KATE CRAWFORD

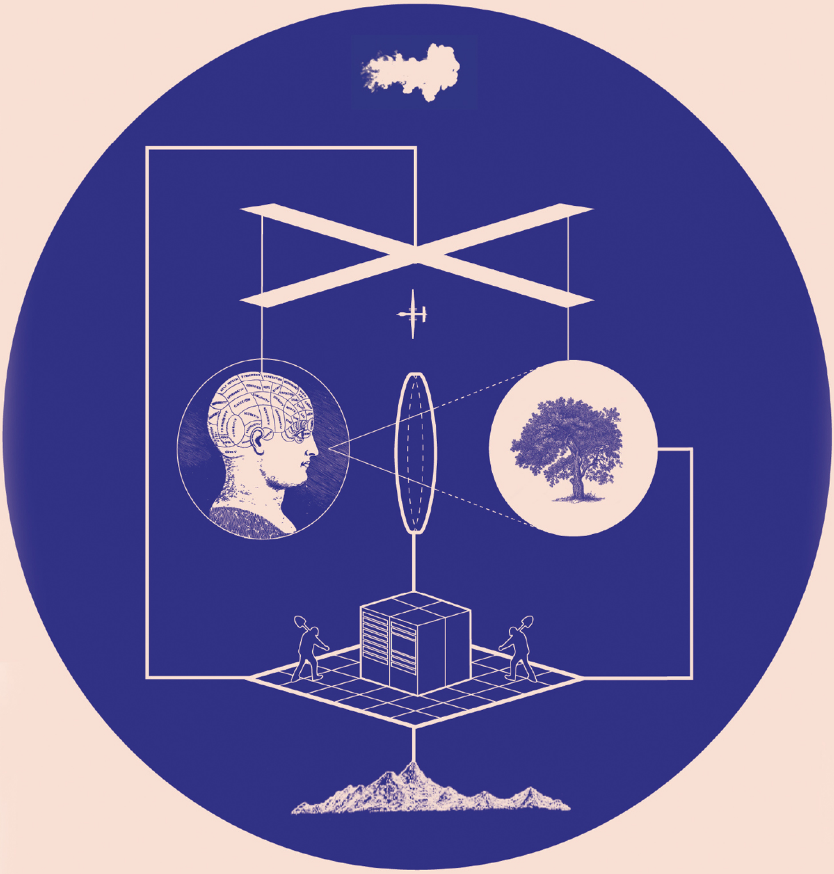

ATLAS OF AI

Atlas of AI

Power, Politics, and the Planetary Costs of Artificial Intelligence

KATE CRAWFORD

Yale UNIVERSITY PRESS

New Haven and London

Yale University Press books may be purchased in quantity for educational, business, or promotional use. For information, please e-mail sales.press@yale.edu (U.S. office) or sales@yaleup.co.uk (U.K. office).

Cover design and chapter opening illustrations by Vladan Joler.
Set in Minion by Tseng Information Systems, Inc.
Printed in the United States of America.

Library of Congress Control Number: 2020947842
ISBN 978-0-300-20957-0 (hardcover : alk. paper)

A catalogue record for this book is available from the British Library.

This paper meets the requirements of ANSI/NISO Z39.48-1992 (Permanence of Paper).

10 9 8 7 6 5 4 3 2 1

For Elliott and Margaret

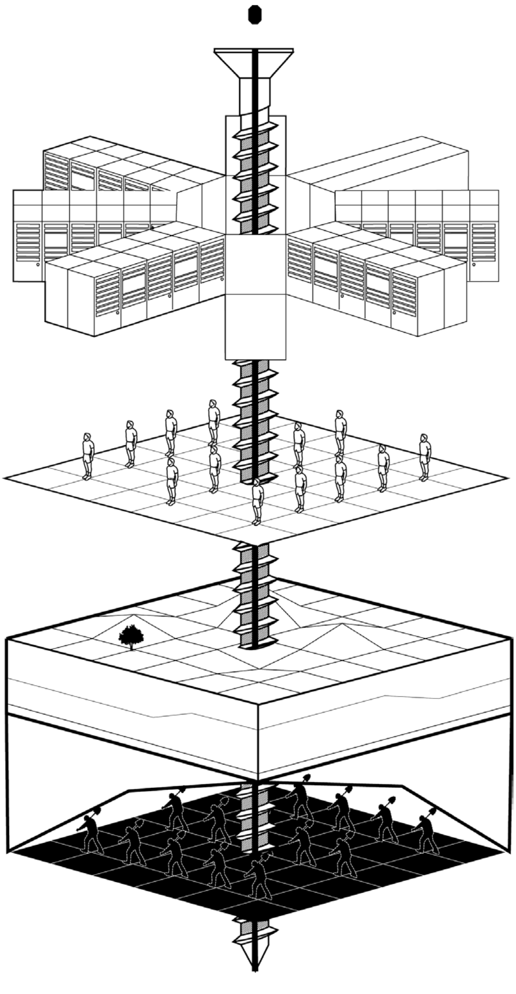

1
Earth

The Boeing 757 banks right over San Jose on its final approach to San Francisco International Airport. The left wing drops as the plane lines up with the runway, revealing an aerial view of the tech sector's most iconic location. Below are the great empires of Silicon Valley. The gigantic black circle of Apple's headquarters is laid out like an uncapped camera lens, glistening in the sun. Then there's Google's head office, nestled close to NASA's Moffett Federal Airfield. This was once a key site for the U.S. Navy during World War II and the Korean War, but now Google has a sixty-year lease on it, and senior executives park their private planes here. Arrayed near Google are the large manufacturing sheds of Lockheed Martin, where the aerospace and weapons manufacturing company builds hundreds of orbital satellites destined to look down on the activities of Earth. Next, by the Dumbarton Bridge, appears a collection of squat buildings that are home to Facebook, ringed with massive parking lots close to the sulfuric salt ponds of the Ravenswood Slough. From this vantage point, the nondescript suburban cul-de-sacs and industrial midrise skyline of Palo Alto betray little of its true wealth, power, and influence. There are only a few hints of its centrality in the global economy and in the computational infrastructure of the planet.

I'm here to learn about artificial intelligence and what it is made from. To see that, I will need to leave Silicon Valley altogether.

From the airport, I jump into a van and drive east. I cross the San Mateo–Hayward Bridge and pass by the Lawrence Livermore National Laboratory, where Edward Teller directed his research into thermonuclear weapons in the years after World War II. Soon the Sierra Nevada foothills rise beyond the Central Valley towns of Stockton and Manteca. Here the roads start winding up through the tall granite cliffs of the Sonora Pass and down the eastern side of the mountains toward grassy valleys dotted with golden poppies. Pine forests give way to the alkaline waters of Mono Lake and the parched desert landforms of the Basin and Range. To refuel, I pull into Hawthorne, Nevada, site of the world's biggest ammunition depot, where the U.S. Army stores armaments in dozens of dirt-covered ziggurats that populate the valley in neat rows. Driving along Nevada State Route 265 I see a lone VORTAC in the distance, a large bowling pin–shaped radio tower that was designed for the era before GPS. It has a single function: it broadcasts "I am here" to all passing aircraft, a fixed point of reference in a lonely terrain.

My destination is the unincorporated community of Silver Peak in Nevada's Clayton Valley, where about 125 people live, depending on how you count. The mining town, one of the oldest in Nevada, was almost abandoned in 1917 after the ground was stripped bare of silver and gold. A few gold rush buildings still stand, eroding under the desert sun. The town may be small, with more junked cars than people, but it harbors something exceedingly rare. Silver Peak is perched on the edge of a massive underground lake of lithium. The valuable

Silver Peak Lithium Mine. Photograph by Kate Crawford

lithium brine under the surface is pumped out of the ground and left in open, iridescent green ponds to evaporate. From miles away, the ponds can be seen when they catch the light and shimmer. Up close, it's a different view. Alien-looking black pipes erupt from the ground and snake along the salt-encrusted earth, moving in and out of shallow trenches, ferrying the salty cocktail to its drying pans.

Here, in a remote pocket of Nevada, is a place where the stuff of AI is made.

Mining for AI

Clayton Valley is connected to Silicon Valley in much the way that the nineteenth-century goldfields were to early San Francisco. The history of mining, like the devastation it leaves in its wake, is commonly overlooked in the strategic amnesia that accompanies stories of technological progress. As historical geographer Gray Brechin points out, San Francisco was built from the gains of pulling gold and silver out of the lands of California and Nevada in the 1800s.[1] The city is made from mining. Those same lands had been taken from Mexico under the Treaty of Guadalupe Hidalgo in 1848 at the end of the Mexican-American War, when it was already clear to the settlers that these would be highly valuable goldfields. It was a textbook example, Brechin observes, of the old adage that "commerce follows the flag, but the flag follows the pick."[2] Thousands of people were forced from their homes during this substantial territorial expansion of the United States. After America's imperial invasion, the miners moved in. The land was stripped until the waterways were contaminated and the surrounding forests destroyed.

Since antiquity, the business of mining has only been profitable because it does not have to account for its true costs: including environmental damage, the illness and death of miners, and the loss to the communities it displaces. In 1555, Georgius Agricola, known as the father of mineralogy, observed that "it is clear to all that there is greater detriment from mining than the value of the metals which the mining produces."[3] In other words, those who profit from mining do so only because the costs must be sustained by others, those living and those not yet born. It is easy to put a price on precious metals, but what is the exact value of a wilderness, a clean stream, breathable air, the health of local communities? It was never estimated, and thus an easy calculus emerged: extract everything as rapidly as possible. It was the "move fast and break things" of a different time. The result was that the Central Valley was decimated, and as one tourist observed in 1869, "Tornado, flood, earthquake and volcano combined could hardly make greater havoc, spread wider ruin and wreck than [the] gold-washing operations. . . . There are no rights which mining respects in California. It is the one supreme interest."[4]

As San Francisco drew enormous wealth from the mines, it was easy for its populace to forget where it all came from. The mines were located far from the city they enriched, and this remoteness allowed city dwellers to remain ignorant of what was happening to the mountains, rivers, and laborers that fed their fortunes. But small reminders of the mines are all around. The city's new buildings used the same technology that came from deep within the Central Valley for transport and life support. The pulley systems that carried miners down into the mine shafts were adapted and turned upside down to transport people in elevators to the top of the city's high-rises.[5] Brechin suggests that we should think of the skyscrapers of San Francisco as inverted minescapes. The ores extracted from holes in the ground were sold to create the stories in the air; the deeper the extractions went, the higher the great towers of office work stretched into the sky.

San Francisco is enriched once more. Once it was gold ore that underwrote fortunes; now it is the extraction of substances like white lithium crystal. It's known in mineral markets as "gray gold."[6] The technology industry has become a new supreme interest, and the five biggest companies in the world by market capitalization have offices in this city: Apple, Microsoft, Amazon, Facebook, and Google. Walking past the start-up warehouses in the SoMa district where miners in tents once lived, you can see luxury cars, venture capital–backed coffee chains, and sumptuous buses with tinted windows running along private routes, carrying workers to their offices in Mountain View or Menlo Park.[7] But only a short walk away is Division Street, a multilane thoroughfare between SoMa and the Mission district, where rows of tents have returned to shelter people who have nowhere to go. In the wake of the tech boom, San Francisco now has one of the highest rates of street homelessness in the United States.[8] The United Nations special rapporteur on adequate housing called it an "unacceptable" human rights violation, due to the thousands of homeless residents denied basic necessities of water, sanitation, and health services in contrast to the record number of billionaires who live nearby.[9] The greatest benefits of extraction have been captured by the few.

In this chapter we'll traverse across Nevada, San Jose, and San Francisco, as well as Indonesia, Malaysia, China, and Mongolia: from deserts to oceans. We'll also walk the spans of historical time, from conflict in the Congo and artificial black lakes in the present day to the Victorian passion for white latex. The scale will shift, telescoping from rocks to cities, trees to megacorporations, transoceanic shipping lanes to the atomic bomb. But across this planetary supersystem we will see the logics of extraction, a constant drawdown of minerals, water, and fossil fuels, undergirded by the violence of wars, pollution, extinction, and depletion. The effects of large-scale computation can be found in the atmosphere, the oceans, the earth's crust, the deep time of the planet, and the brutal impacts on disadvantaged populations around the world. To understand it all, we need a panoramic view of the planetary scale of computational extraction.

Landscapes of Computation

I'm driving through the desert valley on a summer afternoon to see the workings of this latest mining boom. I ask my phone to direct me to the perimeter of the lithium ponds, and it replies from its awkward perch on the dashboard, tethered by a white USB cable. Silver Peak's large, dry lake bed was formed millions of years ago during the late Tertiary Period. It's surrounded by crusted stratifications pushing up into ridgelines containing dark limestones, green quartzites, and gray and red slate.[10] Lithium was discovered here after the area was scoped for strategic minerals like potash during World War II. This soft, silvery metal was mined in only modest quantities for the next fifty years, until it became highly valuable material for the technology sector.

In 2014, Rockwood Holdings, Inc., a lithium mining operation, was acquired by the chemical manufacturing company Albemarle Corporation for $6.2 billion. It is the only operating lithium mine in the United States. This makes Silver Peak a site of intense interest to Elon Musk and the many other tech tycoons for one reason: rechargeable batteries. Lithium is a crucial element for their production. Smartphone batteries, for example, usually contain about three-tenths of an ounce of it. Each Tesla Model S electric car needs about one hundred thirty-eight pounds of lithium for its battery pack.[11] These kinds of batteries were never intended to supply a machine as power hungry as a car, but lithium batteries are currently the only mass-market option available.[12] All of these batteries have a limited lifespan; once degraded, they are discarded as waste.

About two hundred miles north of Silver Peak is the Tesla Gigafactory. This is the world's largest lithium battery plant. Tesla is the number-one lithium-ion battery consumer in the world, purchasing them in high volumes from Panasonic and Samsung and repackaging them in its cars and home chargers. Tesla is estimated to use more than twenty-eight thousand tons of lithium hydroxide annually—half of the planet's total consumption.[13] In fact, Tesla could more accurately be described as a battery business than a car company.[14] The imminent shortage of such critical minerals as nickel, copper, and lithium poses a risk for the company, making the lithium lake at Silver Peak highly desirable.[15] Securing control of the mine would mean controlling the U.S. domestic supply.

As many have shown, the electric car is far from a perfect solution to carbon dioxide emissions.[16] The mining, smelting, export, assemblage, and transport of the battery supply chain has a significant negative impact on the environment and, in turn, on the communities affected by its degradation. A small number of home solar systems produce their own energy. But for the majority of cases, charging an electric car necessitates taking power from the grid, where currently less than a fifth of all electricity in the United States comes from renewable energy sources.[17] So far none of this has dampened the determination of auto manufacturers to compete with Tesla, putting increasing pressure on the battery market and accelerating the removal of diminishing stores of the necessary minerals.

Global computation and commerce rely on batteries. The term "artificial intelligence" may invoke ideas of algorithms, data, and cloud architectures, but none of that can function without the minerals and resources that build computing's core components. Rechargeable lithium-ion batteries are essential for mobile devices and laptops, in-home digital assistants, and data center backup power. They undergird the internet and every commerce platform that runs on it, from banking to retail to stock market trades. Many aspects of modern life have been moved to "the cloud" with little consideration of these material costs. Our work and personal lives, our medical histories, our leisure time, our entertainment, our political interests—all of this takes place in the world of networked computing architectures that we tap into from devices we hold in one hand, with lithium at their core.

The mining that makes AI is both literal and metaphorical. The new extractivism of data mining also encompasses and propels the old extractivism of traditional mining. The stack required to power artificial intelligence systems goes well beyond the multilayered technical stack of data modeling, hardware, servers, and networks. The full-stack supply chain of AI reaches into capital, labor, and Earth's resources—and from each, it demands an enormous amount.[18] The cloud is the backbone of the artificial intelligence industry, and it's made of rocks and lithium brine and crude oil.

In his book *A Geology of Media,* theorist Jussi Parikka suggests we think of media not from Marshall McLuhan's point of view—in which media are extensions of the human senses—but rather as extensions of Earth.[19] Computational media now participate in geological (and climatological) processes, from the transformation of the earth's materials into infrastructures and devices to the powering of these new systems with oil and gas reserves. Reflecting on media and technology as geological processes enables us to consider the radical depletion of nonrenewable resources required to drive the technologies of the present moment. Each object in the extended network of an AI system, from network routers to batteries to data centers, is built using elements that required billions of years to form inside the earth.

From the perspective of deep time, we are extracting Earth's geological history to serve a split second of contemporary technological time, building devices like the Amazon Echo and the iPhone that are often designed to last for only a few years. The Consumer Technology Association notes that the average smartphone life span is a mere 4.7 years.[20] This obsolescence cycle fuels the purchase of more devices, drives up profits, and increases incentives for the use of unsustainable extraction practices. After a slow process of development, these minerals, elements, and materials then go through an extraordinarily rapid period of excavation, processing, mixing, smelting, and logistical transport—crossing thousands of miles in their transformation. What begins as ore removed from the ground, after the spoil and the tailings are discarded, is then made into devices that are used and discarded. They ultimately end up buried in e-waste dumping grounds in places like Ghana and Pakistan. The lifecycle of an AI system from birth to death has many fractal supply chains: forms of exploitation of human labor and natural resources and massive concentrations of corporate and geopolitical power. And all along the chain, a continual, large-scale consumption of energy keeps the cycle going.

The extractivism on which San Francisco was built is echoed in the practices of the tech sector based there today.[21] The massive ecosystem of AI relies on many kinds of extraction: from harvesting the data made from our daily activities and expressions, to depleting natural resources, and to exploiting labor around the globe so that this vast planetary network can be built and maintained. And AI extracts far more from us and the planet than is widely known. The Bay Area is a central node in the mythos of AI, but we'll need to traverse far beyond the United States to see the many-layered legacies of human and environmental damage that have powered the tech industry.

The Mineralogical Layer

The lithium mines in Nevada are just one of the places where the materials are extracted from the earth's crust to make AI. There are many such sites, including the Salar in southwest Bolivia—the richest site of lithium in the world and thus a site of ongoing political tension—as well as places in central Congo, Mongolia, Indonesia, and the Western Australia deserts. These are the other birthplaces of AI in the greater geography of industrial extraction. Without the minerals from these locations, contemporary computation simply does not work. But these materials are in increasingly short supply.

In 2020, scientists at the U.S. Geological Survey published a short list of twenty-three minerals that are a high "supply risk" to manufacturers, meaning that if they became unavailable, entire industries—including the tech sector—would grind to a halt.[22] The critical minerals include the rare earth elements dysprosium and neodymium, which are used inside iPhone speakers and electric vehicle motors; germanium, which is used in infrared military devices for soldiers and in drones; and cobalt, which improves performance for lithium-ion batteries.

There are seventeen rare earth elements: lanthanum, cerium, praseodymium, neodymium, promethium, samarium, europium, gadolinium, terbium, dysprosium, holmium, erbium, thulium, ytterbium, lutetium, scandium, and yttrium. They are processed and embedded in laptops and smartphones, making those devices smaller and lighter. The elements can be found in color displays, speakers, camera lenses, rechargeable batteries, hard drives, and many other components. They are key elements in communication systems, from fiber-optic cables and signal amplification in mobile communication towers to satellites and GPS technology. But extracting these minerals from the ground often comes with local and geopolitical violence. Mining is and always has been a brutal undertaking. As Lewis Mumford writes, "Mining was the key industry that furnished the sinews of war and increased the metallic contents of the original capital hoard, the war chest: on the other hand, it furthered the industrialization of arms, and enriched the financier by both processes."[23] To understand the business of AI, we must reckon with the war, famine, and death that mining brings with it.

Recent U.S. legislation that regulates some of those seventeen rare earth elements only hints at the devastation associated with their extraction. The 2010 Dodd-Frank Act focused on reforming the financial sector in the wake of the 2008 financial crisis. It included a specific provision about so-called *conflict minerals,* or natural resources extracted in a conflict zone and then sold to fund the conflict. Companies using gold, tin, tungsten, and tantalum from the region around the Democratic Republic of the Congo now had a reporting requirement to track where those minerals came from and whether the sale was funding armed militia in the region.[24] Like "conflict diamonds," the term "conflict minerals" masks the profound suffering and prolific killing in the mining sector. Mining profits have financed military operations in the decades-long Congo-area conflict, fueling the deaths of thousands and the displacement of millions.[25] Furthermore, working conditions inside the mines have often amounted to modern slavery.[26]

It took Intel more than four years of sustained effort to develop basic insight into its own supply chain.[27] Intel's supply chain is complex, with more than sixteen thousand suppliers in over a hundred countries providing direct materials for the company's production processes, tools, and machines for their factories, as well as their logistics and packaging services.[28] In addition, Intel and Apple have been criticized for auditing only smelters—not the actual mines—to determine the conflict-free status of minerals. The tech giants were assessing smelting plants outside of Congo, and the audits were often performed by locals. So even the conflict-free certifications of the tech industry are now under question.[29]

Dutch-based technology company Philips has also claimed that it was working to make its supply chain "conflict-free." Like Intel, Philips has tens of thousands of suppliers, each of which provides component parts for the company's manufacturing processes.[30] Those suppliers are themselves linked downstream to thousands of component manufacturers acquiring treated materials from dozens of smelters. The smelters in turn buy their materials from an unknown number of traders who deal directly with both legal and illegal mining operations to source the various minerals that end up in computer components.[31]

According to the computer manufacturer Dell, the complexities of the metals and mineral supply chains pose almost insurmountable challenges to the production of conflict-free electronics components. The elements are laundered through such a vast number of entities along the chain that sourcing their provenance proves impossible—or so the end-product manufacturers claim, allowing them a measure of plausible deniability for any exploitative practices that drive their profits.[32]

Just like the mines that served San Francisco in the nineteenth century, extraction for the technology sector is done by keeping the real costs out of sight. Ignorance of the supply chain is baked into capitalism, from the way businesses protect themselves through third-party contractors and suppliers to the way goods are marketed and advertised to consumers. More than plausible deniability, it has become a well-practiced form of bad faith: the left hand cannot know what the right hand is doing, which requires increasingly lavish, baroque, and complex forms of distancing.

While mining to finance war is one of the most extreme cases of harmful extraction, most minerals are not sourced from direct war zones. This doesn't mean, however, that they are free from human suffering and environmental destruction. The focus on conflict minerals, though important, has also been used to avert focus from the harms of mining writ large.

If we visit the primary sites of mineral extraction for computational systems, we find the repressed stories of acid-bleached rivers and deracinated landscapes and the extinction of plant and animal species that were once vital to the local ecology.

