later point: The Consentful Tech Project has created a guide for which we can think about consent in technology. Consent should be FRIES: Freely Given, Reversible, Informed, Enthusiastic, and Specific. See “What Is Consentful Tech?,” Consentful Tech Project, accessed August 29, 2024, <https://www.consentfultech.io/>.
 72. “make justice irresistible”: Ruha Benjamin, *Viral Justice: How We Grow the World We Want* (Princeton, NJ: Princeton University Press, 2022), 11.
 73. disability justice advocates: James I. Charlton, *Nothing About Us Without Us: Disability Oppression and Empowerment* (Berkeley: University of California Press, 1998).
 74. Luddite movement views technology: Jathan Sadowski, “I’m a Luddite. You Should Be One Too,” *The Conversation*, August 9, 2021, <https://theconversation.com/im-a-luddite-you-should-be-one-too-163172>.
 75. full members of humanity: Audra Simpson, “On Ethnographic Refusal: Indigeneity, ‘Voice’ and Colonial Citizenship,” *Junctures* 9 (2007), <https://junctures.org/index.php/junctures/article/view/66>; Ruha Benjamin, “Informed Refusal: Toward a Justice-Based Bioethics,” *Science, Technology, & Human Values* 41, no. 6 (2016): 967–90, <https://www.jstor.org/stable/24778299>.
 76. Feminist Data Manifest-No: Marika Cifor et al., “Feminist Data Manifest-No.,” 2019, <https://www.manifestno.com/>.
 77. Elon Musk said: Robert Hart, “Elon Musk Predicts Tesla Self-Driving Cars Will Arrive ‘This Year,’” *Forbes*, July 6, 2023, <https://www.forbes.com/sites/roberthart/2023/07/06/elon-musk-predicts-tesla-self-driving-cars-will-arrive-this-year/>.
 78. Geoff Hinton proclaimed: “Geoff Hinton: On Radiology,” Creative Destruction Lab, YouTube video, 01:24, November 24, 2016, <https://www.youtube.com/watch?v=2HMPRXstSvQ&t=29s>.
 79. hundreds of crashes: Siddiqui and Merrill, “17 Fatalities, 736 Crashes.”
 80. has projected: “Radiologic and MRI Technologists,” *Occupational Outlook Handbook*, Bureau of Labor Statistics, U.S. Department of Labor, accessed August 29, 2024, <https://www.bls.gov/ooh/healthcare/radiologic-technologists.htm#tab-6>.
 81. has called this: Jenna Burrell, “Artificial Intelligence and the Ever-Receding Horizon of the Future,” *Tech Policy Press*, June 6, 2023, <https://www.techpolicy.press/artificial-intelligence-and-the-ever-receding-horizon-of-the-future/>.
 82. has critiqued: Anna Lauren Hoffmann, “Bringing Up the Bot: Child, Potential, and the Domesticated Subject of AI Ethics.” Society for Social Studies of Science (4S) Conference, Honolulu, HI, November 9, 2023.
 83. the landmark research: Buolamwini and Gebru, “Gender Shades.”
 84. never be made: Zoé Samudzi, “Bots Are Terrible at Recognizing Black Faces. Let’s Keep It That Way,” *Daily Beast*, February 8, 2019, <https://www.thedailybeast.com/bots-are-terrible-at>

recognizing-black-faces-lets-keep-it-that-way.

85. echoes this sentiment: “Community Defense: Sarah T. Hamid on Abolishing Carceral Technologies,” *Logic(s)* 11 (August 31, 2020), <https://logicmag.io/care/community-defense-sarah-t-hamid-on-abolishing-carceral-technologies/>.
86. encourage travelers: “TSA Is Expanding Its Facial Recognition Program. You Can Opt Out,” Algorithmic Justice League, accessed August 30, 2024, <https://www.ajl.org/campaigns/fly>.
87. at refugee camps: Molnár, *The Walls Have Eyes*, especially chapter 3.
88. investment giant Goldman Sachs: Hatzius et al., “The Potentially Large Effects of Artificial Intelligence on Economic Growth (Briggs/Kodnani)” ; Allison Nathan, Jenny Grimberg, and Ashley Rhodes, “Gen AI: Too Much Spend, Too Little Benefit?,” Goldman Sachs Global Macro Research, June 25, 2024, <https://www.goldmansachs.com/intelligence/pages/gs-research/gen-ai-too-much-spend-too-little-benefit/report.pdf>.
89. Daron Acemoglu: Daron Acemoglu, “The Simple Macroeconomics of AI,” National Bureau of Economic Research Working Paper, May 2024, <https://www.nber.org/papers/w32487>.
90. Alphabet CEO Sundar Pichai was grilled: Gerrit De Vynck, “Big Tech Says AI Is Booming. Wall Street Is Starting to See a Bubble,” *Washington Post*, July 24, 2024, <https://www.washingtonpost.com/technology/2024/07/24/ai-bubble-big-tech-stocks-goldman-sachs/>.
91. David Cahn at VC giant Sequoia Capital: David Cahn, “AI’s \$600B Question,” Sequoia Capital, June 20, 2024, <https://www.sequoiacap.com/article/ais-600b-question/>.
92. author and technology critic Cory Doctorow: Cory Doctorow, “What Kind of Bubble Is AI?,” *Locus*, December 18, 2023, <https://locusmag.com/2023/12/commentary-cory-doctorow-what-kind-of-bubble-is-ai/>.
93. Google and Microsoft have sheepishly admitted: St. John, “Google Falling Short”; Justine Calma, “Microsoft’s AI Obsession Is Jeopardizing Its Climate Ambitions,” *The Verge*, May 15, 2024, <https://www.theverge.com/2024/5/15/24157496/microsoft-ai-carbon-footprint-greenhouse-gas-emissions-grow-climate-pledge>.
94. Indigenous Americans: Elizabeth Paras and Emily Seymour, “Black Leaders Say Development Infringing on Historical Cemeteries in Prince William,” *Inside Nova*, June 27, 2024, https://www.insidenova.com/headlines/black-leaders-say-development-infringing-on-historical-cemeteries-in-prince-william/article_4a75e8f8-349d-11ef-894f-43fd29f4c462.html.
95. the hardware called H100s: All the tech companies are buying enormous amounts of processors, mostly from Nvidia, to run their so-called AI systems. For example, *PC Mag* reported in January 2024 that Meta planned to spend billions of dollars to purchase 350,000 H100 GPUs from Nvidia: Michael Kan, “Zuckerberg’s Meta Is Spending Billions to Buy 350,000 Nvidia H100 GPUs,” *PC Mag*, January 18, 2024, <https://www.pcmag.com/news/zuckerbergs-meta-is-spending-billions-to-buy-350000-nvidia-h100-gpus>.
96. OpenAI CTO Mira Murati: Alex Heath, “OpenAI Exec: ‘Some Creative Jobs Maybe Will Go Away, but Maybe They Shouldn’t Have Been There in the First Place,’” *The Verge*, June 21, 2024, <https://www.theverge.com/2024/6/21/24183265/openai-exec-some-creative-jobs-maybe-will-go-away-but-maybe-they-shouldnt-have-been-there-in-the-fir>.
97. over 10,000 layoffs: Merchant, “AI Is Already Taking Jobs in the Video Game Industry.”
98. being human together: We learned this turn of phrase from ethnomusicologist and music historian Gabriel Solis.