Black Lakes and White Latex

In Baotou, the largest city in Inner Mongolia, there is an artificial lake filled with toxic black mud. It reeks of sulfur and stretches as far as the eye can see, covering more than five and a half miles in diameter. The black lake contains more than 180 million tons of waste powder from ore processing.[33] It was created by the waste runoff from the nearby Bayan Obo mines, which is estimated to contain almost 70 percent of the world's reserves of rare earth minerals. It is the largest deposit of rare earth elements on the planet.[34]

China supplies 95 percent of the world's rare earth minerals. China's market domination, as the writer Tim Maughan observes, owes far less to geology than to the country's willingness to take on the environmental damage of extraction.[35] Although rare earth minerals like neodymium and cerium are relatively common, making them usable requires the hazardous process of dissolving them in large volumes of sulfuric and nitric acid. These acid baths yield reservoirs of poisonous waste that fill the dead lake in Baotou. This is just one of the places that are brimming with what environmental studies scholar Myra Hird calls "the waste we want to forget."[36]

To date, the unique electronic, optical, and magnetic uses of rare earth elements cannot be matched by any other metals, but the ratio of usable minerals to waste toxins is extreme. Natural resource strategist David Abraham describes the mining in Jiangxi, China, of dysprosium and terbium, which are used in a variety of high-tech devices. He writes, "Only 0.2

percent of the mined clay contains the valuable rare earth elements. This means that 99.8 percent of earth removed in rare earth mining is discarded as waste, called 'tailings,' that are dumped back into the hills and streams," creating new pollutants like ammonium.[37] In order to refine one ton of these rare earth elements, "the Chinese Society of Rare Earths estimates that the process produces 75,000 liters of acidic water and one ton of radioactive residue."[38]

About three thousand miles south of Baotou are the small Indonesian islands of Bangka and Belitung, off the coast of Sumatra. Bangka and Belitung produce 90 percent of Indonesia's tin, used in semiconductors. Indonesia is the world's second-largest producer of the metal, behind China. Indonesia's national tin corporation, PT Timah, supplies companies such as Samsung directly, as well as solder makers Chernan and Shenmao, which in turn supply Sony, LG, and Foxconn—all suppliers for Apple, Tesla, and Amazon.[39]

On these small islands, gray-market miners who are not officially employed sit on makeshift pontoons, using bamboo poles to scrape the seabed before diving underwater to suck tin from the surface by drawing their breath through giant, vacuumlike tubes. The miners sell the tin they find to middlemen, who also collect ore from miners working in authorized mines, and they mix it together to sell to companies like Timah.[40] Completely unregulated, the process unfolds beyond any formal worker or environmental protections. As investigative journalist Kate Hodal reports, "Tin mining is a lucrative but destructive trade that has scarred the island's landscape, bulldozed its farms and forests, killed off its fish stocks and coral reefs, and dented tourism to its pretty palm-lined beaches. The damage is best seen from the air, as pockets of lush forest huddle amid huge swaths of barren orange earth. Where not dominated by mines, this is pockmarked with graves, many holding the bodies of miners who have died over the centuries digging for tin."[41] The mines are everywhere: in backyards, in the forest, by the side of the road, on the beaches. It is a landscape of ruin.

It is a common practice of life to focus on the world immediately before us, the one we see and smell and touch every day. It grounds us where we are, with our communities and our known corners and concerns. But to see the full supply chains of AI requires looking for patterns in a global sweep, a sensitivity to the ways in which the histories and specific harms are different from place to place and yet are deeply interconnected by the multiple forces of extraction.

We can see these patterns across space, but we can also find them across time. Transatlantic telegraph cables are the essential infrastructure that ferries data between the continents, an emblem of global communication and capital. They are also a material product of colonialism, with its patterns of extraction, conflict, and environmental destruction. At the end of the nineteenth century, a particular Southeast Asian tree called *Palaquium gutta* became the center of a cable boom. These trees, found mainly in Malaysia, produce a milky white natural latex called gutta-percha. After English scientist Michael Faraday published a study in the *Philosophical Magazine* in 1848 about the use of this material as an electrical insulator, gutta-percha rapidly became the darling of the engineering world. Engineers saw gutta-percha as the solution to the problem of insulating telegraphic cables to withstand harsh and varying conditions on the ocean floor. The twisted strands of copper wire needed four layers of the soft, organic tree sap to protect them from water incursion and carry their electrical currents.

As the global submarine telegraphy business grew, so did demand for *Palaquium gutta* tree trunks. The historian John Tully describes how local Malay, Chinese, and Dayak workers were paid little for the dangerous work of felling the trees and slowly collecting the latex.[42] The latex was processed and then sold through Singapore's trade markets into the British market, where it was transformed into, among other things, lengths upon lengths of submarine cable sheaths that wrapped around the globe. As media scholar Nicole Starosielski writes, "Military strategists saw cables as the most efficient and secure mode of communication with the colonies—and, by implication, of control over them."[43] The routes of submarine cables today still mark out the early colonial networks between the centers and the peripheries of empire.[44]

A mature *Palaquium gutta* could yield around eleven ounces of latex. But in 1857, the first transatlantic cable was around eighteen hundred miles long and weighed two thousand tons—requiring about 250 tons of gutta-percha. To produce just one ton of this material required around nine hundred thousand tree trunks. The jungles of Malaysia and Singapore were stripped; by the early 1880s, the *Palaquium gutta* had vanished. In a last-ditch effort to save their supply chain, the British passed a ban in 1883 to halt harvesting the latex, but the tree was all but extinct.[45]

The Victorian environmental disaster of gutta-percha, at the dawn of the global information society, shows how the relations between technology and its materials, environments, and labor practices are interwoven.[46] Just as Victorians precipitated ecological disaster for their early cables, so do contemporary mining and global supply chains further imperil the delicate ecological balance of our era.

There are dark ironies in the prehistories of planetary computation. Currently large-scale AI systems are driving forms of environmental, data, and human extraction, but from the Victorian era onward, algorithmic computation emerged out of desires to manage and control war, population, and cli-

Palaquium gutta

mate change. The historian Theodora Dryer describes how the founding figure of mathematical statistics, English scientist Karl Pearson, sought to resolve uncertainties of planning and management by developing new data architectures including standard deviations and techniques of correlation and regression. His methods were, in turn, deeply imbricated with race science, as Pearson—along with his mentor, the statistician and founder of eugenics Sir Francis Galton—believed that statistics could be "the first step in an enquiry into the possible effect of a selective process upon any character of a race."[47]

As Dryer writes, "By the end of the 1930s, these data architectures—regression techniques, standard deviation, and correlations—would become dominant tools used in interpreting social and state information on the world stage. Tracking the nodes and routes of global trade, the interwar 'mathematical-statistics movement' became a vast enterprise."[48] This enterprise kept expanding after World War II, as new computational systems were used in domains such as weather forecasting during periods of drought to eke out more productivity from large-scale industrial farming.[49] From this perspective, algorithmic computing, computational statistics, and artificial intelligence were developed in the twentieth century to address social and environmental challenges but would later be used to intensify industrial extraction and exploitation and further deplete environmental resources.

The Myth of Clean Tech

Minerals are the backbone of AI, but its lifeblood is still electrical energy. Advanced computation is rarely considered in terms of carbon footprints, fossil fuels, and pollution; metaphors like "the cloud" imply something floating and delicate within a natural, green industry.[50] Servers are hidden in nondescript data centers, and their polluting qualities are far less visible than the billowing smokestacks of coal-fired power stations. The tech sector heavily publicizes its environmental policies, sustainability initiatives, and plans to address climate-related problems using AI as a problem-solving tool. It is all part of a highly produced public image of a sustainable tech industry with no carbon emissions. In reality, it takes a gargantuan amount of energy to run the computational infrastructures of Amazon Web Services or Microsoft's Azure, and the carbon footprint of the AI systems that run on those platforms is growing.[51]

As Tung-Hui Hu writes in *A Prehistory of the Cloud,* "The cloud is a resource-intensive, extractive technology that converts water and electricity into computational power, leaving a sizable amount of environmental damage that it then displaces from sight."[52] Addressing this energy-intensive infrastructure has become a major concern. Certainly, the industry has made significant efforts to make data centers more energy efficient and to increase their use of renewable energy. But already, the carbon footprint of the world's computational infrastructure has matched that of the aviation industry at its height, and it is increasing at a faster rate.[53] Estimates vary, with researchers like Lotfi Belkhir and Ahmed Elmeligi estimating that the tech sector will contribute 14 percent of global greenhouse emissions by 2040, while a team in Sweden predicts that the electricity demands of data centers alone will increase about fifteenfold by 2030.[54]

By looking closely at the computational capacity needed to build AI models, we can see how the desire for exponential increases in speed and accuracy is coming at a high cost to the planet. The processing demands of training AI models, and thus their energy consumption, is still an emerging area of investigation. One of the early papers in this field came from AI researcher Emma Strubell and her team at the University of Massachusetts Amherst in 2019. With a focus on trying to understand the carbon footprint of natural language processing (NLP) models, they began to sketch out potential estimates by running AI models over hundreds of thousands of computational hours.[55] The initial numbers were striking. Strubell's team found that running only a single NLP model produced more than 660,000 pounds of carbon dioxide emissions, the equivalent of five gas-powered cars over their total lifetime (including their manufacturing) or 125 round-trip flights from New York to Beijing.[56]

Worse, the researchers noted that this modeling is, at minimum, a baseline optimistic estimate. It does not reflect the true commercial scale at which companies like Apple and Amazon operate, scraping internet-wide datasets and feeding their own NLP models to make AI systems like Siri and Alexa sound more human. But the exact amount of energy consumption produced by the tech sector's AI models is unknown; that information is kept as highly guarded corporate secrets. Here, too, the data economy is premised on maintaining environmental ignorance.

In the AI field, it is standard practice to maximize computational cycles to improve performance, in accordance with a belief that bigger is better. As Rich Sutton of DeepMind describes it: "Methods that leverage computation are ultimately the most effective, and by a large margin."[57] The computational technique of brute-force testing in AI training runs, or systematically gathering more data and using more computational cycles until a better result is achieved, has driven a steep increase in energy consumption. OpenAI estimated that since 2012, the amount of compute used to train a single AI model has increased by a factor of ten every year. That's due to developers "repeatedly finding ways to use more chips in parallel, and being willing to pay the economic cost of doing so."[58] Thinking only in terms of economic cost narrows the view on the wider local and environmental price of burning computation cycles as a way to create incremental efficiencies. The tendency toward "compute maximalism" has profound ecological impacts.

Data centers are among the world's largest consumers of electricity.[59] Powering this multilevel machine requires grid electricity in the form of coal, gas, nuclear, or renewable energy. Some corporations are responding to growing alarm about the energy consumption of large-scale computation, with Apple and Google claiming to be carbon neutral (which means they offset their carbon emissions by purchasing credits) and Microsoft promising to become carbon negative by 2030. But workers within the companies have pushed for reductions in emissions across the board, rather than what they see as buying indulgences out of environmental guilt.[60] Moreover, Microsoft, Google, and Amazon all license their AI platforms, engineering workforces, and infrastructures to fossil fuel companies to help them locate and extract fuel from the ground, which further drives the industry most responsible for anthropogenic climate change.

Beyond the United States, more clouds of carbon dioxide are rising. China's data center industry draws 73 percent of its power from coal, emitting about 99 million tons of CO2 in 2018.[61] And electricity consumption from China's data center infrastructure is expected to increase by two-thirds by 2023.[62] Greenpeace has raised the alarm about the colossal energy demands of China's biggest technology companies, arguing that "China's leading tech companies, including Alibaba, Tencent, and GDS, must dramatically scale up clean energy procurement and disclose energy use data."[63] But the lasting impacts of coal-fired power are everywhere, exceeding any national boundaries. The planetary nature of resource extraction and its consequences goes well beyond what the nation-state was designed to address.

Water tells another story of computation's true cost. The history of water use in the United States is full of battles and secret deals, and as with computation, the deals made over water are kept close. One of the biggest U.S. data centers belongs to the National Security Agency (NSA) in Bluffdale, Utah. Open since late 2013, the Intelligence Community Comprehensive National Cybersecurity Initiative Data Center is impossible to visit directly. But by driving up through the adjacent suburbs, I found a cul-de-sac on a hill thick with sagebrush, and from there I was afforded a closer view of the sprawling 1.2-million-square-foot facility. The site has a kind of symbolic power of the next era of government data capture, having been featured in films like *Citizenfour* and pictured in thousands of news stories about the NSA. In person, though, it looks nondescript and prosaic, a giant storage container combined with a government office block.

The struggle over water began even before the data center was officially open, given its location in drought-parched Utah.[64] Local journalists wanted to confirm whether the estimated consumption of 1.7 million gallons of water per day was accurate, but the NSA initially refused to share usage data, redacted all details from public records, and claimed that its water use was a matter of national security. Antisurveillance activists created handbooks encouraging the end of material support of water and energy to surveillance, and they strategized that legal controls over water usage could help shut down the facility.[65] But the city of Bluffdale had already made a multiyear deal with the NSA, in which the city would sell water at rates well below the average in return for the promise of economic growth the facility might bring to the region.[66] The geopolitics of water are now deeply combined with the mechanisms and politics of data centers, computation, and power—in every sense. From the dry hillside that overlooks the NSA's data repository, all the contestation and obfuscation about water makes sense: this is a landscape with a limit, and water that is used to cool servers is being taken away from communities and habitats that rely on it to live.

Just as the dirty work of the mining sector was far removed from the companies and city dwellers who profited most, so the majority of data centers are far removed from major population hubs, whether in the desert or in semi-industrial exurbs. This contributes to our sense of the cloud being out of sight and abstracted away, when in fact it is material, affecting the environment and climate in ways that are far from being fully recognized and accounted for. The cloud is of the earth, and to keep it growing requires expanding resources and layers of logistics and transport that are in constant motion.

The Logistical Layer

So far, we have considered the material stuff of AI, from rare earth elements to energy. By grounding our analysis in the specific materialities of AI—the things, places, and people—we can better see how the parts are operating within broader systems of power. Take, for example, the global logistical machines that move minerals, fuel, hardware, workers, and consumer AI devices around the planet.[67] The dizzying spectacle of logistics and production displayed by companies like Amazon would not be possible without the development and widespread acceptance of a standardized metal object: the cargo container. Like submarine cables, cargo containers bind the industries of global communication, transport, and capital, a material exercise of what mathematicians call "optimal transport"—in this case, as an optimization of space and resources across the trade routes of the world.

Standardized cargo containers (themselves built from the basic earth elements of carbon and iron forged as steel) enabled the explosion of the modern shipping industry, which in turn made it possible to envision and model the planet as a single massive factory. The cargo container is the single unit of value—like a piece of Lego—that can travel thousands of miles before meeting its final destination as a modular part of a greater system of delivery. In 2017, the capacity of container ships in seaborne trade reached nearly 250 million deadweight tons of cargo, dominated by giant shipping companies including Maersk of Denmark, the Mediterranean Shipping Company of Switzerland, and France's CMA CGM Group, each owning hundreds of container vessels.[68] For these commercial ventures, cargo shipping is a relatively cheap way to navigate the vascular system of the global factory, yet it disguises far larger external costs. Just as they tend to neglect the physical realities and costs of AI infrastructure, popular culture and media rarely cover the shipping industry. The author Rose George calls this condition "sea blindness."[69]

In recent years, shipping vessels produced 3.1 percent of yearly global carbon dioxide emissions, more than the total produced by Germany.[70] In order to minimize their internal costs, most container shipping companies use low-grade fuel in enormous quantities, which leads to increased amounts of airborne sulfur and other toxic substances. One container ship is estimated to emit as much pollution as fifty million cars, and sixty thousand deaths every year are attributed indirectly to cargo-ship-industry pollution.[71]

Even industry-friendly sources like the World Shipping Council admit that thousands of containers are lost each year, sinking to the ocean floor or drifting loose.[72] Some carry toxic substances that leak into the oceans; others release thousands of yellow rubber ducks that wash ashore around the world over decades.[73] Typically, workers spend almost six months at sea, often with long working shifts and without access to external communications.

Here, too, the most severe costs of global logistics are borne by the Earth's atmosphere, the oceanic ecosystem and low-paid workers. The corporate imaginaries of AI fail to depict the lasting costs and long histories of the materials needed to build computational infrastructures or the energy required to power them. The rapid growth of cloud-based computation, portrayed as environmentally friendly, has paradoxically driven an expansion of the frontiers of resource extraction. It is only by factoring in these hidden costs, these wider collections of actors and systems, that we can understand what the shift toward increasing automation will mean. This requires working against the grain of how the technological imaginary usually works, which is completely untethered from earthly matters. Like running an image search of "AI," which returns dozens of pictures of glowing brains and blue-tinted binary code floating in space, there is a powerful resistance to engaging with the materialities of these technologies. Instead, we begin with the earth, with extraction, and with the histories of industrial power and then consider how these patterns are repeated in systems of labor and data.

AI as Megamachine

In the late 1960s, the historian and philosopher of technology Lewis Mumford developed the concept of the *megamachine* to illustrate how all systems, no matter how immense, consist of the work of many individual human actors.[74] For Mumford, the Manhattan Project was the defining modern megamachine whose intricacies were kept not only from the public but even from the thousands of people who worked on it at discrete, secured sites across the United States. A total of 130,000 workers operated in complete secrecy under the direction of the military, developing a weapon that would kill (by conservative estimates) 237,000 people when it hit Hiroshima and Nagasaki in 1945. The atomic bomb depended on a complex, secret chain of supply, logistics, and human labor.

Artificial intelligence is another kind of megamachine, a set of technological approaches that depend on industrial infrastructures, supply chains, and human labor that stretch around the globe but are kept opaque. We have seen how AI is much more than databases and algorithms, machine learning models and linear algebra. It is metamorphic: relying on manufacturing, transportation, and physical work; data centers and the undersea cables that trace lines between the continents; personal devices and their raw components; transmission signals passing through the air; datasets produced by scraping the internet; and continual computational cycles. These all come at a cost.

We have looked at the relations between cities and mines, companies and supply chains, and the topographies of extraction that connect them. The fundamentally intertwined nature of production, manufacturing, and logistics reminds us that the mines that drive AI are everywhere: not only sited in discrete locations but diffuse and scattered across the geography of the earth, in what Mazen Labban has called the "planetary mine."[75] This is not to deny the many specific locations where technologically driven mining is taking place. Rather, Labban observes that the planetary mine expands and reconstitutes extraction into novel arrangements, extending the practices of mines into new spaces and interactions around the world.

Finding fresh methods for understanding the deep material and human roots of AI systems is vital at this moment in history, when the impacts of anthropogenic climate change are already well under way. But that's easier said than done. In part, that's because many industries that make up the AI system chain conceal the ongoing costs of what they do. Furthermore, the scale required to build artificial intelligence systems is too complex, too obscured by intellectual property law, and too mired in logistical and technical complexity for us to see into it all. But the aim here is not to try and make these complex assemblages transparent: rather than trying to see *inside* them, we will be connecting *across* multiple systems to understand how they work in relation to each other.[76] Thus, our path

The ruins at Blair. Photograph by Kate Crawford

will follow the stories about the environmental and labor costs of AI and place them in context with the practices of extraction and classification braided throughout everyday life. It is by thinking about these issues together that we can work toward greater justice.

I make one more trip to Silver Peak. Before I reach the town, I pull the van over to the side of the road to read a weather-beaten sign. It's Nevada Historical Marker 174, dedicated to the creation and destruction of a small town called Blair. In 1906, the Pittsburgh Silver Peak Gold Mining Company bought up the mines in the area. Anticipating a boom, land speculators purchased all of the available plots near Silver Peak along with its water rights, driving prices to record artificial highs. So the mining company surveyed a couple of miles north and declared it the site for a new town: Blair. They built a hundred-stamp cyanide mill for leach mining, the biggest in the state, and laid the Silver Peak railroad that ran from Blair

Junction to the Tonopah and Goldfield main line. Briefly, the town thrived. Many hundreds of people came from all over for the jobs, despite the harsh working conditions. But with so much mining activity, the cyanide began to poison the ground, and the gold and silver seams began to falter and dry up. By 1918, Blair was all but deserted. It was all over within twelve years. The ruins are marked on a local map—just a forty-five-minute walk away.

It's a blazing hot day in the desert. The only sounds are the metallic reverberations of cicadas and the rumble of an occasional passenger jet. I decide to start up the hill. By the time I reach the collection of stone buildings at the top of the long dirt road, I'm exhausted from the heat. I take shelter inside the collapsed remains of what was once a gold miner's house. Not much is left: some broken crockery, shards of glass bottles, a few rusted tins. Back in Blair's lively years, multiple saloons thrived nearby and a two-story hotel welcomed visitors. Now it's a cluster of broken foundations.

Through the space where a window used to be, the view stretches all the way down the valley. I'm struck by the realization that Silver Peak will also be a ghost town soon. The current draw on the lithium mine is aggressive in response to the high demand, and no one knows how long it will last. The most optimistic estimate is forty years, but the end may come much sooner. Then the lithium pools under the Clayton Valley will be exsanguinated—extracted for batteries that are destined for landfill. And Silver Peak will return to its previous life as an empty and quiet place, on the edge of an ancient salt lake, now drained.

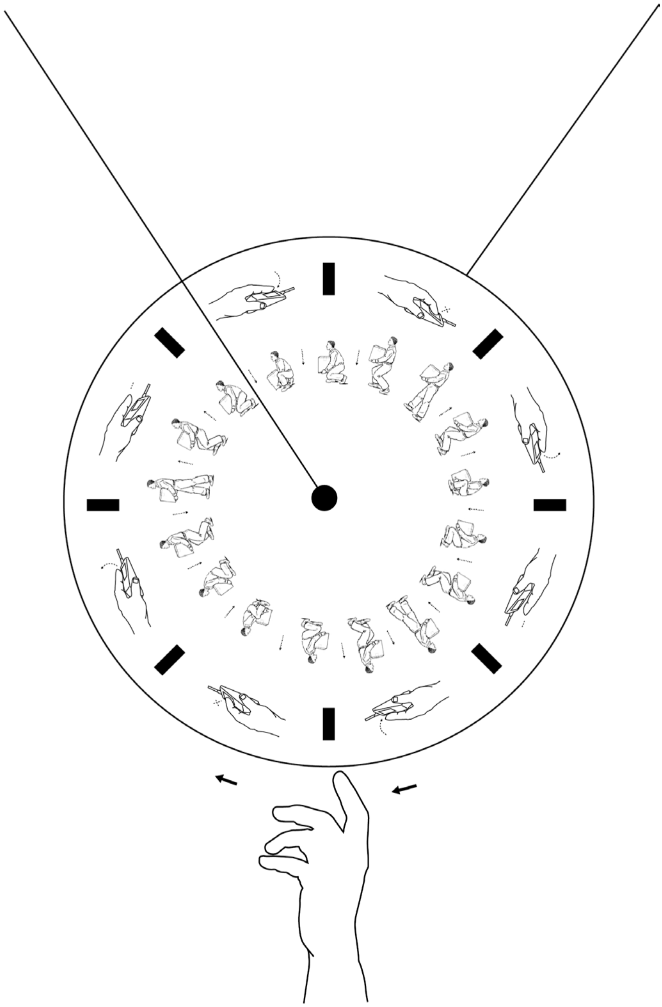

2
Labor

When I enter Amazon's vast fulfillment center in Robbinsville, New Jersey, the first thing I see is a large sign that reads "Time Clock." It juts out from one of the bright yellow concrete pylons spanning across the vast factory space of 1.2 million square feet. This is a major distribution warehouse for smaller objects—a central distribution node for the Northeastern United States. It presents a dizzying spectacle of contemporary logistics and standardization, designed to accelerate the delivery of packages. Dozens of time-clock signs appear at regular intervals along the entryway. Every second of work is being monitored and tallied. Workers—known as "associates"—must scan themselves in as soon as they arrive. The sparse, fluorescent-lit break rooms also feature time clocks—with more signs to underscore that all scans in and out of the rooms are tracked. Just as packages are scanned in the warehouse, so too are workers monitored for the greatest possible efficiency: they can only be off-task for fifteen minutes per shift, with an unpaid thirty-minute meal break. Shifts are ten hours long.

This is one of the newer fulfillment centers that feature robots to move the heavy shelving units laden with products in trays. The bright orange Kiva robots glide smoothly across the concrete floors like vivid water bugs, following a programmed logic that causes them to spin in lazy circles and then lock onto a path toward the next worker awaiting the trays. Then they move forward, carrying on their backs a tower of purchases that can weigh up to three thousand pounds. This shuffling army of ground-hugging robots presents a kind of effortless efficiency: they carry, they rotate, they advance, they repeat. They make a low, whirring hum, but it is almost entirely drowned out by the deafening sound of fast-moving conveyor belts that act as the factory's arteries. There are fourteen miles of conveyor belts moving without pause in this space. The result is a constant roar.

While the robots perform their coordinated algorithmic ballet behind bare chain-link fences, the workers in the factory are far less serene. The anxiety of making the "picking rate"—the number of items they must select and pack within the allocated time—is clearly taking a toll. Many of the workers I encounter on my visits are wearing some kind of support bandage. I see knee braces, elbow bandages, wrist guards. When I observe that many people seem to have injuries, the Amazon worker guiding me through the factory points to the vending machines spaced at regular intervals that are "stocked with over-the-counter painkillers for anyone who needs them."

Robotics has become a key part of Amazon's logistical armory, and while the machinery seems well tended, the corresponding human bodies seem like an afterthought. They are there to complete the specific, fiddly tasks that robots cannot: picking up and visually confirming all of the oddly shaped objects that people want delivered to their homes, from phone cases to dishwashing detergent, within the shortest amount of time. Humans are the necessary connective tissue to get

Workers and time clocks at the Amazon fulfillment center in Robbinsville Township, New Jersey. AP Photo/Julio Cortez

ordered items into containers and trucks and delivered to consumers. But they aren't the most valuable or trusted component of Amazon's machine. At the end of the day, all associates must exit through a row of metal detectors. This is an effective antitheft measure, I am told.

Within the layers of the internet, one of the most common units of measurement is the network packet—a basic unit of data to be sent from one destination and delivered to another. At Amazon, the basic unit of measurement is the brown cardboard box, that familiar domestic cargo vessel emblazoned with a curved arrow simulating a human smile. Network packets each have a timestamp known as a *time to live.* Data has to reach its destination before the time to live expires. At Amazon, the cardboard box also has a *time to live* driven by the customer's shipping demands. If the box is late, this affects Amazon's brand and ultimately its profits. So enormous attention has been devoted to the machine learning algorithm that is tuned to the data regarding the best size, weight, and strength of corrugated boxes and paper mailers. Apparently without irony, the algorithm is called "the matrix."[1] Whenever a person reports a broken item, it becomes a data point about what sort of box should be used in the future. The next time that product is mailed, it will automatically be assigned a new type of box by the matrix, without human input. This prevents breakages, which saves time, which increases profits. Workers, however, are forced continually to adapt, which makes it harder to put their knowledge into action or habituate to the job.

The control over time is a consistent theme in the Amazon logistical empire, and the bodies of workers are run according to the cadences of computational logics. Amazon is America's second-largest private employer, and many companies strive to emulate its approach. Many large corporations are heavily investing in automated systems in the attempt to extract ever-larger volumes of labor from fewer workers. Logics of efficiency, surveillance, and automation are all converging in the current turn to computational approaches to managing labor. The hybrid human-robotic distribution warehouses of Amazon are a key site to understand the trade-offs being made in this commitment to automated efficiency. From there, we can begin to consider the question of how labor, capital, and time are entwined in AI systems.

Rather than debating whether humans will be replaced by robots, in this chapter I focus on how the experience of work is shifting in relation to increased surveillance, algorithmic assessment, and the modulation of time. Put another way, instead of asking whether robots will replace humans, I'm interested in how humans are increasingly treated like robots and what this means for the role of labor. Many forms of work are shrouded in the term "artificial intelligence," hiding the fact that people are often performing rote tasks to shore up the impression that machines can do the work. But large-scale computation is deeply rooted in and running on the exploitation of human bodies.

If we want to understand the future of work in the context of artificial intelligence, we need to begin by understanding the past and present experience of workers. Approaches to maximizing the extraction of value from workers vary from reworkings of the classical techniques used in Henry Ford's factories to a range of machine learning–assisted tools designed to increase the granularity of tracking, nudging, and assessment. This chapter maps geographies of labor past and present, from Samuel Bentham's inspection houses to Charles Babbage's theories of time management and to Frederick Winslow Taylor's micromanagement of human bodies. Along the way, we will see how AI is built on the very human efforts of (among other things) crowdwork, the privatization of time, and the seemingly never-ending reaching, lifting, and toiling of putting boxes into order. From the lineage of the mechanized factory, a model emerges that values increased conformity, standardization, and interoperability—for products, processes, and humans alike.

Prehistories of Workplace AI

Workplace automation, though often told as a story of the future, is already a long-established experience of contemporary work. The manufacturing assembly line, with its emphasis on consistent and standardized units of production, has analogues in the service industries, from retail to restaurants. Secretarial labor has been increasingly automated since the 1980s and now is emulated by highly feminized AI assistants such as Siri, Cortana, and Alexa.[2] So-called knowledge workers, those white-collar employees assumed to be less threatened by the forces driving automation, find themselves increasingly subjected to workplace surveillance, process automation, and collapse between the distinction of work and leisure time (although women have rarely experienced such clear distinctions, as feminist theorists of work like Silvia Federici and Melissa Gregg have shown).[3] Work of all stripes has had to significantly adapt itself in order to be interpretable and understood by software-based systems.[4]

The common refrain for the expansion of AI systems and process automation is that we are living in a time of beneficial human-AI collaboration. But this collaboration is not fairly negotiated. The terms are based on a significant power asymmetry—is there ever a choice *not* to collaborate with algorithmic systems? When a company introduces a new AI platform, workers are rarely allowed to opt out. This is less of a collaboration than a forced engagement, where workers are expected to re-skill, keep up, and unquestioningly accept each new technical development.

Rather than representing a radical shift from established forms of work, the encroachment of AI into the workplace should properly be understood as a return to older practices of industrial labor exploitation that were well established in the 1890s and the early twentieth century. That was a time when factory labor was already seen in relation to machines and work tasks were increasingly subdivided into smaller actions requiring minimal skill but maximum exertion. Indeed, the current expansion of labor automation continues the broader historical dynamics inherent in industrial capitalism. Since the appearance of the earliest factories, workers have encountered ever more powerful tools, machines, and electronic systems that play a role in changing how labor is managed while transferring more value to their employers. We are witnessing new refrains on an old theme. The crucial difference is that employers now observe, assess, and modulate intimate parts of the work cycle and bodily data—down to the last micromovement—that were previously off-limits to them.

There are many prehistories of workplace AI; one is the Industrial Revolution's widespread automation of common productive activities. In his *Wealth of Nations,* the eighteenth-century political economist Adam Smith first pointed to the division and subdivision of manufacturing tasks as the basis of both improved productivity and increasing mechanization.[5] He observed that by identifying and analyzing the various steps involved in manufacturing any given item, it was possible to divide them into ever-smaller steps, so that a product once made entirely by expert craftspeople could now be built by a team of lower-skill workers equipped with tools purpose-built for a particular task. Thus, a factory's output could be scaled up significantly without an equivalent increase in labor cost.

Developments in mechanization were important, but it was only when combined with a growing abundance of energy derived from fossil fuels that they could drive a massive increase in the productive capacities of industrial societies. This increase in production occurred in tandem with a major transformation of the role of labor vis-à-vis machinery in the workplace. Initially conceived as labor-saving devices, factory machines were meant to assist workers with their daily activities but quickly became the center of productive activity, shaping the speed and character of work. Steam engines powered by coal and oil could drive continuous mechanical actions that influenced the pace of work in the factory. Work ceased to be primarily a product of human labor and took on an increasingly machinelike character, with workers adapting to the needs of the machine and its particular rhythms and cadences. Building on Smith, Karl Marx noted as early as 1848 that automation abstracts labor from the production of finished objects and turns a worker into "an appendage of the machine."[6]

The integration of workers' bodies with machines was sufficiently thorough that early industrialists could view their employees as a raw material to be managed and controlled like any other resource. Factory owners, using both their local political clout and paid muscle, sought to direct and restrict how their workers moved around within factory towns, sometimes even preventing workers from emigrating to less mechanized regions of the world.[7]

This also meant increasing control over time. The historian E. P. Thompson's formative essay explores how the Industrial Revolution demanded greater synchronization of work and stricter time disciplines.[8] The transition to industrial capitalism came with new divisions of labor, oversight, clocks, fines, and time sheets—technologies that also influenced the way people experienced time. Culture was also a powerful tool. During the eighteenth and nineteenth centuries, the propaganda about hard work came in the forms of pamphlets and essays on the importance of discipline and sermons on the virtues of early rising and working diligently for as long as possible.[9] The use of time came to be seen in both moral and economic terms: understood as a currency, time could be well spent or squandered away. But as more rigid time disciplines were imposed in workshops and factories, the more workers began to push back—campaigning over time itself. By the 1800s, labor movements were strongly advocating for reducing the working day, which could run as long as sixteen hours. Time itself became a key site for struggle.

Maintaining an efficient and disciplined workforce in the early factory necessitated new systems of surveillance and control. One such invention from the early years of industrial manufacturing was the inspection house, a circular arrangement that placed all of a factory's workers within sight of their supervisors, who worked from an office placed on a raised platform at the center of the building. Developed in the 1780s in Russia by the English naval engineer Samuel Bentham while under the employ of Prince Potemkin, this arrangement allowed expert supervisors to keep an eye on their untrained subordinates—mostly Russian peasants loaned to Bentham by Potemkin—for signs of poor working habits. It also allowed Bentham himself to keep an eye on the supervisors for signs of ill-discipline. The supervisors, mostly master shipbuilders recruited from England, caused Bentham great annoyance due to their tendency to drink and get into petty disagreements with one another. "Morning after morning I am taken up chiefly with disputes amongst my Officers," Bentham complained.[10] As his frustrations grew, he embarked on a redesign that would maximize his ability to keep a watchful eye on them, and on the system as a whole. With a visit from his elder brother, the utilitarian philosopher Jeremy Bentham, Samuel's inspection house became the inspiration for the famous panopticon, a design for a model prison featuring a central watchtower from which guards could supervise the prisoners in their cells.[11]

Since Michel Foucault's *Discipline and Punish,* it has become commonplace to consider the prison as the origin point of today's surveillance society, with the elder Bentham as its ideological progenitor. In fact, the panoptic prison owes its origins to the work of the younger Bentham in the context of the early manufacturing facility.[12] The panopticon began as a workplace mechanism well before it was conceptualized for prisons.

While Samuel Bentham's work on the inspection house has largely faded from our collective memory, the story behind it remains part of our shared lexicon. The inspection house was part of a strategy coordinated by Bentham's employer, Prince Potemkin, who wished to gain favor in Catherine the Great's court by demonstrating the potential for modernizing rural Russia and transforming the peasantry into a modern manufacturing workforce. The inspection house was built to serve as a spectacle for visiting dignitaries and financiers, much like the so-called Potemkin villages, which were little more than decorated facades designed to distract observers from the impoverished rural village landscapes discreetly obscured from view.

And this is just one genealogy. Many other histories of labor shaped these practices of observation and control. The plantation colonies of the Americas used forced labor to maintain cash crops like sugar, and slave owners depended on systems of constant surveillance. As Nicholas Mirzoeff describes in *The Right to Look,* a central role in the plantation economy was the overseer, who watched over the flow of production on the colonial slave plantation, and their oversight meant ordering the work of the slaves within a system of extreme violence.[13] As one planter described in 1814, the role of the overseer was "to never leave the slave for an instant in inaction; he keeps the fabrication of sugar under surveillance, never leaving the sugar-mill for an instant."[14] This regime of oversight also relied on bribing some slaves with food and clothing to enlist them as an expanded surveillance network and to maintain discipline and speed of work when the overseer was occupied.[15]

Now the role of oversight in the modern workplace is primarily deputized to surveillance technologies. The managerial class employs a wide range of technologies to surveil employees, including tracking their movements with apps, analyzing their social media feeds, comparing the patterns of replying to emails and booking meetings, and nudging them with suggestions to make them work faster and more efficiently. Employee data is used to make predictions about who is most likely to succeed (according to narrow, quantifiable parameters), who might be diverging from company goals, and who might be organizing other workers. Some use the techniques of machine learning, and others are more simplistic algorithmic systems. As workplace AI becomes more prevalent, many of the more basic monitoring and tracking systems are being expanded with new predictive capacities to become increasingly invasive mechanisms of worker management, asset control, and value extraction.

Potemkin AI and the Mechanical Turks

One of the less recognized facts of artificial intelligence is how many underpaid workers are required to help build, maintain, and test AI systems. This unseen labor takes many forms—supply-chain work, on-demand crowdwork, and traditional service-industry jobs. Exploitative forms of work exist at all stages of the AI pipeline, from the mining sector, where resources are extracted and transported to create the core infrastructure of AI systems, to the software side, where distributed workforces are paid pennies per microtask. Mary Gray and Sid Suri refer to such hidden labor as "ghost work."[16] Lilly Irani calls it "human-fueled automation."[17] These scholars have drawn attention to the experiences of crowdworkers or microworkers who perform the repetitive digital tasks that underlie AI systems, such as labeling thousands of hours of training data and reviewing suspicious or harmful content. Workers do the repetitive tasks that backstop claims of AI magic—but they rarely receive credit for making the systems function.[18]

Although this labor is essential to sustaining AI systems, it is usually very poorly compensated. A study from the United Nations International Labour Organization surveyed

3,500 crowdworkers from seventy-five countries who routinely offered their labor on popular task platforms like Amazon Mechanical Turk, Figure Eight, Microworkers, and Clickworker. The report found that a substantial number of people earned below their local minimum wage even though the majority of respondents were highly educated, often with specializations in science and technology.[19] Likewise, those who do content moderation work—assessing violent videos, hate speech, and forms of online cruelty for deletion—are also paid poorly. As media scholars such as Sarah Roberts and Tarleton Gillespie have shown, this kind of work can leave lasting forms of psychological trauma.[20]

But without this kind of work, AI systems won't function. The technical AI research community relies on cheap, crowd-sourced labor for many tasks that can't be done by machines. Between 2008 and 2016, the term "*crowdsourcing*" went from appearing in fewer than a thousand scientific articles to more than twenty thousand—which makes sense, given that Mechanical Turk launched in 2005. But during the same time frame, there was far too little debate about what ethical questions might be posed by relying on a workforce that is commonly paid far below the minimum wage.[21]

Of course, there are strong incentives to ignore the dependency on underpaid labor from around the world. All the work they do—from tagging images for computer-vision systems to testing whether an algorithm is producing the right results—refines AI systems much more quickly and cheaply, particularly when compared to paying students to do these tasks (as was the earlier tradition). So the issue has generally been ignored, and as one crowdwork research team observed, clients using these platforms "expect cheap, 'frictionless' completion of work without oversight, as if the platform were not an interface to human workers but a vast computer without living expenses."[22] In other words, clients treat human employees as little more than machines, because to recognize their work and compensate it fairly would make AI more expensive and less "efficient."

Sometimes workers are directly asked to pretend to be an AI system. The digital personal assistant start-up x.ai claimed that its AI agent, called Amy, could "magically schedule meetings" and handle many mundane daily tasks. But a detailed Bloomberg investigation by journalist Ellen Huet revealed that it wasn't artificial intelligence at all. "Amy" was carefully being checked and rewritten by a team of contract workers pulling long shifts. Similarly, Facebook's personal assistant, M, was relying on regular human intervention by a group of workers paid to review and edit every message.[23]

Faking AI is an exhausting job. The workers at x.ai were sometimes putting in fourteen-hour shifts of annotating emails in order to sustain the illusion that the service was automated and functioning 24/7. They couldn't leave at the end of the night until the queues of emails were finished. "I left feeling totally numb and absent of any sort of emotion," one employee told Huet.[24]

We could think of this as a kind of Potemkin AI—little more than facades, designed to demonstrate to investors and a credulous media what an automated system would look like while actually relying on human labor in the background.[25] In a charitable reading, these facades are an illustration of what the system might be capable of when fully realized, or a "minimum viable product" designed to demonstrate a concept. In a less charitable reading, Potemkin AI systems are a form of deception perpetrated by technology vendors eager to stake a claim in the lucrative tech space. But until there is another way to create large-scale AI that doesn't use extensive behind-the-curtain work by humans, this is a core logic of how AI works.

The writer Astra Taylor has described the kind of overselling of high-tech systems that aren't actually automated as "fauxtomation."[26] Automated systems appear to do work previously performed by humans, but in fact the system merely coordinates human work in the background. Taylor cites the examples of self-service kiosks in fast-food restaurants and self-checkout systems in supermarkets as places where an employee's labor appears to have been replaced by an automated system but where in fact the data-entry labor has simply been relocated from a paid employee to the customer. Meanwhile, many online systems that provide seemingly automated decisions, such as removing duplicated entries or deleting offensive content, are actually powered by humans working from home on endless queues of mundane tasks.[27] Much like Potemkin's decorated villages and model workshops, many valuable automated systems feature a combination of underpaid digital pieceworkers and consumers taking on unpaid tasks to make systems function. Meanwhile, companies seek to convince investors and the general public that intelligent machines are doing the work.

What is at stake in this artifice? The true labor costs of AI are being consistently downplayed and glossed over, but the forces driving this performance run deeper than merely marketing trickery. It is part of a tradition of exploitation and deskilling, where people must do more tedious and repetitive work to back-fill for automated systems, for a result that may be less effective or reliable than what it replaced. But this approach can *scale*—producing cost reductions and profit increases while obscuring how much it depends on remote workers being paid subsistence wages and off-loading additional tasks of maintenance or error-checking to consumers.

Fauxtomation does not directly replace human labor; rather, it relocates and disperses it in space and time. In so doing it increases the disconnection between labor and value and thereby performs an ideological function. Workers, having been alienated from the results of their work as well as disconnected from other workers doing the same job, are liable to be more easily exploited by their employers. This is evident from the extremely low rates of compensation crowdworkers receive around the world.[28] They and other kinds of fauxtomation laborers face the very real fact that their labor is interchangeable with any of the thousands of other workers who compete with them for work on platforms. At any point they could be replaced by another crowdworker, or possibly by a more automated system.

In 1770, Hungarian inventor Wolfgang von Kempelen constructed an elaborate mechanical chess player. He built a cabinet of wood and clockwork, behind which was seated a life-size mechanical man who could play chess against human opponents and win. This extraordinary contraption was first shown in the court of Empress Maria Theresa of Austria, then to visiting dignitaries and government ministers, all of whom were utterly convinced that this was an intelligent automaton. The lifelike machine was dressed in a turban, wide-legged pants, and a fur-trimmed robe to give the impression of an "oriental sorcerer."[29] This racialized appearance signaled exotic otherness, at a time when the elites of Vienna would drink Turkish coffee and dress their servants in Turkish costumes.[30] It came to be known as the Mechanical Turk. But the chess-playing automaton was an elaborate illusion: it had a human chess master hiding inside an internal chamber, operating the machine from within and completely out of sight.

Some 250 years later, the hoax lives on. Amazon chose to name its micropayment-based crowdsourcing platform "Amazon Mechanical Turk," despite the association with racism and trickery. On Amazon's platform, real workers remain out of sight in service of an illusion that AI systems are autonomous and magically intelligent.[31] Amazon's initial motivation to build Mechanical Turk emerged from the failures of its own artificial intelligence systems that could not adequately detect duplicate product pages on its retail site. After a series of futile and expensive attempts to solve the problem, the project engineers enlisted humans to fill the gaps in its streamlined systems.[32] Now Mechanical Turk connects businesses with an unseen and anonymous mass of workers who bid against one another for the opportunity to work on a series of microtasks. Mechanical Turk is a massively distributed workshop where humans emulate and improve on AI systems by checking and correcting algorithmic processes. This is what Amazon chief executive Jeff Bezos brazenly calls "*artificial* artificial intelligence."[33]

These examples of Potemkin AI are all around. Some are directly visible: when we see one of the current crop of self-driving cars on the streets, we also see a human operator in the driver's seat, ready to take control of the vehicle at the first sign of trouble. Others are less visible, as when we interact with a web-based chat interface. We engage only with the facades that obscure their inner workings, designed to hide the various combinations of machine and human labor in each interaction. We aren't informed whether we are receiving a response from the system itself or from a human operator paid to respond on its behalf.

If there is growing uncertainty about whether we are engaging with an AI system or not, the feeling is mutual. In a paradox that many of us have experienced, and ostensibly in order to prove true human identity when reading a website, we are required to convince Google's reCAPTCHA of our humanity. So we dutifully select multiple boxes containing street numbers, or cars, or houses. We are training Google's image recognition algorithms for free. Again, the myth of AI as affordable and efficient depends on layers of exploitation, including the extraction of mass unpaid labor to fine-tune the AI systems of the richest companies on earth.

Contemporary forms of artificial intelligence are neither artificial nor intelligent. We can—and should—speak instead of the hard physical labor of mine workers, the repetitive factory labor on the assembly line, the cybernetic labor in the cognitive sweatshops of outsourced programmers, the poorly paid crowdsourced labor of Mechanical Turk workers, and the unpaid immaterial work of everyday users. These are the places where we can see how planetary computation depends on the exploitation of human labor, all along all the supply chain of extraction.