Index

A specific form of pagination for this digital edition has been developed to match the print edition from which the index was created. If the application you are reading this on supports this feature, the page references noted in this index should align. At this time, however, not all digital devices support this functionality. Therefore, we encourage you to please use your device's search capabilities to locate a specific entry.

- ableism, 33–36
- accelerationism, 147–148, 239
 - anticapitalist, 239
- accessibility, claims of, 68–69, 249
- accountability, 182–184, 253
- Accountable Tech, 176, 246, 250
- Acemoglu, Daron, 193, 257
- ACLU, 98, 199, 254
- acoustic models, 24
- Adams, Douglas, 117
- Adams, Eric, 72, 76
- advertising, 6, 50–51, 126–129, 132, 135, 211. *See also* marketing
 - false, 175, 191
- AdVon Commerce, 125–126, 130
- Adweek* (periodical), 132, 236
- Affordable Care Act (ACA), 85
- AFL-CIO, 137
- African Content Moderators Union, 65
- Agüera y Arcas, Blaise, x, 21, 22, 155–156
- Ahmed, Shazeda, 239
- AI (artificial intelligence)
 - authors' use of term, 5
 - as a con, 4
 - emoji signifying, 171
 - history of, 11–14
 - as a marketing term, 5, 7–8

- technologies sold as, 5–7
- unethical funding networks, 13
- AI Act (EU), 183
- AI bubble, xi, 164, 193–195
- AI hype
 - as a con, 4
 - FOMO and, 9, 13, 164, 195–196
 - as a funding incentive, 13–14
 - harms of, 14–17
 - history of, 11–14
 - hype cycles, 4
 - information literacy as resistance to, 171–174
 - libraries as resistance to, 173–174
 - promises of, 10, 14–15, 18
 - purposes of, 164–165
 - questioning, 164–170
 - reduction of the human condition by, 32–33, 206
 - refusing, 190–193
 - resisting, 170–171, 196
 - sentience claims, 22–23, 39–40
 - spotting, 17–20
 - worst-case scenario promotion, 10
 - See also* Boosterism; Doomerism
- AI Now, 176, 246, 253
- AI Overviews, 159, 173, 180, 227
- “AI Pause” letter, 140, 161–162
- AI registers, 181
- AI Safety Institute, 162
- AI safety research, 142–146, 149, 151, 238, 239, 245
- AI Safety Summit, 74, 145
- AI Scientist, 117–118, 120, 231–232
- AI winter, 7–8
- Alden Capital, 128
- Algorithmic Justice League, 192
- alignment problem, 138, 142–145
- Allegheny Family Screening Tool (AFST), 70, 216
- Allen Institute for Artificial Intelligence, 177
- Alphabet, 55, 193
- AltaVista, 51
- Altman, Sam, 9, 10, 23, 31–32, 52, 98, 141, 171, 242
- alt text, 108
- Amazon, 6, 46–47, 89, 104, 106, 126, 175
- Amazon Mechanical Turk (AMT), 60–61, 63, 65
- American Association for the Advancement of Science, 121, 220
- American Association of University Professors (AAUP), 226, 234
- American Data Privacy and Protection Act, 254
- American Servicemembers Protection Act (2002), 144
- Andreessen, Marc, 36, 38–39, 77, 146–148, 208
- Andreessen Horowitz, 38, 39, 77, 89, 112, 223
- Angwin, Julia, 71

Anthropic, 1, 13, 92, 138, 149
 anthropomorphization, 166–167, 182, 219
 AP Stylebook, 169
 Arena Group, 125
 art and artists, 103–112
 authors, 105–107, 111
 citational practice, 110
 content theft, 19, 31, 64, 104, 110, 135, 152
 copyright, 111–112, 229
 creativity and, 101–103, 135–136
 credit, consent, and compensation, 104–105
 functions of, 109
 social nature of, 109–110
 used for training AI, 107–109
 visual artists, 64–65, 103–105, 194
 artificial general intelligence (AGI), 22, 33–34, 36–39, 151, 155, 196
 artificial intelligence. *See* AI (artificial intelligence)
 Asilomar AI Principles, 143
 Asimov, Isaac, 10
 Associated Press (AP), 169, 238, 244
 asylum cases, 15, 72, 73–74
 Atleson, Michael, 175, 250
 austerity measures, 68, 69–71, 84, 85
 Authors Against Book Bans, 174
 Automating Inequality (Eubanks), 69, 216
 automatic control, 45–46
 automatic decision systems, 6, 14, 19, 69–71, 180–181
 automatic transcription, 6, 81–82
 automation, 6–7, 17, 54, 55, 66, 76, 102–103, 137, 139, 143, 165, 168, 169, 175, 184, 189, 190
 disclosure of, 180–182
 illusion of, 58
 imposed on workers, 18, 42–43, 47–49, 187, 193
 in journalism, 129–131
 in science, 113, 123
 Luddite resistance to, 43–47, 66
 to replace social services, 19, 69, 84–86
 avatars, 52, 89–90

Baidu, 8
 bail, 70–71
 Ballmer, Steve, 9
 Bankman-Fried, Sam, 13, 151
 Barbrook, Richard, 240
 Bard, 76, 92, 173, 252
 Bards and Sages, 105–106
The Bear (TV show), 187
The Bell Curve (Hernstein and Murray), 34
 benchmarking, 81–83, 199
 Bengio, Yoshua, 1, 138, 140, 205
 Benjamin, Ruha, 190, 255, 256