Visions of Disassembly and Workplace Automation: Babbage, Ford, and Taylor

Charles Babbage is well known as the inventor of the first mechanical computer. In the 1820s, he developed the idea for the *Difference Engine,* a mechanical calculating machine designed to generate mathematical and astronomical tables in a fraction of the time required to calculate them by hand. By the 1830s, he had a viable conceptual design for the *Analytical Engine,* a programmable general-purpose mechanical computer, complete with a system of punch cards for providing it with instructions.[34]

Babbage also had a strong interest in liberal social theory and wrote extensively on the nature of labor—the combination of his interests in computation and worker automation. Following Adam Smith, he noted the division of labor as a means of streamlining factory work and generating efficiencies. He went further, however, arguing that the industrial corporation could be understood as an analogue to a computational system. Just like a computer, it included multiple specialized units performing particular tasks, all coordinated to produce a given body of work, but with the labor content of the finished product rendered largely invisible by the process as a whole.

In Babbage's more speculative writing, he imagined perfect flows of work through the system that could be visualized as data tables and monitored by pedometers and repeating clocks.[35] Through a combination of computation, surveillance, and labor discipline, he argued, it would be possible to enforce ever-higher degrees of efficiency and quality control.[36] It was a strangely prophetic vision. Only in very recent years, with the adoption of artificial intelligence in the workplace, has Babbage's unusual twin goals of computation and worker automation become possible at scale.

Babbage's economic thought extended outward from Smith's but diverged in one important way. For Smith, the economic value of an object was understood in relation to the cost of the labor required to produce it. In Babbage's rendering, however, value in a factory was derived from investment in the design of the manufacturing process rather than from the labor force of its employees. The real innovation was the logistical process, while workers simply enacted the tasks defined for them and operated the machines as instructed.

For Babbage, labor's role in the value production chain was largely negative: workers might fail to perform their tasks in the timely manner prescribed by the precision machines they operated, whether through poor discipline, injury, absenteeism, or acts of resistance. As noted by historian Simon Schaffer, "Under Babbage's gaze, factories looked like perfect engines and calculating machines like perfect computers. The workforce might be a source of trouble—it could make tables err or factories fail—but it could not be seen as a source of value."[37] The factory is conceived as a rational calculating machine with only one weakness: its frail and untrustworthy human labor force.

Babbage's theory was, of course, heavily inflected with a kind of financial liberalism, causing him to view labor as a problem that needed to be contained by automation. There was little consideration of the human costs of this automation or of how automation might be put to use to improve the working lives of factory employees. Instead, Babbage's idealized machinery aimed primarily to maximize financial returns to the plant owners and their investors. In a similar vein, today's proponents of workplace AI present a vision of production that prioritizes efficiency, cost-cutting, and higher profits instead of, say, assisting their employees by replacing repetitive drudge work. As Astra Taylor argues, "The kind of efficiency to which techno-evangelists aspire emphasizes standardization, simplification, and speed, not diversity, complexity, and interdependence."[38] This should not surprise us: it is a necessary outcome of the standard business model of for-profit companies where the highest responsibility is to shareholder value. We are living the result of a system in which companies must extract as much value as possible. Meanwhile, 94 percent of all new American jobs created between 2005 and 2015 were for "alternative work"—jobs that fall outside of full-time, salaried employment.[39] As companies reap the benefits of increasing automation, people are, on average, working longer hours, in more jobs, for less pay, in insecure positions.

The Meat Market

Among the first industries to implement the type of mechanized production line Babbage envisioned was the Chicago meat-packing industry in the 1870s. Trains brought livestock to the stockyard gates; the animals were funneled toward their slaughter in adjacent plants; and the carcasses were transported to various butchering and processing stations by means of a mechanized overhead trolley system, forming what came to be known as the *disassembly line.* The finished products could be shipped to faraway markets in specially designed refrigerated rail cars.[40] Labor historian Harry Braverman noted that the Chicago stockyards realized Babbage's vision of automation and division of labor so completely that the human techniques required at any point on the disassembly line could be performed by just about anyone.[41] Low-skill laborers could be paid the bare minimum and replaced at the first sign of trouble, themselves becoming as thoroughly commoditized as the packaged meats they produced.

When Upton Sinclair wrote *The Jungle,* his harrowing novel about working-class poverty, he set it in the meat-packing plants of Chicago. Although his intended point was to highlight the hardships of working immigrants in support of a socialist political vision, the book had an entirely different effect. The depictions of diseased and rotting meat prompted a public outcry over food safety and resulted in the passing of the Meat Inspection Act in 1906. But the focus on workers was lost. Powerful institutions from the meat-packing industry to Congress were prepared to intervene to improve the methods of production, but addressing the more fundamental exploitative labor dynamics that propped up the entire system was off limits. The persistence of this pattern underscores how power responds to critique: whether the product is cow carcasses or

Armour Beef dressing floor, 1952.
Courtesy Chicago Historical Society

facial recognition, the response is to accept regulation at the margins but to leave untouched the underlying logics of production.

Two other figures loom large in the history of workplace automation: Henry Ford, whose moving assembly line from the early twentieth century was inspired by Chicago's disassembly lines, and Frederick Winslow Taylor, the founder of scientific management. Taylor forged his career in the latter years of the nineteenth century developing a systematic approach to workplace management, one that focused on the minute movements of workers' bodies. Whereas Smith's and Babbage's notion of the division of labor was intended to provide a way to distribute work between people and tools, Taylor narrowed his focus to include microscopic subdivisions in the actions of each worker.

As the latest technology for precisely tracking time, the stopwatch was to become a key instrument of workplace surveillance for shop-floor supervisors and production engineers alike. Taylor used stopwatches to perform studies of workers that included detailed breakdowns of the time taken to perform the discrete physical motions involved in any given task. His *Principles of Scientific Management* established a system to quantify the movements of workers' bodies, with a view to deriving an optimally efficient layout of tools and working processes. The aim was maximum output at minimal cost.[42] It exemplified Marx's description of the domination of clock time, "Time is everything, man is nothing; he is, at most, time's carcass."[43]

Foxconn, the largest electronics manufacturing company in the world, which makes Apple iPhones and iPads, is a vivid example of how workers are reduced to animal bodies performing tightly controlled tasks. Foxconn became notorious for its rigid and militaristic management protocols after a spate of suicides in 2010.[44] Just two years later, the company's chairman, Terry Gou, described his more than one million employees this way: "As human beings are also animals, to manage one million animals gives me a headache."[45]

Controlling time becomes another way to manage bodies. In service and fast-food industries, time is measured down to the second. Assembly line workers cooking burgers at McDonald's are assessed for meeting such targets as five seconds to process screen-based orders, twenty-two seconds to assemble a sandwich, and fourteen seconds to wrap the food.[46] Strict adherence to the clock removes margin for error from the system. The slightest delay (a customer taking too long to order, a coffee machine failing, an employee calling in sick) can result in a cascading ripple of delays, warning sounds, and management notifications.

Even before McDonald's workers join the assembly line, their time is being managed and tracked. An algorithmic scheduling system incorporating historical data analysis and demand-prediction models determines workers' shift allocations, resulting in work schedules that can vary from week to week and even day to day. A 2014 class action lawsuit against McDonald's restaurants in California noted that franchisees are led by software that gives algorithmic predictions regarding employee-to-sales ratios and instructs managers to reduce staff quickly when demand drops.[47] Employees reported being told to delay clocking in to their shifts and instead to hang around nearby, ready to return to work if the restaurant started getting busy again. Because employees are paid only for time clocked in, the suit alleged that this amounted to significant wage theft on the part of the company and its franchisees.[48]

Algorithmically determined time allocations will vary from extremely short shifts of an hour or less to very long ones during busy times—whatever is most profitable. The algorithm doesn't factor in the human costs of waiting or getting to work only to be sent home or being unable to predict one's schedule and plan one's life. This time theft helps the efficiency of the company, but it comes at the direct cost of the employees.

Managing Time, Privatizing Time

Fast-food entrepreneur Ray Kroc, who helped turn McDonald's into a global franchise, joined the lineage of Smith, Babbage, Taylor, and Ford when he designed the standard sandwich assembly line and made his employees follow it unthinkingly. Surveillance, standardization, and the reduction of individual craft were central to Kroc's method. As labor researchers Clare Mayhew and Michael Quinlan argue with regard to the McDonald's standardized process, "The Fordist management system documented work and production tasks in minuscule detail. It required on-going documented participation and entailed detailed control of each individual's work process. There was an almost total removal of all conceptual work from execution of tasks."[49]

Minimizing the time spent at each station, or cycle time, became an object of intense scrutiny within the Fordist factory, with engineers dividing work tasks into ever-smaller pieces so they could be optimized and automated, and with supervisors disciplining workers whenever they fell behind. Supervisors, even Henry Ford himself, could often be seen walking the length of the factory, stopwatch in hand, recording cycle times and noting any discrepancies in a station's productivity.[50]

Now employers can passively surveil their workforce without walking out onto the factory floor. Instead, workers clock in to their shifts by swiping access badges or by presenting their fingerprints to readers attached to electronic time clocks. They work in front of timing devices that indicate the minutes or seconds left to perform the current task before a manager is notified. They sit at workstations fitted with sensors that continuously report on their body temperature, their physical distance from colleagues, the amount of time they spend browsing websites instead of performing assigned tasks, and so on. WeWork, the coworking behemoth that burned itself out over the course of 2019, quietly fitted its work spaces with surveillance devices in an effort to create new forms of data monetization. Its 2019 acquisition of the spatial analytics startup Euclid raised concerns, with the suggestion that it planned to track its paying members as they moved through their facilities.[51] Domino's Pizza has added to its kitchens machine-vision systems that inspect a finished pizza to ensure the staff made it according to prescribed standards.[52] Surveillance apparatuses are justified for producing inputs for algorithmic scheduling systems that further modulate work time, or to glean behavioral signals that may correlate with signs of high or low performance, or merely sold to data brokers as a form of insight.

In her essay "How Silicon Valley Sets Time," sociology professor Judy Wajcman argues that the aims of time-tracking tools and the demographic makeup of Silicon Valley are no coincidence.[53] Silicon Valley's elite workforce "is even younger, more masculine and more fully committed to working all hours," while also creating productivity tools that are premised on a kind of ruthless, winner-takes-all race to maximal efficiency.[54] This means that young, mostly male engineers, often unencumbered by time-consuming familial or community responsibilities, are building the tools that will police very different workplaces, quantifying the productivity and desirability of employees. The workaholism and round-the-clock hours often glorified by tech start-ups become an implicit benchmark against which other workers are measured, producing a vision of a standard worker that is masculinized, narrow, and reliant on the unpaid or underpaid care work of others.

Private Time

The coordination of time has become ever more granular in the technological forms of workplace management. For example, General Motors' Manufacturing Automation Protocol (MAP) was an early attempt to provide standard solutions to common manufacturing robot coordination problems, including clock synchronization.[55] In due course, other, more generic time synchronization protocols that could be delivered over ethernet and TCP/IP networks emerged, including the Network Time Protocol (NTP), and, later, the Precision Time

Protocol (PTP), each of which spawned a variety of competing implementations across various operating systems. Both NTP and PTP function by establishing a hierarchy of clocks across a network, with a “master” clock driving the “slave” clocks.

The master-slave metaphor is riddled throughout engineering and computation. One of the earliest uses of this racist metaphor dates back to 1904 describing astronomical clocks in a Cape Town observatory.[56] But it wasn’t until 1960s that the master-slave terminology spread, particularly after it was used in computing, starting with the Dartmouth timesharing system. Mathematicians John Kemeny and Thomas Kurtz developed a time-sharing program for access to computing resources after a suggestion by one of the early founders of AI, John McCarthy. As they wrote in *Science* in 1968, “First, all computing for users takes place in the slave computer, while the executive program (the ‘brains’ of the system) resides in the master computer. It is thus impossible for an erroneous or runaway user program in the slave computer to ‘damage’ the executive program and thereby bring the whole system to a halt.”[57] The problematic implication that control is equivalent to intelligence would continue to shape the AI field for decades. And as Ron Eglash has argued, the phrasing has a strong echo of the pre–Civil War discourse on runaway slaves.[58]

The master-slave terminology has been seen as offensive by many and has been removed from Python, a coding language common in machine learning, and Github, a software development platform. But it persists in one of the most expansive computational infrastructures in the world. Google’s Spanner—named as such because it spans the entire planet—is a massive, globally distributed, synchronously replicated database. It is the infrastructure that supports Gmail, Google search, advertising, and all of Google’s distributed services.

At this scale, functioning across the globe, Spanner synchronizes time across millions of servers in hundreds of data centers. Every data center has a "time master" unit that is always receiving GPS time. But because servers were polling a variety of master clocks, there was slight network latency and clock drift. How to resolve this uncertainty? The answer was to create a new distributed time protocol—a proprietary form of time—so that all servers could be in sync regardless of where they were across the planet. Google called this new protocol, without irony, TrueTime.

Google's TrueTime is a distributed time protocol that functions by establishing trust relationships between the local clocks of data centers so they can decide which peers to synchronize with. Benefiting from a sufficiently large number of reliable clocks, including GPS receivers and atomic clocks that provide an extremely high degree of precision, and from sufficiently low levels of network latency, TrueTime allows a distributed set of servers to guarantee that events can occur in a determinate sequence across a wide area network.[59]

What's most remarkable in this system of privatized Google time is how TrueTime manages uncertainty when there is clock drift on individual servers. "If the uncertainty is large, Spanner slows down to wait out that uncertainty," Google researchers explain.[60] This embodies the fantasy of slowing down time, of moving it at will, and of bringing the planet under a single proprietary time code. If we think of the human experience of time as something shifting and subjective, moving faster or slower depending on where we are and whom we are with, then this is a social experience of time. TrueTime is the ability to create a shifting timescale under the control of a centralized master clock. Just as Isaac Newton imagined an absolute form of time that exists independently of any perceiver, Google has invented its own form of universal time.

Proprietary forms of time have long been used to make machines run smoothly. Railroad magnates in the nineteenth century had their own forms of time. In New England in 1849, for example, all trains were to adopt "true time at Boston as given by William Bond & Son, No. 26 Congress Street."[61] As Peter Galison has documented, railroad executives weren't fond of having to switch to other times depending on which state their trains traveled to, and the general manager of the New York & New England Railroad Company called switching to other times "a nuisance and great inconvenience and no use to anybody I can see."[62] But after a head-on train collision killed fourteen people in 1853, there was immense pressure to coordinate all of the clocks using the new technology of the telegraph.

Like artificial intelligence, the telegraph was hailed as a unifying technology that would expand the capabilities of human beings. In 1889 Lord Salisbury boasted that the telegraph had "assembled all mankind upon one great plane."[63] Businesses, governments, and the military used the telegraph to compile time into a coherent grid, erasing more local forms of timekeeping. And the telegraph was dominated by one of the first great industrial monopolies, Western Union. In addition to altering the temporal and spatial boundaries of human interaction, communications theorist James Carey argues that the telegraph also enabled a new form of monopoly capitalism: "a new body of law, economic theory, political arrangements, management techniques, organizational structures, and scientific rationales with which to justify and make effective the development of a privately owned and controlled monopolistic corporation."[64] While this interpretation implies a kind of technological determinism in what was a complex series of developments, it is fair to say that the telegraph—paired with the transatlantic cable—enabled imperial powers to maintain more centralized control over their colonies.

The telegraph made time a central focus for commerce. Rather than traders exploiting the difference in prices between regions by buying low and selling high in varying locations, now they traded between time zones: in Carey's terms, a shift from space to time, from arbitrage to futures.[65] The privatized time zones of data centers are just the latest example. The infrastructural ordering of time acts as a kind of "macrophysics of power," determining new logics of information at a planetary level.[66] Such power is necessarily centralizing, creating orders of meaning that are extremely difficult to see, let alone disrupt.

Defiance of centralized time is a vital part of this history. In the 1930s, when Ford wanted more control over his global supply chain, he set up a rubber plantation and processing facility deep in the Brazilian rain forest, in a town he named Fordlandia. He employed local workers to process rubber for shipping back to Detroit, but his attempts to impose his tightly controlled manufacturing process on the local population backfired. Rioting workers tore apart the factory's time clocks, smashing the devices used to track the entry and exit of each worker in the plant.

Other forms of insurgence have centered on adding friction to the work process. The French anarchist Émile Pouget used the term "sabotage" to mean the equivalent of a "go slow" on the factory floor, when workers intentionally reduce their pace of work.[67] The objective was to withdraw efficiency, to reduce the value of time as a currency. Although there will always be ways to resist the imposed temporality of work, with forms of algorithmic and video monitoring, this becomes much harder—as the relation between work and time is observed at ever closer range.

From the fine modulations of time within factories to the big modulations of time at the scale of planetary computation networks, defining time is an established strategy for centralizing power. Artificial intelligence systems have allowed for greater exploitation of distributed labor around the world to take advantage of uneven economic topologies. Simultaneously, the tech sector is creating for itself a smooth global terrain of time to strengthen and speed its business objectives. Controlling time—whether via the clocks for churches, trains, or data centers—has always been a function of controlling the political order. But this battle for control has never been smooth, and it is a far-reaching conflict. Workers have found ways to intervene and resist, even when technological developments were forced on them or presented as desirable improvements, particularly if the only refinements were to increase surveillance and company control.

Setting the Rate

Amazon goes to great lengths to control what members of the public can see when they enter a fulfillment center. We are told about the fifteen-dollar-an-hour minimum wage and the perks for employees who can last longer than a year, and we are shown brightly lit break rooms that have Orwellian corporate slogans painted on the walls: “Frugality,” “Earn trust of others,” and “Bias for action.” The official Amazon guide cheerily explains what is happening at predetermined stops with rehearsed vignettes. Any questions about labor conditions are carefully answered to paint the most positive picture. But there are signs of unhappiness and dysfunction that are much harder to manage.

Out on the picking floor, where associates must pick up gray containers (known as “totes”) full of purchases to ship,

Fordlandia Time Clock, destroyed in the riot of December 1930.
From the Collections of The Henry Ford

whiteboards bear the marks of recent meetings. One had multiple complaints that the totes were stacked too high and that constantly reaching up to grab them was causing considerable pain and injuries. When asked about this, the Amazon guide quickly responded that this concern was being addressed by lowering the height of the conveyor belt in key sections. This was seen as a success: a complaint had been registered and action would be taken. The guide took this opportunity to explain for the second time that this was why unions were unnecessary here, because "associates have many opportunities to interface with their managers," and unionization only interferes with communication.[68]

But on the way out of the facility, I walked past a live feed of messages from workers on a large flat screen, with a sign above it that read, “The Voice of the Associates.” This was far less varnished. Messages scrolled rapidly past with complaints about arbitrary scheduling changes, the inability to book vacation time near holidays, and missing family occasions and birthdays. Pat responses from management seemed to be multiple variations on the theme of “We value your feedback.”

“Enough is enough. Amazon, we want you to treat us like humans, and not like robots.”[69] These are the words of Abdi Muse, executive director of the Awood Center in Minneapolis, a community organization that advocates for the working conditions of Minnesota’s East African populations. Muse is a soft-spoken defender of Amazon warehouse workers who are pushing for better working conditions. Many workers in his Minnesota community have been hired by Amazon, which actively recruited them and added sweeteners to the deal, such as free busing to work.

What Amazon didn’t advertise was “the rate”—the worker productivity metric driving the fulfillment centers that quickly became unsustainable and, according to Muse, inhumane. Workers began suffering high stress, injuries, and illness. Muse explained that if their rate went down three times they would be fired, no matter how long they had worked at the warehouse. Workers talked about having to skip bathroom breaks for fear that they would underperform.

But the day we met, Muse was optimistic. Even though Amazon explicitly discourages unions, informal groups of workers were springing up across the United States and staging protests. He smiled widely as he reported that the organizing was starting to have an impact. “Something incredible is happening,” he told me. “Tomorrow a group of Amazon workers will be walking off the job. It’s such a courageous group of women, and they are the real heroes.”[70] Indeed, that night, approximately sixty warehouse workers walked out of a delivery center in Eagan, Minnesota, wearing their mandated yellow vests. They were mostly women of Somali descent, and they held up signs in the rain, demanding such improvements as increased wages for night shifts and weight restrictions on boxes.[71] Only a few days earlier, Amazon workers in Sacramento, California, had protested the firing of an employee who had gone one hour over her bereavement leave after a family member died. Two weeks before that, more than a thousand Amazon workers staged the first ever white-collar walkout in the company's history over its massive carbon footprint.

Eventually, Amazon's representatives in Minnesota came to the table. They were happy to discuss many issues but never "the rate." "They said forget about 'the rate,'" recounted Muse. "We can talk about other issues, but the rate is our business model. We cannot change that."[72] The workers threatened to walk away from the table, and still Amazon would not budge. For both sides, "the rate" was the core issue, but it was also the hardest to alter. Unlike other local labor disputes where the on-the-ground supervisors might have been able to make concessions, the rate was set based on what the executives and tech workers in Seattle—far removed from the warehouse floor—had decided and had programmed Amazon's computational distribution infrastructure to optimize for. If the local warehouses were out of sync, Amazon's ordering of time was threatened. Workers and organizers started to see this as the real issue. They are shifting their focus accordingly toward building a movement across different factories and sectors of Amazon's workforce to address the core issues of power and centralization represented by the relentless rhythm of "the rate" itself.

These fights for time sovereignty, as we've seen, have a history. AI and algorithmic monitoring are simply the latest technologies in the long historical development of factories, timepieces, and surveillance architectures. Now many more sectors—from Uber drivers to Amazon warehouse workers to highly paid Google engineers—perceive themselves in this shared fight. This was strongly articulated by the executive director of the New York Taxi Workers Alliance, Bhairavi Desai, who put it this way: "Workers always know. They are out there building solidarity with each other, at red lights or in restaurants or in hotel queues, because they know that in order to prosper they have to band together."[73] Technologically driven forms of worker exploitation are a widespread problem in many industries. Workers are fighting against the logics of production and the order of time they must work within. The structures of time are never completely inhumane, but they are maintained right at the outer limit of what most people can tolerate.