- Bentham, Jeremy, 152
- BERT system, 27, 205
- beta testing, 178, 251
- bias, 62, 76, 162, 178, 180
 - in chatbots, 27–28, 92, 126
 - in criminal justice, 71, 99, 165–166, 192, 217
 - in IQ tests, 35–36
 - in medicine, 92
- Biden, Joe, 161–162, 181, 246
- bigram language models, 23–24
- Big Tech, 13, 40, 55, 62, 88, 104, 120, 125, 127, 132, 159, 179, 191
- Bihar Prabha, 126–127
- Bill & Melinda Gates Foundation, 98
- Binet, Alfred, 35, 37
- Biometric Information Privacy Act (2008), 185
- Birhane, Abeba, 252
- birth rates, 37, 38
- Black people
 - child welfare system and, 69–70, 144
 - criminal justice system and, 71
 - data center construction impacts, 194
 - facial recognition tools and, 2–3, 192, 200
 - health care prediction systems and, 86–87
 - medical bias and, 92
 - racial intelligence claims about, 34–35
 - refusal by, 190–191
- BlackRock, 42
- blank page problem, 49
- Bloom, Raina, 174, 249
- Bloomberg, 129, 160, 219, 228, 244, 247
- Blueprint for an AI Bill of Rights (2022), 162
- book bans, 174
- Boosterism, 19–20, 138–139, 141, 142, 146–152, 153, 155, 160, 164, 174. *See also* Doomerism
- Borning, Alan, 157
- bossware, 175
- Bostrom, Nick, 140
- Bouie, Jamelle, 39, 208
- Bourgon, Malo, 138, 141
- Browder, Joshua, 77
- Bubeck, Sébastien, 21, 35, 36
- Buolamwini, Joy, 33, 192, 251
- Bureau of Labor Statistics, 128, 187, 191
- Burke, Garance, 169
- Burrell, Jenna, 191
- Buschek, Christo, 107–108, 200
- Business Insider (periodical), 200, 208, 251
- BuzzFeed, 130

- Cahn, David, 193
- California Consumer Privacy Act (2018), 185

- California Department of Tax and Fee Administration, 73
- California Ideology, 149, 240
- California Government Operations Agency, 73
- Cameroon, Andy, 240
- “Can Machines Learn How to Behave?” (Agüera y Arcas), x, 204
- capabilities, 12, 13, 78, 141, 162
- capitalism, 35, 42, 149, 159, 239
- carbon emissions, 158–160, 194
- Carceral Tech Resistance Network, 192
- Center for AI Safety, 140–141, 142
- Center for Democracy and Technology, 94
- Centre for Effective Altruism, 151
- Centre for the Study of Existential Risk, 142, 151
- Chai Research, 67, 90, 91
- charter schools, 95
- chatbots
 - anthropomorphizing, 166
 - ChatGPT (*see* ChatGPT)
 - city and state government, 72–73
 - disordered eating advice, 57, 91
 - ELIZA, 12–13, 28–29, 206
 - GPT-4, 21–22, 48, 62, 92, 158, 221
 - health services, 74, 89
 - LaMDA, 21, 155
 - language model-driven, 27
 - legal services, 76–78
 - legislation drafted with, 78, 79
 - OpenAI’s, 47–48 (*see also* ChatGPT)
 - psychotherapy, 67–68, 90–91
 - suicide encouragement, 67, 91
 - treated as data sources, 76–77, 172–173
 - See also* large language models (LLMs)
- ChatGPT, 7, 171
 - in the classroom, 15–16, 93–94, 225
 - development of, 27–28, 47
 - in legal practice, 16, 75, 78, 219, 221
 - in medical practice, 75
 - in peer review, 116–117
 - in policymaking, 78–79
 - in publishing, 105–107
 - racial bias, 92
 - in scientific papers, 121–122, 231
 - training, 111–112
 - in the workplace, 49–52, 65
- Chiang, Ted, 210–211, 228
- chess skill, 154–155
- child welfare system, 69–70, 144, 168–169
- Christian, Brian, 142
- Christiano, Paul, 162
- Chronicle of Higher Education* (periodical), 93, 96, 97

citational practice, 110
 The City (periodical), 72
 civil rights, 1, 145, 175, 178
 Clarkesworld (periodical), 105–106
 classification, 6, 60, 165, 192
 Claude, 92
 Clever Hans effect, 83, 221
 climate crisis, 14, 118, 139, 156–161, 194
 climate refugees, 118, 145, 160
 cloud computing, 157–158
 CNET, 130
 CNN, 128, 221, 236, 241, 245
 Coalition for Critical Technology, 247
 Cohen, Michael, 76–77
 Cold War, 12, 45
 colleges and universities, 93–94, 95–98, 119, 124, 234
 colonialism, 37, 215, 255
 computational linguistics, 25
 “Computing Machinery and Intelligence” (Turing), 153
 con, use of term, 4
 Congress of Industrial Organizations (CIO), 46. *See also* AFL-CIO
 Connected by Data, 75
 Conover, Adam, 54
 consciousness, 21–23, 25, 32, 39–40, 149
 consent, 10, 64, 65, 68, 104, 179, 185, 187, 190, 251, 255
 Consentful Tech Project, 255
 content mills, 128, 132
 content moderation, 63–64, 65
 Copilot, 52, 53
 copyright, 53, 110–112, 229
 corpora, 23–27, 31, 204
 Corrado, Greg, 88–89
 Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), 71
 cosmism, 150, 151
 Costanza-Chock, Sasha, 33
 Consumer Privacy Act (2018), 185
 COVID-19 pandemic, 15, 91, 175
 Craigslist, 127
 creativity, 4, 5, 14, 19, 43, 66, 84, 94, 101–103, 111, 120, 135–136
 credentialing systems, 80, 83–84
 criminal justice system, 70–71, 75, 99, 165–166, 192, 217
 critical harm, 141
 Crockett, M. J., 118–119
 crowdworking, 59–62, 65, 66
 Cruise, 55, 56, 58–59
 cybernetics, 11

 DALL-E, 7, 31, 59
 Dartmouth College, 11, 22, 191
 Data & Society, 137, 191, 212

- data minimization, 185–186
- data rights, 75, 184–186, 189–190
- dataset documentation, 178–180
- data workers, 59–61, 63–64, 65, 169
- Dawson, Julie Ann, 105–106
- Deadspin, 128, 134
- decision making, xi, 4, 6, 14, 18, 69–71, 153, 169, 183
- deepfakes
 - fashion industry, 57–58
 - pornography, 3
- deep learning, 7–8, 59, 61
- DeepMind, 8, 13, 88, 120, 149
- Defector Media, 134
- Defense Advanced Research Projects Agency (DARPA), 81
- Deloitte, 42
- digitization, 8, 84–85, 90
- disclosure, 180–182
- DNNresearch, 8
- Dockum, Rikker, 114
- Doctorow, Cory, 51, 193–194
- Documented, 72
- DoNotPay.com, 77
- Doomerism, 19–20, 138–146, 148–150, 152–153, 155, 156, 160, 161–162, 163, 174, 242. *See also* Boosterism
- doomsday scenarios, 1–4, 136, 138
- DoorDash, 55, 187
- Dotdash Meredith, 126
- Dowden, Oliver, 74
- dpchallenge.com, 108
- “Driving U.S. Innovation in Artificial Intelligence” (Schumer), 177
- Duke, David, 207
- Dungeons and Dragons (game), 104
- Dupré, Maggie Harrison, 125–126