Cross-sector solidarity in labor organizing is nothing new. Many movements, such as those led by traditional labor unions, have connected workers in different industries to win the victories of paid overtime, workplace safety, parental leave, and weekends. But as powerful business lobbies and neoliberal governments have chipped away at labor rights and protections over the past several decades and limited the avenues for worker organizing and communications, cross-sector support has become more difficult.[74] Now AI-driven systems of extraction and surveillance have become a shared locus for labor organizers to fight as a unified front.[75]

"We are all tech workers" has become a common sign at tech-related protests, carried by programmers, janitors, cafeteria workers, and engineers alike.[76] It can be read in multiple ways: it demands that the tech sector recognize the wide labor force it draws on to make its products, infrastructures, and workplaces function. It also reminds us that so many workers use laptops and mobile devices for work, engage on platforms like Facebook or Slack, and are subject to forms of workplace

AI systems for standardization, tracking, and assessment. This has set the stage for a form of solidarity built around tech work. But there are risks in centering tech workers and technology in what are more generalized and long-standing labor struggles. All kinds of workers are subject to the extractive technical infrastructures that seek to control and analyze time to its finest grain—many of whom have no identification with the technology sector or tech work at all. The histories of labor and automation remind us that what is at stake is producing more just conditions for every worker, and this broader goal should not depend on expanding the definition of tech work in order to gain legitimacy. We all have a collective stake in what the future of work looks like.

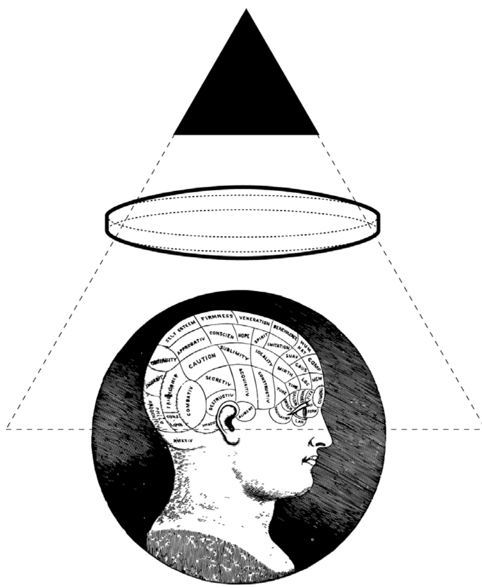
SELF ESTEEM
FIRMNESS
VENERATION
CONSCIEN
HOPE
APPROBATIV
CAUTION
SUBLIMITY
IDEALITY
IMITATION
MIRTH
SECRETIV
COMBATIV
DESTRUCTIV
ACQUISITIV
CONSTRUCTIV

4
Classification

I am surrounded by human skulls. This room contains almost five hundred, collected in the early decades of the 1800s. All are varnished, with numbers inscribed in black ink on the frontal bone. Delicate calligraphic circles mark out areas of the skull associated in phrenology with particular qualities, including "Benevolence" and "Veneration." Some bear descriptions in capital letters, with words like "Dutchman," "Peruvian of the Inca Race," or "Lunatic." Each was painstakingly weighed, measured, and labeled by the American craniologist Samuel Morton. Morton was a physician, natural historian, and member of the Academy of Natural Sciences of Philadelphia. He gathered these human skulls from around the world by trading with a network of scientists and skull hunters who brought back specimens for his experiments, sometimes by robbing graves.[1] By the end of his life in 1851, Morton had amassed more than a thousand skulls, the largest collection in the world at the time.[2] Much of the archive is now held in storage at the Physical Anthropology Section of the Penn Museum in Philadelphia.

Morton was not a classical phrenologist in that he didn't believe that human character could be read through examin-

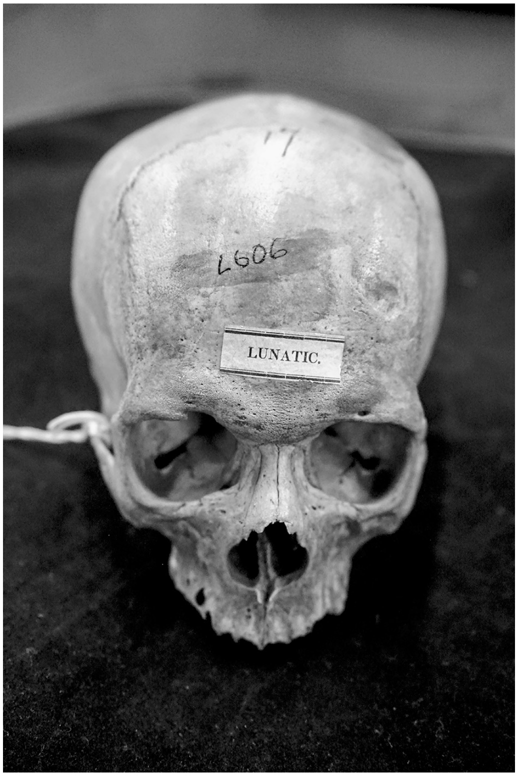

A skull from the Morton cranial collection marked "Lunatic."
Photograph by Kate Crawford

ing the shape of the head. Rather, his aim was to classify and rank human races "objectively" by comparing the physical characteristics of skulls. He did this by dividing them into the five "races" of the world: African, Native American, Caucasian, Malay, and Mongolian—a typical taxonomy of the time and a reflection of the colonialist mentality that dominated its geopolitics.[3] This was the viewpoint of polygenism—the belief that distinct human races had evolved separately at different times—legitimized by white European and American scholars and hailed by colonial explorers as a justification for racist violence and dispossession.[4] Craniometry grew to be one of their leading methods since it purported to assess human difference and merit accurately.[5]

Many of the skulls I see belong to people who were born in Africa but who died enslaved in the Americas. Morton measured these skulls by filling the cranial cavities with lead shot, then pouring the shot back into cylinders and gauging the volume of lead in cubic inches.[6] He published his results, comparing them to skulls he acquired from other locations: for example, he claimed that white people had the largest skulls, while Black people were on the bottom of the scale. Morton's tables of average skull volume by race were regarded as the cutting edge of science of the time. His work was cited for the rest of the century as objective, hard data that proved the relative intelligence of human races and biological superiority of the Caucasian race. This research was instrumented in the United States to maintain the legitimacy of slavery and racial segregation.[7] Considered the scientific state of the art at the time, it was used to authorize racial oppression long after the studies were no longer cited.

But Morton's work was not the kind of evidence it claimed to be. As Stephen Jay Gould describes in his landmark book *The Mismeasure of Man:*

> In short, and to put it bluntly, Morton's summaries are a patchwork of fudging and finagling in the clear interest of controlling *a priori* convictions. Yet—and this is the most intriguing aspect of his case—I find no evidence of conscious fraud. . . . The prevalence of unconscious finagling, on the other hand, suggests a general conclusion about the so-

> cial context of science. For if scientists can be honestly self-deluded to Morton's extent, then prior prejudice may be found anywhere, even in the basics of measuring bones and toting sums.[8]

Gould, and many others since, has reweighed the skulls and reexamined Morton's evidence.[9] Morton made errors and miscalculations, as well as procedural omissions, such as ignoring the basic fact that larger people have larger brains.[10] He selectively chose samples that supported his belief of white supremacy and deleted the subsamples that threw off his group averages. Contemporary assessments of the skulls at the Penn Museum show no significant differences among people—even when using Morton's data.[11] But prior prejudice—a way of seeing the world—had shaped what Morton believed was objective science and was a self-reinforcing loop that influenced his findings as much as the lead-filled skulls themselves.

Craniometry was, as Gould notes, "the leading numerical science of biological determinism during the nineteenth century" and was based on "egregious errors" in terms of the core underlying assumptions: that brain size equated to intelligence, that there are separate human races which are distinct biological species, and that those races could be placed in a hierarchy according to their intellect and innate character.[12] Ultimately, this kind of race science was debunked, but as Cornel West has argued, its dominant metaphors, logics, and categories not only supported white supremacy but also made specific political ideas about race possible while closing down others.[13]

Morton's legacy foreshadows epistemological problems with measurement and classification in artificial intelligence. Correlating cranial morphology with intelligence and claims to legal rights acts as a technical alibi for colonialism and slavery.[14] While there is a tendency to focus on the errors in skull measurements and how to correct for them, the far greater error is in the underlying worldview that animated this methodology. The aim, then, should be not to call for more accurate or "fair" skull measurements to shore up racist models of intelligence but to condemn the approach altogether. The practices of classification that Morton used were *inherently* political, and his invalid assumptions about intelligence, race, and biology had far-ranging social and economic effects.

The politics of classification is a core practice in artificial intelligence. The practices of classification inform how machine intelligence is recognized and produced from university labs to the tech industry. As we saw in the previous chapter, artifacts in the world are turned into data through extraction, measurement, labeling, and ordering, and this becomes—intentionally or otherwise—a slippery ground truth for technical systems trained on that data. And when AI systems are shown to produce discriminatory results along the categories of race, class, gender, disability, or age, companies face considerable pressure to reform their tools or diversify their data. But the result is often a narrow response, usually an attempt to address technical errors and skewed data to make the AI system appear more fair. What is often missing is a more fundamental set of questions: How does classification function in machine learning? What is at stake when we classify? In what ways do classifications interact with the classified? And what unspoken social and political theories underlie and are supported by these classifications of the world?

In their landmark study of classification, Geoffrey Bowker and Susan Leigh Star write that "classifications are powerful technologies. Embedded in working infrastructures they become relatively invisible without losing any of their power."[15] Classification is an act of power, be it labeling images in AI training sets, tracking people with facial recognition, or pouring lead shot into skulls. But classifications can disappear, as Bowker and Star observe, "into infrastructure, into habit, into the taken for granted."[16] We can easily forget that the classifications that are casually chosen to shape a technical system can play a dynamic role in shaping the social and material world.

The tendency to focus on the issue of bias in artificial intelligence has drawn us away from assessing the core practices of classification in AI, along with their attendant politics. To see that in action, in this chapter we'll explore some of the training datasets of the twenty-first century and observe how their schemas of social ordering naturalize hierarchies and magnify inequalities. We will also look at the limits of the bias debates in AI, where mathematical parity is frequently proposed to produce "fairer systems" instead of contending with underlying social, political, and economic structures. In short, we will consider how artificial intelligence uses classification to encode power.

Systems of Circular Logic

A decade ago, the suggestion that there could be a problem of bias in artificial intelligence was unorthodox. But now examples of discriminatory AI systems are legion, from gender bias in Apple's creditworthiness algorithms to racism in the COMPAS criminal risk assessment software and to age bias in Facebook's ad targeting.[17] Image recognition tools miscategorize Black faces, chatbots adopt racist and misogynistic language, voice recognition software fails to recognize female-sounding voices, and social media platforms show more highly paid job advertisements to men than to women.[18] As scholars like Ruha Benjamin and Safiya Noble have shown, there are hundreds of examples throughout the tech ecosystem.[19] Many more have never been detected or publicly admitted.

The typical structure of an episode in the ongoing AI bias narrative begins with an investigative journalist or whistleblower revealing how an AI system is producing discriminatory results. The story is widely shared, and the company in question promises to address the issue. Then either the system is superseded by something new, or technical interventions are made in the attempt to produce results with greater parity. Those results and technical fixes remain proprietary and secret, and the public is told to rest assured that the malady of bias has been "cured."[20] It is much rarer to have a public debate about *why* these forms of bias and discrimination frequently recur and whether more fundamental problems are at work than simply an inadequate underlying dataset or a poorly designed algorithm.

One of the more vivid examples of bias in action comes from an insider account at Amazon. In 2014, the company decided to experiment with automating the process of recommending and hiring workers. If automation had worked to drive profits in product recommendation and warehouse organization, it could, the logic went, make hiring more efficient. In the words of one engineer, "They literally wanted it to be an engine where I'm going to give you 100 resumes, it will spit out the top five, and we'll hire those."[21] The machine learning system was designed to rank people on a scale of one to five, mirroring Amazon's system of product ratings. To build the underlying model, Amazon's engineers used a dataset of ten years' worth of résumés from fellow employees and then trained a statistical model on fifty thousand terms that appeared in those résumés. Quickly, the system began to assign less importance to commonly used engineering terms, like programming languages, because everyone listed them in their job histories. Instead, the models began valuing more subtle cues that recurred on successful applications. A strong preference emerged for particular verbs. The examples the engineers mentioned were "executed" and "captured."[22]

Recruiters starting using the system as a supplement to their usual practices.[23] Soon enough, a serious problem emerged: the system wasn't recommending women. It was actively downgrading résumés from candidates who attended women's colleges, along with any résumés that even included the word "women." Even after editing the system to remove the influence of explicit references to gender, the biases remained. Proxies for hegemonic masculinity continued to emerge in the gendered use of language itself. The model was biased against women not just as a category but against commonly gendered forms of speech.

Inadvertently, Amazon had created a diagnostic tool. The vast majority of engineers hired by Amazon over ten years had been men, so the models they created, which were trained on the successful résumés of men, had learned to recommend men for future hiring. The employment practices of the past and present were shaping the hiring tools for the future. Amazon's system unexpectedly revealed the ways bias already existed, from the way masculinity is encoded in language, in résumés, and in the company itself. The tool was an intensification of the existing dynamics of Amazon and highlighted the lack of diversity across the AI industry past and present.[24]

Amazon ultimately shut down its hiring experiment. But the scale of the bias problem goes much deeper than a single system or failed approach. The AI industry has traditionally understood the problem of bias as though it is a bug to be fixed rather than a feature of classification itself. The result has been a focus on adjusting technical systems to produce greater quantitative parity across disparate groups, which, as we'll see, has created its own problems.

Understanding the relation between bias and classification requires going beyond an analysis of the production of knowledge—such as determining whether a dataset is biased or unbiased—and, instead, looking at the mechanics of knowledge construction itself, what sociologist Karin Knorr Cetina calls the "epistemic machinery."[25] To see that requires observing how patterns of inequality across history shape access to resources and opportunities, which in turn shape data. That data is then extracted for use in technical systems for classification and pattern recognition, which produces results that are perceived to be somehow objective. The result is a statistical ouroboros: a self-reinforcing discrimination machine that amplifies social inequalities under the guise of technical neutrality.

The Limits of Debiasing Systems

To better understand the limitations of analyzing AI bias, we can look to the attempts to fix it. In 2019, IBM tried to respond to concerns about bias in its AI systems by creating what the company described as a more "inclusive" dataset called Diversity in Faces (DiF).[26] DiF was part of an industry response to the groundbreaking work released a year earlier by researchers Joy Buolamwini and Timnit Gebru that had demonstrated that several facial recognition systems—including those by IBM, Microsoft, and Amazon—had far greater error rates for people with darker skin, particularly women.[27] As a result, efforts were ongoing inside all three companies to show progress on rectifying the problem.

"We expect face recognition to work accurately for each of us," the IBM researchers wrote, but the only way that the "challenge of diversity could be solved" would be to build "a data set comprised from the face of every person in the world."[28] IBM's researchers decided to draw on a preexisting dataset of a hundred million images taken from Flickr, the largest publicly available collection on the internet at the time.[29] They then used one million photos as a small sample and measured the craniofacial distances between landmarks in each face: eyes, nasal width, lip height, brow height, and so on. Like Morton measuring skulls, the IBM researchers sought to assign cranial measures and create categories of difference.

The IBM team claimed that their goal was to increase diversity of facial recognition data. Though well intentioned, the classifications they used reveal the politics of what diversity meant in this context. For example, to label the gender and age of a face, the team tasked crowdworkers to make subjective annotations, using the restrictive model of binary gender. Anyone who seemed to fall outside of this binary was removed from the dataset. IBM's vision of diversity emphasized the expansive options for cranial orbit height and nose bridges but discounted the existence of trans or gender nonbinary people. "Fairness" was reduced to meaning higher accuracy rates for machine-led facial recognition, and "diversity" referred to a wider range of faces to train the model. Craniometric analysis functions like a bait and switch, ultimately depoliticizing the idea of diversity and replacing it with a focus on *variation*. Designers get to decide what the variables are and how people are allocated to categories. Again, the practice of classification is centralizing power: the power to decide which differences make a difference.

IBM's researchers go on to state an even more problematic conclusion: "Aspects of our heritage—including race, ethnicity, culture, geography—and our individual identity—age, gender and visible forms of self-expression—are reflected in our faces."[30] This claim goes against decades of research that has challenged the idea that race, gender, and identity are biological categories at all but are better understood as politically, culturally, and socially constructed.[31] Embedding identity claims in technical systems as though they are facts observable from the face is an example of what Simone Browne calls "digital epidermalization," the imposition of race on the body. Browne defines this as the exercise of power when the disembodied gaze of surveillance technologies "do the work of alienating the subject by producing a 'truth' about the body and one's identity (or identities) despite the subject's claims."[32]

The foundational problems with IBM's approach to classifying diversity grow out of this kind of centralized production of identity, led by the machine learning techniques that were available to the team. Skin color detection is done because it can be, not because it says anything about race or produces a deeper cultural understanding. Similarly, the use of cranial measurement is done because it is a method that *can* be done with machine learning. The affordances of the tools become the horizon of truth. The capacity to deploy cranial measurements and digital epidermalization at scale drives a desire to find meaning in these approaches, even if this method has nothing to do with culture, heritage, or diversity. They are used to increase a problematic understanding of accuracy. Technical claims about accuracy and performance are commonly shot through with political choices about categories and norms but are rarely acknowledged as such.[33] These approaches are grounded in an ideological premise of biology as destiny, where our faces become our fate.

The Many Definitions of Bias

Since antiquity, the act of classification has been aligned with power. In theology, the ability to name and divide things was a divine act of God. The word "category" comes from the Ancient Greek *katēgoríā,* formed from two roots: *kata* (against) and *agoreuo* (speaking in public). In Greek, the word can be either a logical assertion or an accusation in a trial—alluding to both scientific and legal methods of categorization.

The historical lineage of "bias" as a term is much more recent. It first appears in fourteenth-century geometry, where it refers to an oblique or diagonal line. By the sixteenth century, it had acquired something like its current popular meaning, of "undue prejudice." By the 1900s, "bias" had developed a more technical meaning in statistics, where it refers to systematic differences between a sample and population, when the sample is not truly reflective of the whole.[34] It is from this statistical tradition that the machine learning field draws its understanding of bias, where it relates to a set of other concepts: generalization, classification, and variance.

Machine learning systems are designed to be able to generalize from a large training set of examples and to correctly classify new observations not included in the training datasets.[35] In other words, machine learning systems can perform a type of induction, learning from specific examples (such as past résumés of job applicants) in order to decide which data points to look for in new examples (such as word groupings in résumés from new applicants). In such cases, the term "bias" refers to a type of error that can occur during this predictive process of generalization—namely, a systematic or consistently reproduced classification error that the system exhibits when presented with new examples. This type of bias is often contrasted with another type of generalization error, variance, which refers to an algorithm's sensitivity to differences in training data. A model with high bias and low variance may be underfitting the data—failing to capture all of its significant features or signals. Alternatively, a model with high variance and low bias may be overfitting the data—building a model too close to the training data so that it potentially captures "noise" in addition to the data's significant features.[36]

Outside of machine learning, "bias" has many other meanings. For instance, in law, bias refers to a preconceived notion or opinion, a judgment based on prejudices, as opposed to a decision come to from the impartial evaluation of the facts of a case.[37] In psychology, Amos Tversky and Daniel Kahneman study "cognitive biases," or the ways in which human judgments deviate systematically from probabilistic expectations.[38] More recent research on implicit biases emphasizes the ways that unconscious attitudes and stereotypes "produce behaviors that diverge from a person's avowed or endorsed beliefs or principles."[39] Here bias is not simply a type of technical error; it also opens onto human beliefs, stereotypes, or forms of discrimination. These definitional distinctions limit the utility of "bias" as a term, especially when used by practitioners from different disciplines.

Technical designs can certainly be improved to better account for how their systems produce skews and discriminatory results. But the harder questions of why AI systems perpetuate forms of inequity are commonly skipped over in the rush to arrive at narrow technical solutions of statistical bias as though that is a sufficient remedy for deeper structural problems. There has been a general failure to address the ways in which the instruments of knowledge in AI reflect and serve the incentives of a wider extractive economy. What remains is a persistent asymmetry of power, where technical systems maintain and extend structural inequality, regardless of the intention of the designers.

Every dataset used to train machine learning systems, whether in the context of supervised or unsupervised machine learning, whether seen to be technically biased or not, contains a worldview. To create a training set is to take an almost infinitely complex and varied world and fix it into taxonomies composed of discrete classifications of individual data points, a process that requires inherently political, cultural, and social choices. By paying attention to these classifications, we can glimpse the various forms of power that are built into the architectures of AI world-building.

Training Sets as Classification Engines: The Case of ImageNet

In the last chapter we looked at the history of ImageNet and how this benchmark training set has influenced computer vision research since its creation in 2009. By taking a closer look at ImageNet's structure, we can begin to see how the dataset is ordered and its underlying logic for mapping the world of objects. ImageNet's structure is labyrinthine, vast, and filled with curiosities. The underlying semantic structure of ImageNet was imported from WordNet, a database of word classifications first developed at Princeton University's Cognitive Science Laboratory in 1985 and funded by the U.S. Office of Naval Research.[40] WordNet was conceived as a machine-readable dictionary, where users would search on the basis of semantic rather than alphabetic similarity. It became a vital source for the fields of computational linguistics and natural language processing. The WordNet team collected as many words as they could, starting with the Brown Corpus, a collection of one million words compiled in the 1960s.[41] The words in the Brown Corpus came from newspapers and a ramshackle collection of books including *New Methods of Parapsychology, The Family Fallout Shelter,* and *Who Rules the Marriage Bed?*[42]

WordNet attempts to organize the entire English language into synonym sets, or synsets. The ImageNet researchers selected only nouns, with the idea that nouns are things that pictures can represent—and that would be sufficient to train machines to automatically recognize objects. So Image-

Net's taxonomy is organized according to a nested hierarchy derived from WordNet, in which each synset represents a distinct concept, with synonyms grouped together (for example, "auto" and "car" are treated as belonging to the same set). The hierarchy moves from more general concepts to more specific ones. For example, the concept "chair" is found under artifact → furnishing → furniture → seat → chair. This classification system unsurprisingly evokes many prior taxonomical ranks, from the Linnaean system of biological classification to the ordering of books in libraries.

But the first indication of the true strangeness of ImageNet's worldview is its nine top-level categories that it drew from WordNet: plant, geological formation, natural object, sport, artifact, fungus, person, animal, and miscellaneous. These are curious categories into which all else must be ordered. Below that, it spawns into thousands of strange and specific nested classes, into which millions of images are housed like Russian dolls. There are categories for apples, apple butter, apple dumplings, apple geraniums, apple jelly, apple juice, apple maggots, apple rust, apple trees, apple turnovers, apple carts, and applesauce. There are pictures of hot lines, hot pants, hot plates, hot pots, hot rods, hot sauce, hot springs, hot toddies, hot tubs, hot-air balloons, hot fudge sauce, and hot water bottles. It is a riot of words, ordered into strange categories like those from Jorge Luis Borges's mythical encyclopedia.[43] At the level of images, it looks like madness. Some images are high-resolution stock photography, others are blurry phone photographs in poor lighting. Some are photos of children. Others are stills from pornography. Some are cartoons. There are pin-ups, religious icons, famous politicians, Hollywood celebrities, and Italian comedians. It veers wildly from the professional to the amateur, the sacred to the profane.

Human classifications are a good place to see these politics of classification at work. In ImageNet the category "human body" falls under the branch Natural Object → Body → Human Body. Its subcategories include "male body," "person," "juvenile body," "adult body," and "female body." The "adult body" category contains the subclasses "adult female body" and "adult male body." There is an implicit assumption here that only "male" and "female" bodies are recognized as "natural." There is an ImageNet category for the term "Hermaphrodite," but it is situated within the branch Person → Sensualist → Bisexual alongside the categories "Pseudohermaphrodite" and "Switch Hitter."[44]

Even before we look at the more controversial categories within ImageNet, we can see the politics of this classificatory scheme. The decisions to classify gender in this way are also naturalizing gender as a biological construct, which is binary, and transgender or gender nonbinary people are either nonexistent or placed under categories of sexuality.[45] Of course, this is not a novel approach. The classification hierarchy of gender and sexuality in ImageNet recalls earlier harmful forms of categorization, such as the classification of homosexuality as a mental disorder in the *Diagnostic and Statistical Manual.*[46] This deeply damaging categorization was used to justify subjecting people to repressive so-called therapies, and it took years of activism before the American Psychiatric Association removed it in 1973.[47]

Reducing humans into binary gender categories and rendering transgender people invisible or "deviant" are common features of classification schemes in machine learning. Os Keyes's study of automatic gender detection in facial recognition shows that almost 95 percent of papers in the field treat gender as binary, with the majority describing gender as immutable and physiological.[48] While some might respond that this can be easily remedied by creating more categories, this fails to address the deeper harm of allocating people into gender or race categories without their input or consent. This practice has a long history. Administrative systems for centuries have sought to make humans legible by applying fixed labels and definite properties. The work of essentializing and ordering on the basis of biology or culture has long been used to justify forms of violence and oppression.

While these classifying logics are treated as though they are natural and fixed, they are moving targets: not only do they affect the people being classified, but how they impact people in turn changes the classifications themselves. Hacking calls this the "looping effect," produced when the sciences engage in "making up people."[49] Bowker and Star also underscore that once classifications of people are constructed, they can stabilize a contested political category in ways that are difficult to see.[50] They become taken for granted unless they are actively resisted. We see this phenomenon in the AI field when highly influential infrastructures and training datasets pass as purely technical, whereas in fact they contain political interventions within their taxonomies: they naturalize a particular ordering of the world which produces effects that are seen to justify their original ordering.

The Power to Define "Person"

To impose order onto an undifferentiated mass, to ascribe phenomena to a category—that is, to name a thing—is in turn a means of reifying the existence of that category.

In the case of the 21,841 categories that were originally in the ImageNet hierarchy, noun classes such as "apple" or "apple butter" might seem reasonably uncontroversial, but not all nouns are created equal. To borrow an idea from linguist George Lakoff, the concept of an "apple" is a more *nouny* noun than the concept of "light," which in turn is more nouny than a concept such as "health."[51] Nouns occupy various places on an axis from the concrete to the abstract, from the descriptive to the judgmental. These gradients have been erased in the logic of ImageNet. Everything is flattened out and pinned to a label, like taxidermy butterflies in a display case. While this approach has the aesthetics of objectivity, it is nonetheless a profoundly ideological exercise.

For a decade, ImageNet contained 2,832 subcategories under the top-level category "Person." The subcategory with the most associated pictures was "gal" (with 1,664 images) followed by "grandfather" (1,662), "dad" (1,643), and chief executive officer (1,614—most of them male). With these highly populated categories, we can already begin to see the outlines of a worldview. ImageNet contains a profusion of classificatory categories, including ones for race, age, nationality, profession, economic status, behavior, character, and even morality.

There are many problems with the way ImageNet's taxonomy purports to classify photos of people with the logics of object recognition. Even though its creators removed some explicitly offensive synsets in 2009, categories remained for racial and national identities including Alaska Native, Anglo-American, Black, Black African, Black Woman (but not White Woman), Latin American, Mexican American, Nicaraguan, Pakistani, South American Indian, Spanish American, Texan, Uzbek, White, and Zulu. To present these as logical categories of organizing people is already troubling, even before they are used to classify people based on their appearance. Other people are labeled by careers or hobbies: there are Boy Scouts, cheerleaders, cognitive neuroscientists, hairdressers, intelligence analysts, mythologists, retailers, retirees, and so on. The existence of these categories suggests that people can be visually ordered according to their profession, in a way that seems reminiscent of such children's books as Richard Scarry's *What Do People Do All Day?* ImageNet also contains categories that make no sense whatsoever for image classification such as Debtor, Boss, Acquaintance, Brother, and Color-Blind Person. These are all nonvisual concepts that describe a relationship, be it to other people, to a financial system, or to the visual field itself. The dataset reifies these categories and connects them to images, so that similar images can be "recognized" by future systems.

Many truly offensive and harmful categories hid in the depths of ImageNet's Person categories. Some classifications were misogynist, racist, ageist, and ableist. The list includes Bad Person, Call Girl, Closet Queen, Codger, Convict, Crazy, Deadeye, Drug Addict, Failure, Flop, Fucker, Hypocrite, Jezebel, Kleptomaniac, Loser, Melancholic, Nonperson, Pervert, Prima Donna, Schizophrenic, Second-Rater, Slut, Spastic, Spinster, Streetwalker, Stud, Tosser, Unskilled Person, Wanton, Waverer, and Wimp. Insults, racist slurs, and moral judgments abound.

These offensive terms remained in ImageNet for ten years. Because ImageNet was typically used for object recognition—with "object" broadly defined—the specific Person category was rarely discussed at technical conferences, nor did it receive much public attention until the ImageNet Roulette project went viral in 2019: led by the artist Trevor Paglen, the project included an app that allowed people to upload images to see how they would be classified based on ImageNet's Person categories.[52] This focused considerable media attention on the influential collection's longtime inclusion of racist and sexist terms. The creators of ImageNet published a paper shortly afterward titled "Toward Fairer Datasets" that sought to "remove unsafe synsets." They asked twelve graduate students to flag any categories that seemed unsafe because they were either "inherently offensive" (for example, containing profanity or "racial or gender slurs") or "sensitive" (not inherently offensive but terms that "may cause offense when applied inappropriately, such as the classification of people based on sexual orientation and religion").[53] While this project sought to assess the offensiveness of ImageNet's categories by asking graduate students, the authors nonetheless continue to support the automated classification of people based on photographs despite the notable problems.

The ImageNet team ultimately removed 1,593 of 2,832 of the People categories—roughly 56 percent—deeming them "unsafe," along with the associated 600,040 images. The remaining half-million images were "temporarily deemed safe."[54] But what constitutes *safe* when it comes to classifying people? The focus on the hateful categories is not wrong, but it avoids addressing questions about the workings of the larger system. The entire taxonomy of ImageNet reveals the complexities and dangers of human classification. While terms like "microeconomist" or "basketball player" may initially seem less concerning than the use of labels like "spastic," "unskilled person," "mulatto," or "redneck," when we look at the people who are labeled in these categories we see many assumptions and stereotypes, including race, gender, age, and ability. In the metaphysics of ImageNet, there are separate image categories for "assistant professor" and "associate professor"—as though once someone gets a promotion, her or his biometric profile would reflect the change in rank.

In fact, there are no neutral categories in ImageNet, because the selection of images always interacts with the meaning of words. The politics are baked into the classificatory logic, even when the words aren't offensive. ImageNet is a lesson, in this sense, of what happens when people are categorized like objects. But this practice has only become more common in recent years, often inside the tech companies. The classification schemes used in companies like Facebook are much harder to investigate and criticize: proprietary systems offer few ways for outsiders to probe or audit how images are ordered or interpreted.

Then there is the issue of where the images in ImageNet's Person categories come from. As we saw in the last chapter, ImageNet's creators harvested images en masse from image search engines like Google, extracted people's selfies and vacation photos without their knowledge, and then paid Mechanical Turk workers to label and repackage them. All the skews and biases in how search engines return results are then informing the subsequent technical systems that scrape and label them. Low-paid crowdworkers are given the impossible task of making sense of the images at the rate of fifty per minute and fitting them into categories based on WordNet sysnets and Wikipedia definitions.[55] Perhaps it is no surprise that when we investigate the bedrock layer of these labeled images, we find that they are beset with stereotypes, errors, and absurdities. A woman lying on a beach towel is a "kleptomaniac," a teenager in a sports jersey is labeled a "loser," and an image of the actor Sigourney Weaver appears, classified as a "hermaphrodite."

Images—like all forms of data—are laden with all sorts of potential meanings, irresolvable questions, and contradictions. In trying to resolve these ambiguities, ImageNet's labels compress and simplify complexity. The focus on making training sets "fairer" by deleting offensive terms fails to contend with the power dynamics of classification and precludes a more thorough assessment of the underlying logics. Even if the worst examples are fixed, the approach is still fundamentally built on an extractive relationship with data that is divorced from the people and places from whence it came. Then it is rendered through a technical worldview that seeks to fuse together a form of singular objectivity from what are complex and varied cultural materials. The worldview of ImageNet is not unusual in this sense. In fact, it is typical of many AI training datasets, and it reveals many of the problems of top-down schemes that flatten complex social, cultural, political, and historical relations into quantifiable entities. This phenomenon is perhaps most obvious and insidious when it comes to the widespread efforts to classify people by race and gender in technical systems.

Constructing Race and Gender

By focusing on classification in AI, we can trace the ways that gender, race, and sexuality are falsely assumed to be natural, fixed, and detectable biological categories. Surveillance scholar Simone Browne observes, "There is a certain assumption with these technologies that categories of gender identity and race are clear cut, that a machine can be programmed to assign gender categories or determine what bodies and body parts should signify."[56] Indeed, the idea that race and gender can be automatically detectable in machine learning is treated as an assumed fact and rarely questioned by the technical disciplines, despite the profound political problems this presents.[57]

The UTKFace dataset (produced by a group at the University of Tennessee at Knoxville), for example, consists of more than twenty thousand images of faces with annotations for age, gender, and race.[58] The dataset's authors state that the dataset can be used for a variety of tasks, including automated face detection, age estimation, and age progression. The annotations for each image include an estimated age for each person, expressed in years from zero to 116. Gender is a forced binary: either zero for male or one for female. Second, race is categorized into five classes: White, Black, Asian, Indian, and Others. The politics of gender and race here are as obvious as they are harmful. Yet these kinds of dangerously reductive categorizations are widely used across many human-classifying training sets and have been part of the AI production pipelines for years.

UTKFace's narrow classificatory schema echoes the problematic racial classifications of the twentieth century, such as South Africa's apartheid system. As Bowker and Star have detailed, the South African government passed legislation in the 1950s that created a crude racial classification scheme to divide citizens into the categories of "Europeans, Asiatics, persons of mixed race or coloureds, and 'natives' or pure-blooded individuals of the Bantu race."[59] This racist legal regime governed people's lives, overwhelmingly those of Black South Africans whose movements were restricted and who were forcibly removed from their land. The politics of racial classification extended into the most intimate parts of people's lives. Interracial sexuality was forbidden, leading to more than 11,500 convictions by 1980, mostly of nonwhite women.[60] The complex centralized database for these classifications was designed and maintained by IBM, but the firm often had to rearrange the system and reclassify people, because in practice there were no singular pure racial categories.[61]

Above all, these systems of classification have caused enormous harm to people, and the concept of a pure "race" signifier has always been in dispute. In her writing about race, Donna Haraway observes, "In these taxonomies, which are, after all, little machines for clarifying and separating categories, the entity that always eluded the classifier was simple: race itself. The pure Type, which animated dreams, sciences, and terrors, kept slipping through, and endlessly multiplying, all the typological taxonomies."[62] Yet in dataset taxonomies, and in the machine learning systems that train on them, the myth of the pure type has emerged once more, claiming the authority of science. In an article on the dangers of facial recognition, media scholar Luke Stark notes that "by introducing a variety of classifying logics that either reify existing racial categories or produce new ones, the automated pattern-generating logics of facial recognition systems both reproduce systemic inequality and exacerbate it."[63]

Some machine learning methods go beyond predicting age, gender, and race. There have been highly publicized efforts to detect sexuality from photographs on dating sites and criminality based on headshots from drivers' licenses.[64] These approaches are deeply problematic for many reasons, not least of which is that characteristics such as "criminality"—like race and gender—are profoundly relational, socially determined categories. These are not inherent features that are fixed; they are contextual and shifting depending on time and place. To make such predictions, machine learning systems are seeking to classify entirely relational things into fixed categories and are rightly critiqued as scientifically and ethically problematic.[65]

Machine learning systems are, in a very real way, *constructing* race and gender: they are defining the world within the terms they have set, and this has long-lasting ramifications for the people who are classified. When such systems are hailed as scientific innovations for predicting identities and future actions, this erases the technical frailties of how the systems were built, the priorities of why they were designed, and the many political processes of categorization that shape them. Disability scholars have long pointed to the ways in which so-called normal bodies are classified and how that has worked to stigmatize difference.[66] As one report notes, the history of disability itself is a "story of the ways in which various systems of classification (i.e., medical, scientific, legal) interface with social institutions and their articulations of power and knowledge."[67] At multiple levels, the act of defining categories and ideas of normalcy creates an outside: forms of abnormality, difference, and otherness. Technical systems are making political and normative interventions when they give names to something as dynamic and relational as personal identity, and they commonly do so using a reductive set of possibilities of what it is to be human. That restricts the range of how people are understood and can represent themselves, and it narrows the horizon of recognizable identities.

As Ian Hacking observes, classifying people is an imperial imperative: subjects were classified by empires when they were conquered, and then they were ordered into "a kind of people" by institutions and experts.[68] These acts of naming were assertions of power and colonial control, and the negative effects of those classifications can outlast the empires themselves. Classifications are technologies that produce and limit ways of knowing, and they are built into the logics of AI.

The Limits of Measurement

So what is to be done? If so much of the classificatory strata in training data and technical systems are forms of power and politics represented as objective measurement, how should we go about redressing this? How should system designers account for, in some cases, slavery, oppression, and hundreds of years of discrimination against some groups to the benefit of others? In other words, how should AI systems make representations of the social?

Making these choices about which information feeds AI systems to produce new classifications is a powerful moment of decision making: but who gets to choose and on what basis? The problem for computer science is that justice in AI systems will never be something that can be coded or computed. It requires a shift to assessing systems beyond optimization metrics and statistical parity and an understanding of where the frameworks of mathematics and engineering are causing the problems. This also means understanding how AI systems interact with data, workers, the environment, and the individuals whose lives will be affected by its use and deciding where AI should not be used.

Bowker and Star conclude that the sheer density of the collisions of classification schemes calls for a new kind of approach, a sensitivity to the "topography of things such as the distribution of ambiguity; the fluid dynamics of how classification systems meet up—a plate tectonics rather than static geology."[69] But it also requires attending to the uneven allocations of advantage and suffering, for "how these choices are made, and how we may think about that invisible matching process, is at the core of the ethical project."[70] Nonconsensual classifications present serious risks, as do normative assumptions about identity, yet these practices have become standard. That must change.

In this chapter we've seen how classificatory infrastructures contain gaps and contradictions: they necessarily reduce complexity, and they remove significant context, in order to make the world more computable. But they also proliferate in machine learning platforms in what Umberto Eco called "chaotic enumeration."[71] At a certain level of granularity, like and unlike things become sufficiently commensurate so that their similarities and differences are machine readable, yet in actuality their characteristics are uncontainable. Here, the issues go far beyond whether something is classified wrong or classified right. We are seeing strange, unpredictable twists as machine categories and people interact and change each other, as they try to find legibility in the shifting terrain, to fit the right categories and be spiked into the most lucrative feeds. In a machine learning landscape, these questions are no less urgent because they are hard to see. What is at stake is not just a historical curiosity or the odd feeling of a mismatch between the dotted-outline profiles we may glimpse in our platforms and feeds. Each and every classification has its consequence.

The histories of classification show us that the most harmful forms of human categorization—from the Apartheid system to the pathologization of homosexuality—did not simply fade away under the light of scientific research and ethical critique. Rather, change also required political organizing, sustained protest, and public campaigning over many years. Classificatory schemas enact and support the structures of power that formed them, and these do not shift without considerable effort. In Frederick Douglass's words, "Power concedes nothing without a demand. It never did and it never will."[72] Within the invisible regimes of classification in machine learning, it is harder to make demands and oppose their internal logics.

The training sets that are made public—such as ImageNet, UTKFace, and DiF—give us some insight into the kinds of categorizations that are propagating across industrial AI systems and research practices. But the truly massive engines of classification are the ones being operated at a global scale by private technology companies, including Facebook, Google, TikTok, and Baidu. These companies operate with little oversight into how they categorize and target users, and they fail to offer meaningful avenues for public contestation. When the matching processes of AI are truly hidden and people are kept unaware of why or how they receive forms of advantage or disadvantage, a collective political response is needed—even as it becomes more difficult.

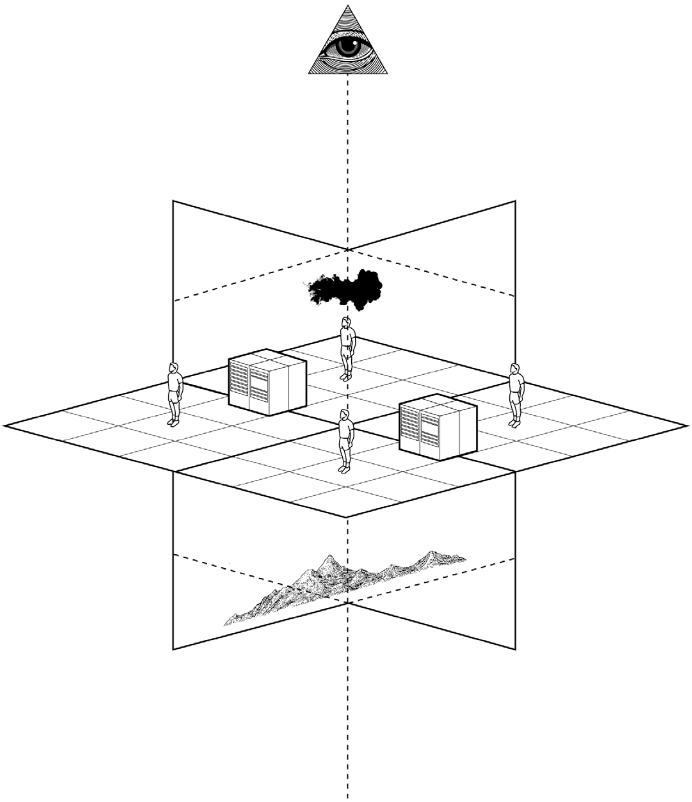

Conclusion

Power

Artificial intelligence is not an objective, universal, or neutral computational technique that makes determinations without human direction. Its systems are embedded in social, political, cultural, and economic worlds, shaped by humans, institutions, and imperatives that determine what they do and how they do it. They are designed to discriminate, to amplify hierarchies, and to encode narrow classifications. When applied in social contexts such as policing, the court system, health care, and education, they can reproduce, optimize, and amplify existing structural inequalities. This is no accident: AI systems are built to see and intervene in the world in ways that primarily benefit the states, institutions, and corporations that they serve. In this sense, AI systems are expressions of power that emerge from wider economic and political forces, created to increase profits and centralize control for those who wield them. But this is not how the story of artificial intelligence is typically told.

The standard accounts of AI often center on a kind of algorithmic exceptionalism—the idea that because AI sys-

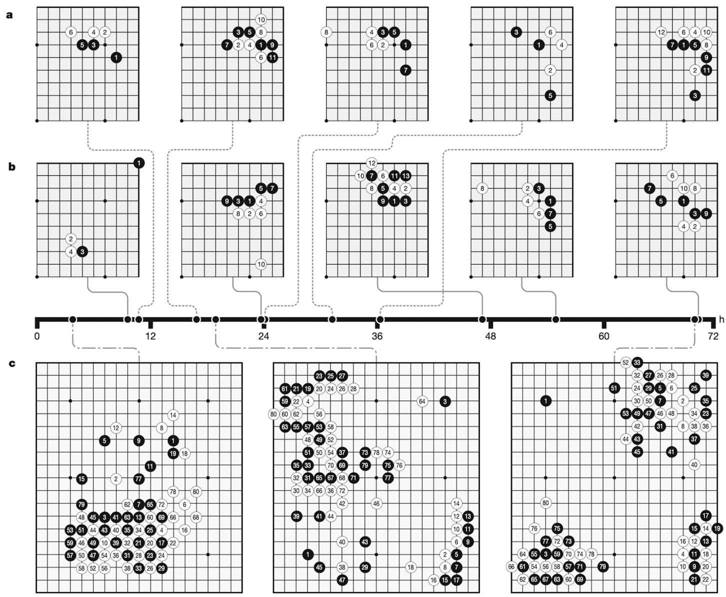

Go knowledge learned by AlphaGo Zero. Courtesy of DeepMind

tems can perform uncanny feats of computation, they must be smarter and more objective than their flawed human creators. Consider this diagram of AlphaGo Zero, an AI program designed by Google's DeepMind to play strategy games.[1] The image shows how it "learned" to play the Chinese strategy game Go by evaluating more than a thousand options per move. In the paper announcing this development, the authors write: "Starting *tabula rasa,* our new program AlphaGo Zero achieved superhuman performance."[2] DeepMind cofounder Demis Hassabis has described these game engines as akin to an alien intelligence. "It doesn't play like a human, but it also doesn't play like computer engines. It plays in a third, almost alien, way. . . . It's like chess from another dimension."[3] When the next iteration mastered Go within three days, Hassabis described it as "rediscovering three thousand years of human knowledge in 72 hours!"[4]

The Go diagram shows no machines, no human workers, no capital investment, no carbon footprint, just an abstract rules-based system endowed with otherworldly skills. Narratives of magic and mystification recur throughout AI's history, drawing bright circles around spectacular displays of speed, efficiency, and computational reasoning.[5] It's no coincidence that one of the iconic examples of contemporary AI is a game.

Games without Frontiers

Games have been a preferred testing ground for AI programs since the 1950s.[6] Unlike everyday life, games offer a closed world with defined parameters and clear victory conditions. The historical roots of AI in World War II stemmed from military-funded research in signal processing and optimization that sought to simplify the world, rendering it more like a strategy game. A strong emphasis on rationalization and prediction emerged, along with a faith that mathematical formalisms would help us understand humans and society.[7] The belief that accurate prediction is fundamentally about reducing the complexity of the world gave rise to an implicit theory of the social: find the signal in the noise and make order from disorder.