- eBay, 127
- ecosystem metaphor, 102
- education, 14, 19, 68, 93–98, 124, 145, 184
 - AI harms, 15–16, 74
 - AI hype, 14
 - charter schools, 95
 - ChatGPT use, 93–94
 - classroom surveillance, 94–95
 - colleges and universities, 93–94, 95–98, 119, 124, 234
- effective accelerationists (e/acc), 148–149
- effective altruism, 19, 148, 150, 151
- electronic health records (EHRs), 86
- Electronic Privacy Information Center (EPIC), 176
- ELIZA, 12–13, 17, 28–29, 202, 206
- embeddings, 26–27
- emergency response systems, 82

empathy, 90, 91
energy consumption, 157–160
engineering. *See* science and engineering
enshittification, 51
enterprise sales, 51–52
environmental impacts, 157–159
Epic Systems, 86
epistemic community, 239
ethics, 17, 147, 168–169, 182, 245
Etzioni, Oren, 177
Eubanks, Virginia, 69, 216
eugenics, 23, 35, 36–39, 151–152, 153, 208
European Union, 48, 183, 185, 189
evaluation methodology, ix, 80–84, 85, 88, 167–168, 220
existential risk, 2, 39, 137–138, 141, 144–145, 146, 147, 160, 245
extropianism, 150

Facebook, 15, 65, 127, 133, 173, 230. *See also* Meta
facial recognition tools, 2–3, 166, 175, 192, 200
factories, 42–46, 60
fair use exception, 111–112, 229
family separation, 69–70, 144, 168–169
fashion industry, 57–58, 169
Federal Trade Commission (FTC), 57, 175, 183
Federated Data Platform, 75
Feminist Data Manifest-No, 191
fine-tuning, 108
First Law of Robotics, 10
Flowers, Jonathan, 109
Ford Motor Company, 45–46
Fortress Investment Group, 134
foundation models, 183
404 Media, 105, 122, 132, 134, 200, 214, 227, 229, 233, 235
Frankenstein (Shelley), 10
freedom-of-information requests, 74–75
friction, in information access, 171–173
Friedman, Batya, 157, 251
Frontiers in Cell and Developmental Biology (periodical), 121, 233
fruit fly (*Drosophila*), 154–155, 241
Future Fund, 151
Future of Life Institute, 140, 143, 151
Futurism, 125–126, 130, 223

Galactica, 16, 113–114, 115, 120, 230
gaming industry, 104, 194
Gannett, 128, 130–131
Garg, Nikhil, 122
GateHouse Media, 131
Gates, Bill, 95, 98, 141
Gawker, 128

Gaza, 3–4
Gebru, Timnit, x, 33, 150, 151, 192, 201, 232, 245, 251, 252
Gemini, 7, 76, 92, 127
General Catalyst, 89
General Data Protection Regulation (GDPR), 185, 189
General Motors, 55
General Purpose Technology (GPT), 47
generative AI, 7, 13, 40, 133, 134, 160, 175, 195. *See also* chatbots
generative language models, 27
Generative Pre-trained Transformer (GPT), 47
Genesis, 132
ghost work, 59–61
gig work, 4, 43, 54–56, 212
gig workers, 187–188
GitHub, 51–52, 53
Glass Health, 89
Glaze, 65
Goddard, Henry, 35
Goldman Sachs, 48–49, 193
G/O Media, 128, 134
Google, 8, 21, 27, 40, 92, 120, 127, 132, 149, 155, 158, 159, 160, 173, 174, 194, 227, 249, 252
Google Health, 88
Google I/O, 88
Google News Initiative, 132
Google Research, 88
Google Search, 51, 128, 129
“The Gospel” target selection system, 3–4
Gottfredson, Linda S., 34–35, 36, 207
government, 71–76, 168–169, 181
GPT-3, 67, 77, 158
GPT-4, 21–22, 48, 62, 92, 158
“GPTs are GPTs” (OpenAI), 47
Great Hills Partners, 128
Grok, 39
ground truth, 82, 220
Guardian (periodical), 32, 200, 202, 203, 210, 217, 243
Gutiérrez, Juan David, 78

Hacker News, 59
Hague Invasion Act (2002), 144
HAL 9000, 10, 140
hallucination, use of term, 167
Hamas, 3
Hamid, Sarah T., 192
Hammer, MC, 148
Hanania, Richard, 39, 208
Hao, Karen, 62, 63, 129, 170, 213, 241
Haraway, Donna, 232
Harder, Delmar S., 45
Harris, Kamala, 145

Harrisburg University of Science and Technology, 165, 247
 hate speech, 27–28, 59
 Health AI, 88
 health care divide, U.S., 222
 health care prediction systems, 87–88
 health insurance, 87–88
 health services, 74–75. *See also* medicine
 Helpline Associates United, 57
Her (film), 10
 Hernández, Andrea Paola, 63
 Herrnstein, Richard, 34
 Hinton, Geoff, 8, 61, 153, 191, 245
 Hippocratic AI, 89–90, 91, 165
The Hitchhiker's Guide to the Galaxy (Adams), 117
 Hoffman, Reid, 13
 Hoffmann, Anna Lauren, 174, 191, 249
 Hollywood, 41–43, 64, 110, 187
 Holocaust, 37, 144
 hosiers, 44–45
 Hugging Face, 179, 205
 Hulu, 41, 187
 human condition, 31–33, 91, 109, 136
 human subjects, 115, 123
 human values, 142–144
 Hutson, Jevan, 247
 Huxley, Aldous, 151
 Huxley, Julian, 151
 hype, 8–9. *See also* AI hype
 hype men, 8–9

IBM, 81, 179, 183, 184, 253
 image classification, 6, 60
 image generation, 3, 7, 18, 59, 88, 103–104
 image labeling algorithms, 26
 image localization, 60
 ImageNet, 60–61
 ImageNet Challenge, 61
 Indigeneity, 215
 Indigenous people

- child welfare system and, 69–70, 144
- data center construction impacts, 194
- data rights, 189
- ecological knowledge, 119, 124
- refusal by, 190–191

Industrial Light & Magic, 104
 Industrial Revolution, 18, 37, 43
 inequality, 17, 36, 69, 133
 Inflection AI, 13
 information ecosystem, 31, 50, 102, 122, 135, 162, 168, 173, 181
 information literacy, 164, 171–174

inputs vs. outputs, 165–166
Inside Higher Ed (periodical), 96, 231
Insight Forum, 1–2, 137–138, 162
intelligence
 general, 18, 22, 33–36, 155
 human, 22, 23
 machine, 7, 21–23
 measuring, 34, 153–155
 superintelligence, 142–143
 See also AI (artificial intelligence); artificial general intelligence (AGI)
Intergovernmental Panel on Climate Change, 156–157
International Committee of the Red Cross, 3
International Conference on Machine Learning, 121
International Criminal Court, 144
intersubjectivity, 30
Invisible Institute, 133
IQ (intelligence quotient) tests, 34–36, 37, 207
Irani, Lilly, 65
Israel Defense Forces (IDF), 3–4

Johansson, Scarlett, 10
Johnson, Khari, 73, 200
Johnson, Mike, 23
journalism, 19, 42, 102, 125–135, 169
 AI coverage, 169–170
 AI hype, ix, x, 135
 content mills, 128
 creativity and, 101–103, 135–136
 decline of, 127
 funding models, 133–134
 local, 127–128, 130–131
 nonprofit, 133
 revenue loss, 127–128, 132–133
 “shoe leather”, 129
 synthetic text extruding machines, 128–131
Journal of the American Medical Association (periodical), 86
Just Treatment, 75

Kaplan, Jared, 1, 138
Khan, Lina, 250
Khan, Sal, 93, 225
Khanna, Ro, 78, 79
Khosla, Vineet, 129
Kimmerer, Robin Wall, 232
Kirchenbauer, John, 252–253
Kitano, Hiroki, 117, 120, 231
Koebler, Jason, 134, 227, 233, 235
Koko, 67–68
Koren, Ariel, 74
Krizhevsky, Alex, 8

Kurzweil, Ray, 149, 151

labor strikes

coal miners', 46

Hollywood, 41–42, 54–55, 64

labor unions, 41–43, 57, 64, 65, 86, 110, 137, 170–171, 186–187, 188

LAION-5B, 107–108, 200

La Libre (periodical), 67

LaMDA, 21, 155

language

interpreting, 29–30

law and, 78–79

learning, 29–30

models, 30–31, 32

understanding, 28–29, 30, 82, 155

language modeling, 23–28

large language models (LLMs)

anthropomorphizing, 167

biases in, 27–28

carbon footprint, 157–159

dataset documentation, 178–180

development of, 23–27

evaluation methodology, 80–84

fabricated output, 16

monetization of, 50–52

plagiarism, 53

red-teaming, 62–63

as a research tool, 122–123

risks of, 124

training, 23–28, 53, 113–114

treated as data sources, 48, 76–77, 121–122, 172–173

Latour, Bruno, 114–115

Lazar, Seth, 238, 245

LeCun, Yann, 113–114

LedeAI, 131

Lee, Victor, 93

legal services, 14, 16, 42, 68, 76–80, 99

legislation, 20, 76, 78–79, 141, 174, 178–188, 246. *See also* regulation

Lemoine, Blake, 21

Lena, Jennifer, 109

Leontief, Wassily, 46

Levi's, 57–58

Lewis, Will, 129, 134

Li, Fei-Fei, 60, 61

libraries, 164, 171–174, 249

licensing processes, 83–84, 88, 91, 222

Lieu, Ted, 23

Lineup Publishing, 128

LinkedIn, 13, 52, 173

literature review, 97, 113, 114–115, 121, 123

Llama 3.1, 27
Local Reporting Network, 133
logic of domains, 232
Logler, Nick, 157
longtermism, 19, 150, 151–152
Lopez, John, 54
Luccioni, Sasha, 158, 159
Ludd, Ned, 45
Luddites, 43–45, 56, 64, 65, 66, 190
Lycos, 51
Lyft, 55, 56, 187

Machine Intelligence Research Institute, 138, 149
machine learning, 13, 14, 26, 107, 116, 118, 149, 178, 220, 252
machine translation, 6, 7, 15, 24, 26, 73–74, 81, 189
Mądry, Aleksander, 1, 138
Magic: The Gathering (game), 104
Mahon, Michelle, 165
Maiberg, Emanuel, 134, 200, 214, 234, 253
Majority World, 37, 43, 69, 91, 110, 126, 130, 143, 144, 152, 170, 209
Malthus, Thomas, 37
Manning, Chris, 206
Marcus, Gary, 161
marketing, 5, 7–8, 9, 95, 96, 118, 120, 125, 127, 132, 165. *See also* advertising
Markov, Andrey, 23
The Markup (periodical), 72
Martin, George R. R., 111
Marvel Studios, 104, 112
Massachusetts Institute of Technology (MIT), 12–13, 22, 193, 205
mass extinction event, 138, 139, 141, 142, 152
Matthew effect, 248
mathy maths, 5, 102, 174, 249
The Matrix (film), 10, 139
Maxwell, Sharon, 57
Mayo Clinic, 88
McCarthy, John, 11–12, 154, 191, 202, 241
McCulloch, Warren, 25
McKinsey, 40, 42, 88
McMillan-Major, Angelina, 215–216, 245, 251
Mechanical Turk (MTurk), 60–61, 62, 63
media synthesis machines, 7, 19, 31, 40, 42, 43, 56, 59, 65, 73, 104, 110, 111, 120, 135, 136, 168, 182, 184. *See also* text synthesis machines
medicine
 AI hype, 14, 84–85, 165, 191
 automation in, 85–87
 bias in, 92
 chatbots, 89–90
 health care prediction systems, 86–88
 health insurance, 87–88
 patient care, 74–75

sepsis detection, 86
 Merchant, Brian, 64, 209, 212, 257
 Messeri, Lisa, 118–119
 Meta, 16, 27, 40, 95, 111, 113–114, 126, 134, 155, 173, 180, 230, 257. *See also* Facebook
 Microsoft, 8, 9, 13, 21, 34, 40, 51, 95, 98, 130, 149, 158, 179, 160, 186, 194, 236
 Microsoft Office, 173
 Midjourney, 31, 104, 108, 111, 121, 171
 mining industry, 46, 157
 Minsky, Marvin, 11–12, 13, 22, 191
Mission Impossible: Dead Reckoning (film), 10
 Mitchell, Margaret “Meg”, x, 161, 179
 Model Alliance, 57
 model cards, 179–180
 Moore, Nicole, 56, 212
 Morehouse College, 166
 Morris, Rob, 67–68
 Mostaque, Emad, 103–104, 140
 Motaung, Daniel, 65
 Motherboard, 67, 134, 210, 212
 MSN, 130
 Mumm, Jared, 15–16
 Murati, Mira, 167, 194
 Murray, Charles, 34–35
 mushroom hunting, fake books on, 106–107
 Musk, Elon, 16, 36, 38–39, 140, 141, 150–151, 191, 208, 251
Mystery AI Hype Theater 3000 (podcast), x, 219, 229, 235, 236, 243, 245, 247
Mystery Science Theater 3000 (MST3K) (TV show), x