This epistemological flattening of complexity into clean signal for the purposes of prediction is now a central logic of machine learning. The historian of technology Alex Campolo and I call this *enchanted determinism:* AI systems are seen as enchanted, beyond the known world, yet deterministic in that they discover patterns that can be applied with predictive certainty to everyday life.[8] In discussions of deep learning systems, where machine learning techniques are extended by layering abstract representations of data on top of each other, enchanted determinism acquires an almost theological quality. That deep learning approaches are often uninterpretable, even to the engineers who created them, gives these systems an aura of being too complex to regulate and too powerful to refuse. As the social anthropologist F. G. Bailey observed, the technique of "obscuring by mystification" is often employed in public settings to argue for a phenomenon's inevitability.[9] We are told to focus on the innovative nature of the method rather than on what is primary: the purpose of the thing itself. Above all, enchanted determinism obscures power and closes off informed public discussion, critical scrutiny, or outright rejection.

Enchanted determinism has two dominant strands, each a mirror image of the other. One is a form of tech utopianism that offers computational interventions as universal solutions applicable to any problem. The other is a tech dystopian perspective that blames algorithms for their negative outcomes as though they are independent agents, without contending with the contexts that shape them and in which they operate. At an extreme, the tech dystopian narrative ends in the singularity, or superintelligence—the theory that a machine intelligence could emerge that will ultimately dominate or destroy humans.[10] This view rarely contends with the reality that so many people around the world are *already* dominated by systems of extractive planetary computation.

These dystopian and utopian discourses are metaphysical twins: one places its faith in AI as a solution to every problem, while the other fears AI as the greatest peril. Each offers a profoundly ahistorical view that locates power solely within technology itself. Whether AI is abstracted as an all-purpose tool or an all-powerful overlord, the result is technological determinism. AI takes the central position in society's redemption or ruin, permitting us to ignore the systemic forces of unfettered neoliberalism, austerity politics, racial inequality, and widespread labor exploitation. Both the tech utopians and dystopians frame the problem with technology always at the center, inevitably expanding into every part of life, decoupled from the forms of power that it magnifies and serves.

When AlphaGo defeats a human grandmaster, it's tempting to imagine that some kind of otherworldly intelligence has arrived. But there's a far simpler and more accurate explanation. AI game engines are designed to play millions of games, run statistical analyses to optimize for winning outcomes, and then play millions more. These programs produce surprising moves uncommon in human games for a straightforward reason: they can play and analyze far more games at a far greater speed than any human can. This is not magic; it is statistical analysis at scale. Yet the tales of preternatural machine intelligence persist.[11] Over and over, we see the ideology of Cartesian dualism in AI: the fantasy that AI systems are disembodied brains that absorb and produce knowledge independently from their creators, infrastructures, and the world at large. These illusions distract from the far more relevant questions: Whom do these systems serve? What are the political economies of their construction? And what are the wider planetary consequences?

The Pipelines of AI

Consider a different illustration of AI: the blueprint for Google's first owned and operated data center, in The Dalles, Oregon. It depicts three 68,680-square-foot buildings, an enormous facility that was estimated in 2008 to use enough energy to power eighty-two thousand homes, or a city the size of Tacoma, Washington.[12] The data center now spreads along the shores of the Columbia River, where it draws heavily on

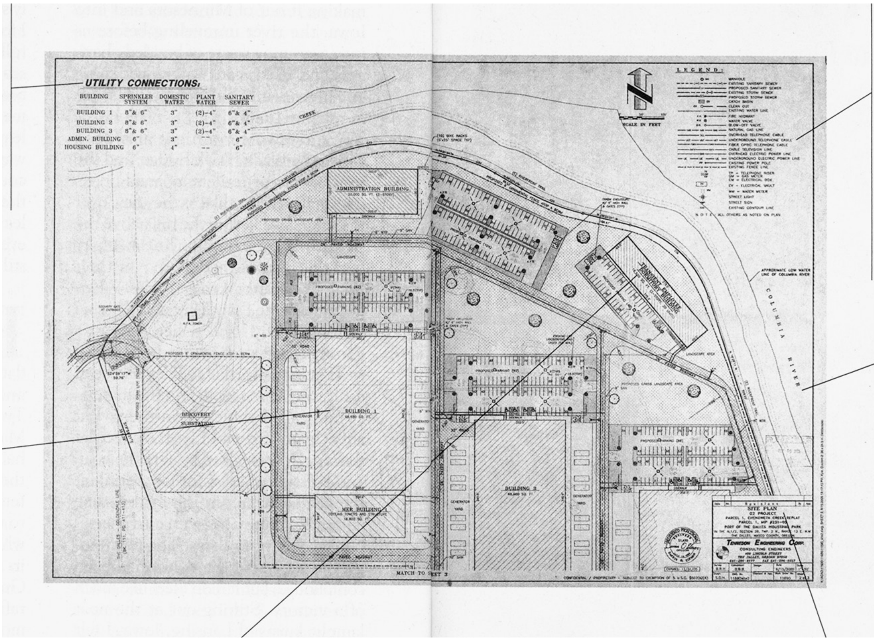

Blueprint of Google Data Center. Courtesy of *Harper's*

some of the cheapest electricity in North America. Google's lobbyists negotiated for six months with local officials to get a deal that included tax exemptions, guarantees of cheap energy, and use of the city-built fiber-optic ring. Unlike the abstract vision of a Go game, the engineering plan reveals how much of Google's technical vision depends on public utilities, including gas mains, sewer pipes, and the high-voltage lines through which the discount electricity would flow. In the words of the writer Ginger Strand, "Through city infrastructure, state give-backs, and federally subsidized power, YouTube is bankrolled by us."[13]

The blueprint reminds us of how much the artificial intelligence industry's expansion has been publicly subsidized:

from defense funding and federal research agencies to public utilities and tax breaks to the data and unpaid labor taken from all who use search engines or post images online. AI began as a major public project of the twentieth century and was relentlessly privatized to produce enormous financial gains for the tiny minority at the top of the extraction pyramid.

These diagrams present two different ways of understanding how AI works. I've argued that there is much at stake in how we define AI, what its boundaries are, and who determines them: it shapes what can be seen and contested. The Go diagram speaks to the industry narratives of an abstract computational cloud, far removed from the earthly resources needed to produce it, a paradigm where technical innovation is lionized, regulation is rejected, and true costs are never revealed. The blueprint points us to the physical infrastructure, but it leaves out the full environmental implications and the political deals that made it possible. These partial accounts of AI represent what philosophers Michael Hardt and Antonio Negri call the "dual operation of *abstraction* and *extraction*" in information capitalism: abstracting away the material conditions of production while extracting more information and resources.[14] The description of AI as fundamentally abstract distances it from the energy, labor, and capital needed to produce it and the many different kinds of mining that enable it.

This book has explored the planetary infrastructure of AI as an extractive industry: from its material genesis to the political economy of its operations to the discourses that support its aura of immateriality and inevitability. We have seen the politics inherent in how AI systems are trained to recognize the world. And we've observed the systemic forms of inequity that make AI what it is today. The core issue is the deep entanglement of technology, capital, and power, of which AI is the latest manifestation. Rather than being inscrutable and alien, these systems are products of larger social and economic structures with profound material consequences.

The Map Is Not the Territory

How do we see the full life cycle of artificial intelligence and the dynamics of power that drive it? We have to go beyond the conventional maps of AI to locate it in a wider landscape. Atlases can provoke a shift in scale, to see how spaces are joined in relation to one another. This book proposes that the real stakes of AI are the global interconnected systems of extraction and power, not the technocratic imaginaries of artificiality, abstraction, and automation. To understand AI for what it is, we need to see the structures of power it serves.

AI is born from salt lakes in Bolivia and mines in Congo, constructed from crowdworker-labeled datasets that seek to classify human actions, emotions, and identities. It is used to navigate drones over Yemen, direct immigration police in the United States, and modulate credit scores of human value and risk across the world. A wide-angle, multiscalar perspective on AI is needed to contend with these overlapping regimes.

This book began below the ground, where the extractive politics of artificial intelligence can be seen at their most literal. Rare earth minerals, water, coal, and oil: the tech sector carves out the earth to fuel its highly energy-intensive infrastructures. AI's carbon footprint is never fully admitted or accounted for by the tech sector, which is simultaneously expanding the networks of data centers while helping the oil and gas industry locate and strip remaining reserves of fossil fuels. The opacity of the larger supply chain for computation in general, and AI in particular, is part of a long-established business model of extracting value from the commons and avoiding restitution for the lasting damage.

Labor represents another form of extraction. In chapter 2, we ventured beyond the highly paid machine learning engineers to consider the other forms of work needed to make artificial intelligence systems function. From the miners extracting tin in Indonesia to crowdworkers in India completing tasks on Amazon Mechanical Turk to iPhone factory workers at Foxconn in China, the labor force of AI is far greater than we normally imagine. Even within the tech companies there is a large shadow workforce of contract laborers, who significantly outnumber full-time employees but have fewer benefits and no job security.[15]

In the logistical nodes of the tech sector, we find humans completing the tasks that machines cannot. Thousands of people are needed to support the illusion of automation: tagging, correcting, evaluating, and editing AI systems to make them appear seamless. Others lift packages, drive for ride-hailing apps, and deliver food. AI systems surveil them all while squeezing the most output from the bare functionality of human bodies: the complex joints of fingers, eyes, and knee sockets are cheaper and easier to acquire than robots. In those spaces, the future of work looks more like the Taylorist factories of the past, but with wristbands that vibrate when workers make errors and penalties given for taking too many bathroom breaks.

The uses of workplace AI further skew power imbalances by placing more control in employers' hands. Apps are used to track workers, nudge them to work longer hours, and rank them in real time. Amazon provides a canonical example of how a microphysics of power—disciplining bodies and their movement through space—is connected to a macrophysics of power, a logistics of planetary time and information. AI systems exploit differences in time and wages across markets to speed the circuits of capital. Suddenly, everyone in urban centers can have—and expects—same day delivery. And the system speeds up again, with the material consequences hidden behind the cardboard boxes, delivery trucks, and "buy now" buttons.

At the data layer, we can see a different geography of extraction. "We are building a mirror of the real world," a Google Street View engineer said in 2012. "Anything that you see in the real world needs to be in our databases."[16] Since then, the harvesting of the real world has only intensified to reach into spaces that were previously hard to capture. As we saw in chapter 3, there has been a widespread pillaging of public spaces; the faces of people in the street have been captured to train facial recognition systems; social media feeds have been ingested to build predictive models of language; sites where people keep personal photos or have online debates have been scraped in order to train machine vision and natural language algorithms. This practice has become so common that few in the AI field even question it. In part, that is because so many careers and market valuations depend on it. The collect-it-all mentality, once the remit of intelligence agencies, is not only normalized but moralized—it is seen as wasteful not to collect data wherever possible.[17]

Once data is extracted and ordered into training sets, it becomes the epistemic foundation by which AI systems classify the world. From the benchmark training sets such as ImageNet, MS-Celeb, or NIST's collections, images are used to represent ideas that are far more relational and contested than the labels may suggest. In chapter 4, we saw how labeling taxonomies allocate people into forced gender binaries, simplistic and offensive racial groupings, and highly normative and stereotypical analyses of character, merit, and emotional state. These classifications, unavoidably value-laden, force a way of seeing onto the world while claiming scientific neutrality.

Datasets in AI are never raw materials to feed algorithms: they are inherently political interventions. The entire practice of harvesting data, categorizing and labeling it, and then using it to train systems is a form of politics. It has brought a shift to what are called operational images—representations of the world made solely for machines.[18] Bias is a symptom of a deeper affliction: a far-ranging and centralizing normative logic that is used to determine how the world should be seen and evaluated.

A central example of this is affect detection, described in chapter 5, which draws on controversial ideas about the relation of faces to emotions and applies them with the reductive logic of a lie detector test. The science remains deeply contested.[19] Institutions have always classified people into identity categories, narrowing personhood and cutting it down into precisely measured boxes. Machine learning allows that to happen at scale. From the hill towns of Papua New Guinea to military labs in Maryland, techniques have been developed to reduce the messiness of feelings, interior states, preferences, and identifications into something quantitative, detectable, and trackable.

What epistemological violence is necessary to make the world readable to a machine learning system? AI seeks to systematize the unsystematizable, formalize the social, and convert an infinitely complex and changing universe into a Linnaean order of machine-readable tables. Many of AI's achievements have depended on boiling things down to a terse set of formalisms based on proxies: identifying and naming some features while ignoring or obscuring countless others. To adapt a phrase from philosopher Babette Babich, machine learning exploits what it does know to predict what it does not know: a game of repeated approximations. Datasets are also *proxies*—stand-ins for what they claim to measure. Put simply, this is transmuting difference into computable sameness. This kind of knowledge schema recalls what Friedrich Nietzsche described as "the falsifying of the multifarious and incalculable into the identical, similar, and calculable."[20] AI systems become deterministic when these proxies are taken as ground truth, when fixed labels are applied to a fluid complexity. We saw this in the cases where AI is used to predict gender, race, or sexuality from a photograph of a face.[21] These approaches resemble phrenology and physiognomy in their desire to essentialize and impose identities based on external appearances.

The problem of ground truth for AI systems is heightened in the context of state power, as we saw in chapter 6. The intelligence agencies led the way on the mass collection of data, where metadata signatures are sufficient for lethal drone strikes and a cell phone location becomes a proxy for an unknown target. Even here, the bloodless language of metadata and surgical strikes is directly contradicted by the unintended killings from drone missiles.[22] As Lucy Suchman has asked, how are "objects" identified as imminent threats? We know that "ISIS pickup truck" is a category based on hand-labeled data, but who chose the categories and identified the vehicles?[23] We saw the epistemological confusions and errors of object recognition training sets like ImageNet; military AI systems and drone attacks are built on the same unstable terrain.

The deep interconnections between the tech sector and the military are now framed within a strong nationalist agenda. The rhetoric about the AI war between the United States and China drives the interests of the largest tech companies to operate with greater government support and few restrictions. Meanwhile, the surveillance armory used by agencies like the NSA and the CIA is now deployed domestically at a municipal level in the in-between space of commercial-military contract-ing by companies like Palantir. Undocumented immigrants are hunted down with logistical systems of total information control and capture that were once reserved for extralegal espionage. Welfare decision-making systems are used to track anomalous data patterns in order to cut people off from unemployment benefits and accuse them of fraud. License plate reader technology is being used by home surveillance systems—a widespread integration of previously separate surveillance networks.[24]

The result is a profound and rapid expansion of surveillance and a blurring between private contractors, law enforcement, and the tech sector, fueled by kickbacks and secret deals. It is a radical redrawing of civic life, where the centers of power are strengthened by tools that see with the logics of capital, policing, and militarization.

Toward Connected Movements for Justice

If AI currently serves the existing structures of power, an obvious question might be: Should we not seek to democratize it? Could there not be an AI for the people that is reoriented toward justice and equality rather than industrial extraction and discrimination? This may seem appealing, but as we have seen throughout this book, the infrastructures and forms of power that enable and are enabled by AI skew strongly toward the centralization of control. To suggest that we democratize AI to reduce asymmetries of power is a little like arguing for democratizing weapons manufacturing in the service of peace. As Audre Lorde reminds us, the master's tools will never dismantle the master's house.[25]

A reckoning is due for the technology sector. To date, one common industry response has been to sign AI ethics principles. As European Union parliamentarian Marietje Schaake observed, in 2019 there were 128 frameworks for AI ethics in Europe alone.[26] These documents are often presented as products of a "wider consensus" on AI ethics. But they are overwhelmingly produced by economically developed countries, with little representation from Africa, South and Central America, or Central Asia. The voices of the people most harmed by AI systems are largely missing from the processes that produce them.[27] Further, ethical principles and statements don't discuss how they should be implemented, and they are rarely enforceable or accountable to a broader public. As Shannon Mattern has noted, the focus is more commonly on the ethical ends for AI, without assessing the ethical means of its application.[28] Unlike medicine or law, AI has no formal professional governance structure or norms—no agreed-upon definitions and goals for the field or standard protocols for enforcing ethical practice.[29]

Self-regulating ethical frameworks allow companies to choose how to deploy technologies and, by extension, to decide what ethical AI means for the rest of the world.[30] Tech companies rarely suffer serious financial penalties when their AI systems violate the law and even fewer consequences when their ethical principles are violated. Further, public companies are pressured by shareholders to maximize return on investment over ethical concerns, commonly making ethics secondary to profits. As a result, ethics is necessary but not sufficient to address the fundamental concerns raised in this book.

To understand what is at stake, we must focus less on ethics and more on power. AI is invariably designed to amplify and reproduce the forms of power it has been deployed to optimize. Countering that requires centering the interests of the communities most affected.[31] Instead of glorifying company founders, venture capitalists, and technical visionaries, we should begin with the lived experiences of those who are disempowered, discriminated against, and harmed by AI systems. When someone says, "AI ethics," we should assess the labor conditions for miners, contractors, and crowdworkers. When we hear "optimization," we should ask if these are tools for the inhumane treatment of immigrants. When there is applause for "large-scale automation," we should remember the resulting carbon footprint at a time when the planet is already under extreme stress. What would it mean to work toward justice across all these systems?

In 1986, the political theorist Langdon Winner described a society "committed to making artificial realities" with no concern for the harms it could bring to the conditions of life: "Vast transformations in the structure of our common world have been undertaken with little attention to what those alterations mean. . . . In the technical realm we repeatedly enter into a series of social contracts, the terms of which are only revealed after signing."[32]

In the four decades since, those transformations are now at a scale that has shifted the chemical composition of the atmosphere, the temperature of Earth's surface, and the contents of the planet's crust. The gap between how technology is judged on its release and its lasting consequences has only widened. The social contract, to the extent that there ever was one, has brought a climate crisis, soaring wealth inequality, racial discrimination, and widespread surveillance and labor exploitation. But the idea that these transformations occurred in ignorance of their possible results is part of the problem. The philosopher Achille Mbembé sharply critiques the idea that we could not have foreseen what would become of the knowledge systems of the twenty-first century, as they were always "operations of abstraction that claim to rationalize the world on the basis of corporate logic."[33] He writes: "It is about extraction, capture, the cult of data, the commodification of human capacity for thought and the dismissal of critical reason in favour of programming. . . . Now more than ever before, what we need is a new critique of technology, of the experience of technical life."[34]

The next era of critique will also need to find spaces beyond technical life by overturning the dogma of inevitability. When AI's rapid expansion is seen as unstoppable, it is possible only to patch together legal and technical restraints on systems after the fact: to clean up datasets, strengthen privacy laws, or create ethics boards. But these will always be partial and incomplete responses in which technology is assumed and everything else must adapt. But what happens if we reverse this polarity and begin with the commitment to a more just and sustainable world? How can we intervene to address interdependent issues of social, economic, and climate injustice? Where does technology serve that vision? And are there places where AI should not be used, where it undermines justice?

This is the basis for a renewed politics of refusal—opposing the narratives of technological inevitability that says, "If it can be done, it will be." Rather than asking where AI will be applied, merely because it can, the emphasis should be on *why* it ought to be applied. By asking, "Why use artificial intelligence?" we can question the idea that everything should be subject to the logics of statistical prediction and profit accumulation, what Donna Haraway terms the "informatics of domination."[35] We see glimpses of this refusal when populations choose to dismantle predictive policing, ban facial recognition, or protest algorithmic grading. So far these minor victories have been piecemeal and localized, often centered in cities with more resources to organize, such as London, San Francisco, Hong Kong, and Portland, Oregon. But they point to the need for broader national and international movements that refuse technology-first approaches and focus on addressing underlying inequities and injustices. Refusal requires rejecting the idea that the same tools that serve capital, militaries, and police are also fit to transform schools, hospitals, cities, and ecologies, as though they were value neutral calculators that can be applied everywhere.

The calls for labor, climate, and data justice are at their most powerful when they are united. Above all, I see the greatest hope in the growing justice movements that address the interrelatedness of capitalism, computation, and control: bringing together issues of climate justice, labor rights, racial justice, data protection, and the overreach of police and military power. By rejecting systems that further inequity and violence, we challenge the structures of power that AI currently reinforces and create the foundations for a different society.[36] As Ruha Benjamin notes, “Derrick Bell said it like this: ‘To see things as they really are, you must imagine them for what they might be.’ We are pattern makers and we must change the content of our existing patterns.”[37] To do so will require shaking off the enchantments of tech solutionism and embracing alternative solidarities—what Mbembé calls “a different politics of inhabiting the Earth, of repairing and sharing the planet.”[38] There are sustainable collective politics beyond value extraction; there are commons worth keeping, worlds beyond the market, and ways to live beyond discrimination and brutal modes of optimization. Our task is to chart a course there.

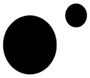

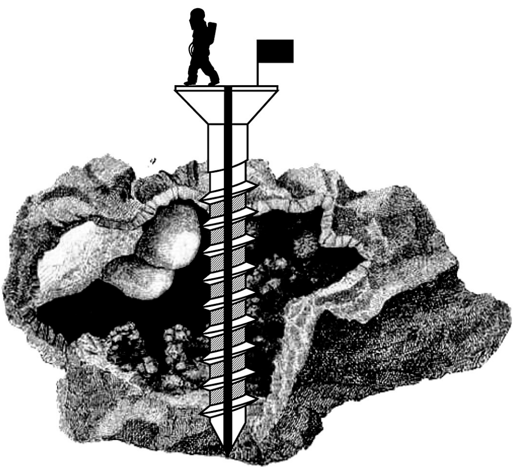

Coda
Space

A countdown begins. File footage starts rolling. Engines at the base of a towering Saturn V ignite, and the rocket begins liftoff. We hear the voice of Jeff Bezos: "Ever since I was five years old—that's when Neil Armstrong stepped onto the surface of the moon—I've been passionate about space, rockets, rocket engines, space travel." A parade of inspirational images appears: mountain climbers at summits, explorers descending into canyons, an ocean diver swimming through a shoal of fish.