Nadella, Satya, 186
 Nagel, Ernest, 46
 Narayan, Shankar, 99
 natalism, 37–39, 147
 National Eating Disorders Association (NEDA), 56–57, 91
 National Health Service (NHS), 74–75
 National Nurses United, 86, 170, 187
 natural language inference, 221
Nature (periodical), 88, 229, 232
 Nelson, Alondra, 238
 neoliberalism, 69, 239
 Netanyahu, Benjamin, 3
 Netflix, 6, 41, 104, 187
 Neural Information Processing systems (NeurIPS), 8, 151
 neural language models, 25, 26–27
 neural networks, 11, 12, 21, 25–26, 47, 61, 155
 Newsom, Gavin, 73, 76
 New York City government, 72, 249
 New York City Police Department, 72
 New York Mycological Society, 106–107
New York Times (periodical), 39, 53, 58–59, 111–112, 132, 199, 208, 211, 213, 217, 220, 221, 225, 228, 232, 234, 236, 237, 239, 255

n-gram language models, 23–25, 26
nH Predict, 87
Nightshade, 65
Nobel Turing Challenge, 117, 118
Noble, Safiya Umoja, 211
Nvidia, 13, 194, 257

Obermayer, Ziad, 86
O’Keefe, Alex, 187
One Medical, 89
“On the Dangers of Stochastic Parrots” (Bender, Gebru), 31–32, 201
OpenAI, 1, 8, 9, 13, 21, 28, 33, 39, 47, 51–52, 59, 62, 65, 75, 77, 92, 93, 111, 138, 142–143, 149, 155, 160, 170, 194. *See also* ChatGPT
Ortiz, Karla, 104, 109, 111–112, 227
Oz, Mehmud, 172

Page, Larry, 149
Palantir, 75
panopticon, 152, 241
paper clip maximizer, 140, 145
paper mills, 122
Papers with Code, 113
palimpsest, 253
Paris Agreement, 157
peer review, 116–117, 118, 121, 122, 123, 124
 lack thereof, 21, 125, 141, 232, 239
perceptrons, 25–26, 27
Pew Research Center, 93
Pichai, Sundar, 179, 193
Picoult, Jodi, 111
Pitchbook, 40
Pitts, Walter, 25
plagiarism, 53, 93–94
policymakers, 78, 152, 153, 156, 161–162, 163, 176, 178. *See also* regulators
Pope, Denise, 93
population trends, 37, 38, 147
pornography
 deepfakes, 3
 filtering of by crowdworkers, 59
potemkin citation, 239
precautionary principle, 242
Prescod-Weinstein, Chanda, 233
pretrial risk assessment, 70–71, 216
privacy, 20, 53, 75, 91, 185–186, 188, 254
private equity, 91, 128, 131, 132–134, 135, 158
probability of doom (p(doom)), 1–2, 4, 19–20, 137–138, 141, 161–162
probability of hope (p(hope)), 1–2
product reviews, 125–126, 130
programming languages, 11, 50, 73
Prolific, 62

prompts, 7, 16, 48, 53, 93, 97, 102, 103, 106, 107, 109, 110, 112, 114, 127, 157, 159, 219, 249
Proposition 22 (2020), 55, 255
ProPublica, 133, 216, 236
Protecting the Right to Organize (PRO) Act, 187–188
psychotherapy, 12–13, 67–68, 90–91, 202, 206
publishing, 105–107, 121
Pulitzer Center, 170
Pulitzer Prize, 130, 133
Pygmalion (Shaw), 12
Pyx Health, 91

Quake III Arena (videogame), 53
Qualtrics, 62
queer (sexual identity), 33
 Turing, Alan, 241
 See also transgender people
queer theory, 114
questioning AI hype, 164–170
Quintarelli, Stefano, 201

race science, 23, 35, 207
racism, ix, 5, 18
 child welfare system, 69–70
 color-blind, 146
 facial recognition tools, 2–3, 192
 general intelligence and, 18, 33–36, 153
 in medicine, 92, 224
 pretrial risk assessment, 70–71
 in scientific texts, 114
Rafalow, Matt, 94–95
Raji, Deb, ix, 220
rationalism, 150
“Recall” feature, 186
recommender systems, 6
recourse, 63, 168–169, 184, 185
red-teaming, 62–63
regulation, 20, 91, 161, 162, 177
 accountability, 182–184
 data minimization, 185–186
 data rights, 184–185, 254
 disclosure, 180–182
 enforcing, 164, 174–176
 innovation and, 178, 188, 251
 labor protections, 186–188
 policymaking, 176–188
 recourse, 184
 transparency, 178–180
 See also legislation
regulators, 17, 162, 175, 176, 179, 180. *See also* policymakers
reinforcement learning from human feedback (RLHF), 28

- Remotasks, 62, 63
- Respond Crisis Translation, 74
- retrieval augmented generation, 249
- Ribes, David, 232
- Rideshare Drivers United, 56
- robotaxis, 55–56, 58–59
- robots
 - in fiction, 10, 40, 140
 - industrial, 46–47
 - police, 72
 - rogue, 10, 21, 40, 140
 - scientist, 103, 113
- robots.txt, 185
- Rogan, Joe, 38
- Rogers, Anna, 221
- Ronacher, Armin, 53
- Rosário, Ramiro, 79, 219
- Rosenblatt, Frank, 12, 25
- Russell, Stuart, 1, 140
- Rutkowski, Greg, 104, 112

- Sacks, David, 208
- “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence” (2023), 161–162, 181
- Safe Street Rebel, 56
- Safe Superintelligence, Inc., 143
- SAG-AFTRA, 41–43, 186–187
- SALAMI, 5, 201
- Sama (Samasource), 62, 65
- Samudzi, Zoé, 192
- Scale AI, 63
- School of Media and Journalism, 127
- Schumer, Charles “Chuck”, 1–2, 137, 162, 177, 237, 246
- Schwartz, Steven A., 16, 76
- Science* (periodical), 121, 222, 233, 238, 248, 251
- science and engineering, 113–125
 - AI hype, 19, 113, 123
 - AI Scientist, 117–118, 120
 - creativity and, 101–103, 135–136
 - damages to scientific ecosystem, 120–123
 - dehumanization of, 120
 - humanity of, 117–120
 - human subjects, 115
 - language of used in LLM training, 113–114
 - literature review, 114–115
 - peer review, 116–117, 121, 124
 - problem-solving needs, 123–125
 - risks of LLMs, 124
 - synthetic scientific papers, 121–122
 - technological advances in, 113

Scientific American (periodical), 46, 224, 244
search engine optimization (SEO), 51, 108, 126, 129, 130, 133
search engines, 8, 51, 76, 80, 171–172, 249
Seetharaman, Deepa, 62
self-driving cars, 39
 AI harms, 16, 177–178
 AI hype, 15, 191
 robotaxis, 55–56, 58–59
self-replication, 161
sentience, 9, 18, 21–23, 31, 39–40, 139, 142, 149
sepsis detection, 86
Sequoia Capital, 193
sexism, ix
sexual harassment, 3
Shah, Munjal, 90
Shanghai Jiao Tong University, 165
Shannon, Claude, 23
Shaw, George Bernard, 12
Shelley, Mary, 10
Shkreli, Martin, 148
ShotSpotter, 167–168, 248
Silberman, Michael (Six), 65
silicon samples, 115
Silicon Valley, ix, 2, 4, 37, 56, 73, 98, 128, 208
Simon, Herbert, 11–12, 241
singularitarianism, 150, 151
Smith, Brad, 160, 243
Smith, Linda Tuhiwai, 215–216
social services
 AI harms, 19
 AI hype, 14, 164
 child welfare system, 69–70
 humans necessary for, 98–99
 psychotherapy, 67–68
Society of Authors, 105
Sora, 103, 187
Soros Fund Management, 134
SoundThinking, 167–168
“Sparks of Artificial General Intelligence” (Bubeck), 21–22, 34
speech to text, 6
Sports Illustrated (periodical), 125
Spotify, 196
Stability AI, 103, 105, 111, 140
Stable Diffusion, 101–102, 103, 104, 108, 109
Stanford-Binet IQ test, 35, 37
Stanford Human-Centered Artificial Intelligence Lab, 60
Stanford University, 35, 37, 60, 61, 116, 183, 206, 221, 224, 225
Stark, Luke, 247
Starr, David Jordan, 37
Star Trek: The Next Generation (TV show), 25

Stat News (periodical), 87, 223
 stochastic parrots, 5, 31–32, 89, 201
 Streamlining Effective Access and Retrieval of Content to Help (SEARCH) Act (2023), 78
 stutter, 220
 suicide, 67, 91
 Suleyman, Mustafa, 13
 Summit Learning, 95
 Sunak, Rishi, 74–75, 76
 superhuman, use of term, 166–167
 superintelligence, 36, 39, 40, 142–143, 160
 surveillance, 144, 152, 164

- biometric, 191–192
- in the classroom, 94–95
- law enforcement, 2, 6, 167–168
- panopticon, 152
- in the workplace, 47, 175–176, 250

 Sutskever, Ilya, 8, 21, 61, 90–91, 143, 155
 Swaak, Taylor, 96, 97
 synthetic image machines, 31. *See also* synthetic media machines
 synthetic text extruding machines, 31, 72, 73, 75, 76, 78, 80, 106, 181. *See also* large language models (LLMs)
 systems safety engineering, 142

T9 system, 24
 Tallinn, Jaan, 151
 Tan, Garry, 148
 Tapper, Jack, 153, 245
 Tarnoff, Ben, 32, 202, 206, 224
 taxi industry, 42, 55–56
 Taylor, Christie, x
 Teach For America, 95
 Tech and Liberty Project, 99
 technology

- in the classroom, 94–95
- clear description of, 201
- consent and, 251, 255
- evaluation methodology, 80–84, 167–168
- General Purpose (GPT), 47
- hype, 9
- intelligent, 7
- labor replaced with, 41–42, 47
- marketed and sold as AI, 5–7
- socially situated, 188–190

 techno-optimism, 39, 66
 “The Techno-Optimist Manifesto” (Andreessen), 38, 146–147
 Te Hiku Media, 189
 Terman, Lewis, 35, 37
 TESCREAL, 150–151
 Tesla, 16, 38, 39, 178, 191, 251
 Texas A&M University, 15–16, 94

text synthesis machines, 7, 49, 68, 73, 99, 103, 106–107, 129, 171. *See also* synthetic text extruding machines
 texting, 24
 Theisen, Lauren, 134
 Thiel, Peter, 75, 128
 thinking machines, 11
 Thorp, Jer, 107–108
Time (periodical), 59, 62, 149, 210, 213, 215, 240
Timis, David Alexandru, 251
tokens, 205
 Tokuda-Hall, Maggie, 174
 Toledo, Adam, 248
 Torres, Émile, 150, 151
 training
 carbon footprint, 157–159
 corpora, 23–24, 27
 data rights and, 184–185
 dataset documentation, 178–180
 image generation, 107–109
 language models, 23–28
 perceptrons, 25–26
 transcription/translation
 automated, 6, 7, 24–25, 26
 benchmarking, 81–82
 harms of, 15
 machine, 6, 15, 73–74, 81–82, 189
 transgender people, 33, 92
 transhumanism, 150, 151
 translation models, 24
 trigram language models, 24
 Tronc/Tribune Publishing, 128
 Trump, Donald, 76
 Tucker, Emily, 201
 Turing, Alan, 153–154, 241
 Turing Test, 117, 154
 Turkoicon, 65, 214
 Turnitin, 94
 Twitter. *See* X/Twitter
2001: A Space Odyssey (film), 10, 140

Uber, 55, 56, 129, 187
 unigram language models, 23, 24
 UnitedHealth Group, 87, 92
 United Kingdom, 15–16, 43, 69, 74–75, 145, 146
 United Nations, 3, 144, 157
 Universal Declaration of Human Rights, 144
 universities and colleges, 93–94, 95–98, 119, 124, 234
 university faculty, casualization of, 124, 226, 234
 University of California, Berkeley, 142
 University of Cambridge, 142, 151

University of Chicago, 65
University of North Carolina, Chapel Hill, 127
University of Pennsylvania, 97
U.S. Customs and Border Protection, 73, 192
U.S. Department of Defense, 81
U.S. Equal Employment Opportunity Commission, 180
U.S. Food and Drug Administration, 91
U.S. House of Representatives, 78
U.S. National Highway Traffic Safety Administration, 251
U.S. National Institute of Standards and Technology (NIST), 162
U.S. National Institutes of Health, 95
U.S. National Labor Relations Board, 175
U.S. National Science Foundation, 95
U.S. Occupational Safety and Health Administration (OSHA), 46–47
U.S. Senate, 1–2, 137–138, 162, 177
U.S. Supreme Court, 77, 133
U.S. Transportation Security Administration, 192

Veale, Michael, 159
vendor lock-in, 51–52
venture capital (VC), 1, 2, 8, 13, 17, 36, 38, 40, 42, 59, 68, 77, 89–90, 91, 92, 112, 120, 128, 134, 141, 148, 158, 188, 189, 193, 195, 208
Vice Media, 134, 208, 215, 250
violence, ix, 2, 15, 59, 63, 70, 133, 144, 145, 147, 168, 186, 191, 248
visual artists, 19, 64–65, 66, 103–105, 110, 194
Vogt, Kyle, 59

Wall Street Journal (periodical), 34, 88, 132, 213, 233
warfare, 3–4, 144
Washington Post (periodical), 21, 79, 129, 132, 134, 159, 200, 203, 207, 215, 216, 240, 243, 251, 255, 257
watermarks, 122, 181–182, 252–253
Waymo, 55, 56
weavers, 44–45
Weizenbaum, Joseph, 12–13, 17, 28–29, 32–33, 67, 90, 182–183, 202, 206
Wharton Business School, 97
white supremacy, ix, 146
Whittaker, Meredith, 239, 252
Wiener, Norbert, 11
Wiener, Scott, 141
Wiley, 122
Williams, Adrienne, 95
Williams, Robert, 2, 3, 200
Woebot, 91, 224
Woodruff, Porcha, 2–3
word error rate, 81–82
workplace
 AI hype, 66
 automation of jobs, 42–43, 47–49, 193
 crowdworking, 59–62

labor exploitation, 43, 49, 61, 62–64, 169
labor replaced with technology, 28, 41–42, 47
red-teaming, 62–63
remote human workers, 43, 49, 58–64
resisting AI hype, 170–171
surveillance, 47, 175–176, 250
worker protections, 18–20, 63, 64–65, 110, 175–176, 186–188
World Economic Forum, 118
World War I, 35
World War II, 144
Writers Guild of America (WGA), 41, 54, 64, 170, 186–187
writing, 41, 49–50, 54–55, 64, 187
Wysa, 91

xAI, 39
X/Twitter, 38, 39, 68, 77, 148, 245, 253

Yahoo!, 51
Yang, Andrew, 140
Y Combinator, 59, 89, 148
Yerkes, Robert, 35–36
Yuan, Eric, 52
Yudkowsky, Eliezer, 149

zero-trust systems, 176, 250–251
Ziff, Sara, 57
Zoom, 52
Zuckerberg, Mark, 95, 133, 155

About the Authors

Dr. Emily M. Bender is a professor of linguistics at the University of Washington, where she is also the faculty director of the Computational Linguistics Master of Science program and affiliate faculty in the School of Computer Science and Engineering and the Information School. In 2023, she was included in the inaugural *Time* 100 list of the most influential people in AI. She is frequently consulted by policymakers, from municipal officials to the federal government to the United Nations, for insight into how to understand so-called AI technologies.

Dr. Alex Hanna is the director of research at the Distributed AI Research Institute (DAIR) and a lecturer in the School of Information at the University of California, Berkeley. She is an outspoken critic of the tech industry, a proponent of community-based uses of technology, and a highly sought-after speaker and expert whose articles have been featured across the media, including in the *Washington Post*, the *Financial Times*, the *Atlantic*, and *Time*.

Discover great authors, exclusive offers, and more at hc.com.

The logo for Bookperk is centered on a black rectangular background. The word "Bookperk" is written in a sans-serif font, with "Book" in yellow and "perk" in white. A faint, thin white circle is visible behind the text.

Bookperk

[Sign up for Bookperk](#) and get e-book bargains, sneak peeks, special offers, and more—delivered straight to your inbox.

SIGN UP NOW

Copyright

THE AI CON. Copyright © 2025 by Emily M. Bender and Alex Hanna. All rights reserved under International and Pan-American Copyright Conventions. By payment of the required fees, you have been granted the nonexclusive, nontransferable right to access and read the text of this e-book on-screen. No part of this text may be reproduced, transmitted, downloaded, decompiled, reverse-engineered, or stored in or introduced into any information storage and retrieval system, in any form or by any means, whether electronic or mechanical, now known or hereafter invented, without the express written permission of HarperCollins e-books. Without limiting the author's and publisher's exclusive rights, any unauthorized use of this publication to train generative artificial intelligence (AI) technologies is expressly prohibited.

FIRST EDITION

Cover design by Kris Potter

Library of Congress Cataloging-in-Publication Data has been applied for.

Digital Edition MAY 2025 ISBN: 978-0-06-341855-4

Print ISBN: 978-0-06-341856-1

About the Publisher

Australia

HarperCollins Publishers Australia Pty. Ltd.
Level 13, 201 Elizabeth Street
Sydney, NSW 2000, Australia
www.harpercollins.com.au

Canada

HarperCollins Publishers Ltd.
Harlequin Enterprises ULC
www.harlequin.com
Bay Adelaide Centre, East Tower
22 Adelaide Street West, 41st Floor
Toronto, Ontario, M5H 4E3
www.harpercollins.ca

India

HarperCollins India
A 75, Sector 57
Noida
Uttar Pradesh 201 301
www.harpercollins.co.in

Ireland

HarperCollins Publishers
Macken House,
39/40 Mayor Street Upper,
Dublin 1, D01 C9W8, Ireland
www.harpercollins.co.uk

New Zealand

HarperCollins Publishers New Zealand
Unit D1, 63 Apollo Drive
Rosedale 0632
Auckland, New Zealand
www.harpercollins.co.nz

United Kingdom

HarperCollins Publishers Ltd.
1 London Bridge Street
London SE1 9GF, UK
www.harpercollins.co.uk

United States

HarperCollins Publishers Inc.
195 Broadway
New York, NY 10007
www.harpercollins.com



Your gateway to knowledge and culture. Accessible for everyone.



z-library.sk

z-lib.gs

z-lib.fm

go-to-library.sk



[Official Telegram channel](#)



[Z-Access](#)



<https://wikipedia.org/wiki/Z-Library>