Cut to Bezos in a control room during a launch, adjusting his headset. His voiceover continues: "This is the most important work I'm doing. It's a simple argument, this is the best planet. And so we face a choice. As we move forward, we're gonna have to decide whether we want a civilization of stasis—we will have to cap population, we will have to cap energy usage per capita—or we can fix that problem, by moving out into space."[1]

The soundtrack soars, and images of deep space are counterposed with shots of the busy freeways of Los Angeles and clogged cloverleaf junctions. "Von Braun said, after the lunar landing, 'I have learned to use the word impossible with great caution.' And I hope you guys take that attitude about your lives."[2]

This scene comes from a promotional video for Bezos's private aerospace company, Blue Origin. The company motto is *Gradatim Ferociter,* Latin for "Step by Step, Ferociously." In the near term, Blue Origin is building reusable rockets and lunar landers, testing them primarily at its facility and suborbital base in West Texas. By 2024, the company wants to be shuttling astronauts and cargo to the Moon.[3] But in the longer term, the company's mission is far more ambitious: to help bring about a future in which millions are living and working in space. Specifically, Bezos has outlined his hopes to build giant space colonies, where people would live in floating manufactured environments.[4] Heavy industry would move off-planet altogether, the new frontier for extraction. Meanwhile, Earth would be zoned for residential building and light industry, left as a "beautiful place to live, a beautiful place to visit"—presumably for those who can afford to be there, rather than working in the off-world colonies.[5]

Bezos possesses extraordinary and growing industrial power. Amazon continues to capture more of U.S. online commerce, Amazon Web Services represents nearly half of the cloud-computing industry, and, by some estimates, Amazon's site has more product searches than Google.[6] Despite all this, Bezos is worried. His fear is that the planet's growing energy demands will soon outstrip its limited supply. For him, the greatest concern "is not necessarily extinction" but *stasis:* "We will have to stop growing, which I think is a very bad future."[7]

Bezos is not alone. He is just one of several tech billionaires focused on space. Planetary Resources, led by the founder of the X Prize, Peter Diamandis, and backed with investment from Google's Larry Page and Eric Schmidt, aimed to create the first commercial mine in space by drilling asteroids.[8] Elon Musk, chief executive of Tesla and SpaceX, has announced his intention to colonize Mars within a hundred years—while admitting that, to do so, the first astronauts must "be prepared to die."[9] Musk has also advocated terraforming the surface of Mars for human settlement by exploding nuclear weapons at the poles.[10] SpaceX made a T-shirt that reads "NUKE MARS." Musk also conducted what is arguably the most expensive public relations exercise in history when he launched a Tesla car into heliocentric orbit on a SpaceX Falcon Heavy rocket. Researchers estimate that the car will remain in space for millions of years, until it finally crashes back to Earth.[11]

The ideology of these space spectacles is deeply interconnected with that of the AI industry. Extreme wealth and power generated from technology companies now enables a small group of men to pursue their own private space race. They depend on exploiting the knowledge and infrastructures of the public space programs of the twentieth century and often rely on government funding and tax incentives as well.[12] Their aim is not to limit extraction and growth but to extend it across the solar system. In actuality, these efforts are as much about an *imaginary* of space, endless growth, and immortality than they are about the uncertain and unpleasant possibilities of actual space colonization.

Bezos's inspiration for conquering space comes, in part, from the physicist and science fiction novelist Gerard K. O'Neill. O'Neill wrote *The High Frontier: Human Colonies in Space,* a 1976 fantasy of space colonization, which includes lush illustrations of moon mining with Rockwellian abundance.[13] Bezos's plan for Blue Origin is inspired by this bucolic vision of permanent human settlement, for which no current technology exists.[14] O'Neill was driven by the "dismay and shock" he felt when he read the 1972 landmark report by the Club of Rome, called *The Limits to Growth*.[15] The report published extensive data and predictive models about the end of nonrenewable resources and the impact on population growth, sustainability, and humanity's future on Earth.[16] As the architecture and planning scholar Fred Scharmen summarizes:

> The Club of Rome models calculate outcomes from different sets of initial assumptions. The baseline scenarios, extrapolated from then-current trends, show resource and population collapse before the year 2100. When the models assume double the known resource reserves, they collapse again, to a slightly higher level but still before 2100. When they assume that technology will make available "unlimited" resources, population collapses *even more sharply* than before due to spikes in pollution. With pollution controls added to the model, population collapses after running out of food. In models that increase agricultural capacity, pollution overruns previous controls and both food and population collapse.[17]

Limits to Growth suggested that moving to sustainable management and reuse of resources was the answer to long-term stability of global society and that narrowing the gap between rich and poor nations was the key to survival. Where *Limits to Growth* fell short was that it did not foresee the larger set of interconnected systems that now make up the global economy and how previously uneconomic forms of mining would be incentivized, driving greater environmental harms, land and water degradation, and accelerated resource depletion.

In writing *The High Frontier,* O'Neill wanted to imagine a different way out of the no-growth model rather than limiting production and consumption.[18] By positing that space was a solution, O'Neill redirected global anxiety in the 1970s over gasoline shortages and oil crises with visions of serene stable space structures that would simultaneously preserve the status quo and offer new opportunities. "If Earth doesn't have enough surface area," O'Neill urged, "then humans should simply build more."[19] The science of how it would work and the economics of how we could afford it were details left for another day; the dream was all that mattered.[20]

That space colonization and frontier mining have become the common corporate fantasies of tech billionaires underscores a fundamentally troubling relationship to Earth. Their vision of the future does not include minimizing oil and gas exploration or containing resource consumption or even reducing the exploitative labor practices that have enriched them. Instead, the language of the tech elite often echoes settler colonialism, seeking to displace Earth's population and capture territory for mineral extraction. Silicon Valley's billionaire space race similarly assumes that the last commons—outer space—can be taken by whichever empire gets there first. This is despite the main convention governing space mining, the 1967 Outer Space Treaty, which recognizes that space is the "common interest of all mankind" and that any exploration or use "should be carried on for the benefit of all peoples."[21]

In 2015, Bezos's Blue Origin and Musk's SpaceX lobbied Congress and the Obama administration to enact the Commercial Space Launch Competitiveness Act.[22] It extends an exemption for commercial space companies from federal regulation until 2023, allowing them to own any mining resources extracted from asteroids and keep the profits.[23] This legislation directly undercuts the idea of space as a commons, and creates a commercial incentive to "go forth and conquer."[24]

Space has become the ultimate imperial ambition, symbolizing an escape from the limits of Earth, bodies, and regulation. It is perhaps no surprise that many of the Silicon Valley tech elite are invested in the vision of abandoning the planet. Space colonization fits well alongside the other fantasies of life-extension dieting, blood transfusions from teenagers, brain-uploading to the cloud, and vitamins for immortality.[25] Blue Origin's high-gloss advertising is part of this dark utopianism. It is a whispered summons to become the Übermensch, to exceed all boundaries: biological, social, ethical, and ecological. But underneath, these visions of brave new worlds seem driven most of all by fear: fear of death—individually and collectively—and fear that time is truly running out.

I'm back in the van for the last leg of my journey. I drive south out of Albuquerque, New Mexico, headed toward the Texas border. On my way, I take a detour past the rocky face of San Augustin Peak and follow the steep drive down to the White Sands Missile Range, where in 1946 the United States launched the first rocket containing a camera into space. That mission was led by Wernher von Braun, who had been the technical director of Germany's missile rocket development program. He defected to the United States after the war, and there he began experimenting with confiscated V-2 rockets—the very missiles he had helped design, which had been fired against the Allies across Europe. But this time he sent them directly upward, into space. The rocket ascended to an altitude of 65 miles, capturing images every 1.5 seconds, before crashing into the New Mexican desert. The film survived inside of a steel cassette, revealing a grainy but distinctly Earthlike curve.[26]

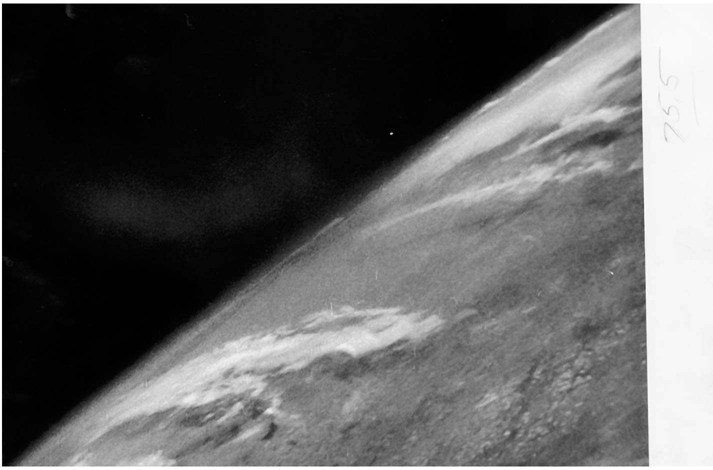

View of Earth from a camera on V-2 #13,
launched October 24, 1946. Courtesy White Sands
Missile Range/Applied Physics Laboratory

That Bezos chose to quote von Braun in his Blue Origin commercial is notable. Von Braun was chief rocket engineer of the Third Reich and admitted using concentration camp slave labor to build his V-2 rockets; some consider him a war criminal.[27] More people died in the camps building the rockets than were killed by them in war.[28] But it is von Braun's work as head of NASA's Marshall Space Flight Center, where he was instrumental in the design of the Saturn V rocket, that is best known.[29] Bathed in the glow of Apollo 11, washed clean of history, Bezos has found his hero—a man who refused to believe in impossibility.

After driving through El Paso, Texas, I take Route 62 toward the Salt Basin Dunes. It's late in the afternoon, and colors are starting to bloom in the cumulus clouds. There's a T-junction, and after turning right, the road begins to trace along the Sierra Diablos. This is Bezos country. The first in-

Blue Origin suborbital launch facility, West Texas.
Photograph by Kate Crawford

dication is a large ranch house set back from the road, with a sign in red letters that reads "Figure 2" on a white gate. It's the ranch that Bezos purchased in 2004, just part of the three hundred thousand acres he owns in the area.[30] The land has a violent colonial history: one of the final battles between the Texas Rangers and the Apaches occurred just west of this site in 1881, and nine years later the ranch was created by the one-time Confederate rider and cattleman James Monroe Daugherty.[31]

Nearby is the turn-off to the Blue Origin suborbital launch facility. The private road is blocked by a bright blue gate with security notices warning of video surveillance and a guard station bristling with cameras. I stay on the highway and pull the van over to the side of the road a few minutes away. From here, the views stretch across the valley to the Blue Origin landing site, where the rockets are being tested for what is expected to be the company's first human mission into space. Cars pass through the boom gates as the workers clock out for the day.

Looking back at the clusters of sheds that mark out the rocket base, it feels very provisional and makeshift in this dry expanse of the Permian Basin. The vast span of the valley is broken with a hollow circle, the landing pad where Blue Origin's reusable rockets are meant to touch down on a feather logo painted in the center. That's all there is to see. It's a private infrastructure-in-progress, guarded and gated, a technoscientific imaginary of power, extraction, and escape, driven by the wealthiest man on the planet. It is a hedge against Earth.

The light is fading now, and steel-gray clouds are moving against the sky. The desert looks silvery, dotted with white sage bushes and clusters of volcanic tuff punctuating what was once the floor of a great inland sea. After taking a photograph, I head back to the van to begin the final drive of the day to the town of Marfa. It's not until I start driving away that I realize I'm being followed. Two matching black Chevrolet pickups begin aggressively tailgating at close range. I pull over in the hope they will pass. They also pull over. No one moves. After waiting a few minutes, I slowly begin to drive again. They maintain their sinister escort all the way to the edge of the darkening valley.

Notes

Introduction

1. Heyn, "Berlin's Wonderful Horse."

2. Pfungst, *Clever Hans.*

3. "'Clever Hans' Again."

4. Pfungst, *Clever Hans.*

5. Pfungst.

6. Lapuschkin et al., "Unmasking Clever Hans Predictors."

7. See the work of philosopher Val Plumwood on the dualisms of intelligence-stupid, emotional-rational, and master-slave. Plumwood, "Politics of Reason."

8. Turing, "Computing Machinery and Intelligence."

9. Von Neumann, *The Computer and the Brain,* 44. This approach was deeply critiqued by Dreyfus, *What Computers Can't Do.*

10. See Weizenbaum, "On the Impact of the Computer on Society," 612. After his death, Minsky was implicated in serious allegations related to convicted pedophile and rapist Jeffrey Epstein. Minsky was one of several scientists who met with Epstein and visited his island retreat where underage girls were forced to have sex with members of Epstein's coterie. As scholar Meredith Broussard observes, this was part of a broader culture of exclusion that became endemic in AI: "As wonderfully creative as Minsky and his cohort were, they also solidified the culture of tech as a billionaire boys' club. Math, physics, and the other 'hard' sciences have never been hospitable to women and people of color; tech followed this lead." See Broussard, *Artificial Unintelligence,* 174.

11. Weizenbaum, *Computer Power and Human Reason,* 202–3.

12. Greenberger, *Management and the Computer of the Future,* 315.

13. Dreyfus, *Alchemy and Artificial Intelligence.*

14. Dreyfus, *What Computers Can't Do.*

15. Ullman, *Life in Code,* 136–37.

16. See, as one of many examples, Poggio et al., "Why and When Can Deep—but Not Shallow—Networks Avoid the Curse of Dimensionality."

17. Quoted in Gill, *Artificial Intelligence for Society,* 3.

18. Russell and Norvig, *Artificial Intelligence,* 30.

19. Daston, "Cloud Physiognomy."

20. Didi-Huberman, *Atlas,* 5.

21. Didi-Huberman, 11.

22. Franklin and Swenarchuk, *Ursula Franklin Reader,* Prelude.

23. For an account of the practices of data colonization, see "Colonized by Data"; and Mbembé, *Critique of Black Reason.*

24. Fei-Fei Li quoted in Gershgorn, "Data That Transformed AI Research."

25. Russell and Norvig, *Artificial Intelligence,* 1.

26. Bledsoe quoted in McCorduck, *Machines Who Think,* 136.

27. Mattern, *Code and Clay, Data and Dirt,* xxxiv–xxxv.

28. Ananny and Crawford, "Seeing without Knowing."

29. Any list will always be an inadequate account of all the people and communities who have inspired and informed this work. I'm particularly grateful to these research communities: FATE (Fairness, Accountability, Transparency and Ethics) and the Social Media Collective at Microsoft Research, the AI Now Institute at NYU, the Foundations of AI working group at the École Normale Supérieure, and the Richard von Weizsäcker Visiting Fellows at the Robert Bosch Academy in Berlin.

30. Saville, "Towards Humble Geographies."

31. For more on crowdworkers, see Gray and Suri, *Ghost Work;* and Roberts, *Behind the Screen.*

32. Canales, *Tenth of a Second.*

33. Zuboff, *Age of Surveillance Capitalism.*

34. Cetina, *Epistemic Cultures,* 3.

35. "Emotion Detection and Recognition (EDR) Market Size."

36. Nelson, Tu, and Hines, "Introduction," 5.

37. Danowski and de Castro, *Ends of the World.*

38. Franklin, *Real World of Technology,* 5.

1
Earth

1. Brechin, *Imperial San Francisco.*

2. Brechin, 29.

3. Agricola quoted in Brechin, 25.

4. Quoted in Brechin, 50.

5. Brechin, 69.

6. See, e.g., Davies and Young, *Tales from the Dark Side of the City;* and "Grey Goldmine."

7. For more on the street-level changes in San Francisco, see Bloomfield, "History of the California Historical Society's New Mission Street Neighborhood."

8. "Street Homelessness." See also "Counterpoints: An Atlas of Displacement and Resistance."

9. Gee, "San Francisco or Mumbai?"

10. H. W. Turner published a detailed geological survey of the Silver Peak area in July 1909. In beautiful prose, Turner extolled the geological variety within what he described as "slopes of cream and pink tuffs, and little hillocks of a bright brick red." Turner, "Contribution to the Geology of the Silver Peak Quadrangle, Nevada," 228.

11. Lambert, "Breakdown of Raw Materials in Tesla's Batteries and Possible Breaknecks."

12. Bullis, "Lithium-Ion Battery."

13. "Chinese Lithium Giant Agrees to Three-Year Pact to Supply Tesla."

14. Wald, "Tesla Is a Battery Business."

15. Scheyder, "Tesla Expects Global Shortage."

16. Wade, "Tesla's Electric Cars Aren't as Green."

17. Business Council for Sustainable Energy, "2019 Sustainable Energy in America Factbook." U.S. Energy Information Administration, "What Is U.S. Electricity Generation by Energy Source?"

18. Whittaker et al., *AI Now Report 2018.*

19. Parikka, *Geology of Media,* vii—viii; McLuhan, *Understanding Media.*

20. Ely, "Life Expectancy of Electronics."

21. Sandro Mezzadra and Brett Neilson use the term "extractivism" to name the relation between different forms of extractive operations in contemporary capitalism, which we see repeated in the context of the AI industry. Mezzadra and Neilson, "Multiple Frontiers of Extraction."

22. Nassar et al., "Evaluating the Mineral Commodity Supply Risk of the US Manufacturing Sector."

23. Mumford, *Technics and Civilization,* 74.

24. See, e.g., Ayogu and Lewis, "Conflict Minerals."

25. Burke, "Congo Violence Fuels Fears of Return to 90s Bloodbath."

26. "Congo's Bloody Coltan."

27. "Congo's Bloody Coltan."

28. "Transforming Intel's Supply Chain with Real-Time Analytics."

29. See, e.g., an open letter from seventy signatories that criticizes the limitations of the so-called conflict-free certification process: "An Open Letter."

30. "Responsible Minerals Policy and Due Diligence."

31. In *The Elements of Power,* David S. Abraham describes the invisible networks of rare metals traders in global electronics supply chains: "The network to get rare metals from the mine to your laptop travels through a murky network of traders, processors, and component manufacturers. Traders are the middlemen who do more than buy and sell rare metals: they help to regulate information and are the hidden link that helps in navigating the network between metals plants and the components in our laptops" (89).

32. "Responsible Minerals Sourcing."

33. Liu, "Chinese Mining Dump."

34. "Bayan Obo Deposit."

35. Maughan, "Dystopian Lake Filled by the World's Tech Lust."

36. Hird, "Waste, Landfills, and an Environmental Ethics of Vulnerability," 105.

37. Abraham, *Elements of Power,* 175.

38. Abraham, 176.

39. Simpson, "Deadly Tin Inside Your Smartphone."

40. Hodal, "Death Metal."

41. Hodal.

42. Tully, "Victorian Ecological Disaster."

43. Starosielski, *Undersea Network,* 34.

44. See Couldry and Mejías, *Costs of Connection,* 46.

45. Couldry and Mejías, 574.

46. For a superb account of the history of undersea cables, see Starosielski, *Undersea Network.*

47. Dryer, "Designing Certainty," 45.

48. Dryer, 46.

49. Dryer, 266–68.

50. More people are now drawing attention to this problem—including researchers at AI Now. See Dobbe and Whittaker, "AI and Climate Change."

51. See, as an example of early scholarship in this area, Ensmenger, "Computation, Materiality, and the Global Environment."

52. Hu, *Prehistory of the Cloud,* 146.

53. Jones, "How to Stop Data Centres from Gobbling Up the World's Electricity." Some progress has been made toward mitigating these concerns through greater energy efficiency practices, but significant long-term challenges remain. Masanet et al., "Recalibrating Global Data Center Energy-Use Estimates."

54. Belkhir and Elmeligi, "Assessing ICT Global Emissions Footprint"; Andrae and Edler, "On Global Electricity Usage."

55. Strubell, Ganesh, and McCallum, "Energy and Policy Considerations for Deep Learning in NLP."

56. Strubell, Ganesh, and McCallum.

57. Sutton, "Bitter Lesson."

58. "AI and Compute."

59. Cook et al., *Clicking Clean.*

60. Ghaffary, "More Than 1,000 Google Employees Signed a Letter." See also "Apple Commits to Be 100 Percent Carbon Neutral"; Harrabin, "Google Says Its Carbon Footprint Is Now Zero"; Smith, "Microsoft Will Be Carbon Negative by 2030."

61. "Powering the Cloud."

62. "Powering the Cloud."

63. "Powering the Cloud."

64. Hogan, "Data Flows and Water Woes."

65. "Off Now."

66. Carlisle, "Shutting Off NSA's Water Gains Support."

67. Materiality is a complex concept, and there is a lengthy literature that contends with it in such fields as STS, anthropology, and media studies. In one sense, materiality refers to what Leah Lievrouw describes as "the physical character and existence of objects and artifacts that makes them useful and usable for certain purposes under particular conditions." Lievrouw quoted in Gillespie, Boczkowski, and Foot, *Media Technologies,* 25. But as Diana Coole and Samantha Frost write, "Materiality is always something more than 'mere' matter: an excess, force, vitality, relationality, or difference that renders matter active, self-creative, productive, unproductive." Coole and Frost, *New Materialisms,* 9.

68. United Nations Conference on Trade and Development, *Review of Maritime Transport, 2017.*

69. George, *Ninety Percent of Everything,* 4.

70. Schlanger, "If Shipping Were a Country."

71. Vidal, "Health Risks of Shipping Pollution."

72. "Containers Lost at Sea—2017 Update."

73. Adams, "Lost at Sea."

74. Mumford, *Myth of the Machine.*

75. Labban, "Deterritorializing Extraction." For an expansion on this idea, see Arboleda, *Planetary Mine.*

76. Ananny and Crawford, "Seeing without Knowing."

2
Labor

1. Wilson, "Amazon and Target Race."

2. Lingel and Crawford, "Alexa, Tell Me about Your Mother."

3. Federici, *Wages against Housework;* Gregg, *Counterproductive.*

4. In *The Utopia of Rules,* David Graeber details the sense of loss experienced by white-collar workers who now have to enter data into the decision-making systems that have replaced specialist administrative support staff in most professional workplaces.

5. Smith, *Wealth of Nations,* 4–5.

6. Marx and Engels, *Marx-Engels Reader,* 479. Marx expanded on this notion of the worker as an "appendage" in *Capital,* vol. 1: "In handicrafts and manufacture, the worker makes use of a tool; in the factory, the machine makes use of him. There the movements of the instrument of labor proceed from him, here it is the movements of the machine that he must follow. In manufacture the workers are parts of a living mechanism. In the factory we have a lifeless mechanism which is independent of the workers, who are incorporated into it as its living appendages." Marx, *Das Kapital,* 548–49.

7. Luxemburg, "Practical Economies," 444.

8. Thompson, "Time, Work-Discipline, and Industrial Capitalism."

9. Thompson, 88–90.

10. Werrett, "Potemkin and the Panopticon," 6.

11. See, e.g., Cooper, "Portsmouth System of Manufacture."

12. Foucault, *Discipline and Punish;* Horne and Maly, *Inspection House.*

13. Mirzoeff, *Right to Look,* 58.

14. Mirzoeff, 55.

15. Mirzoeff, 56.

16. Gray and Suri, *Ghost Work.*

17. Irani, "Hidden Faces of Automation."

18. Yuan, "How Cheap Labor Drives China's A.I. Ambitions"; Gray and Suri, "Humans Working behind the AI Curtain."

19. Berg et al., *Digital Labour Platforms.*

20. Roberts, *Behind the Screen;* Gillespie, *Custodians of the Internet,* 111–40.

21. Silberman et al., "Responsible Research with Crowds."

22. Silberman et al.

23. Huet, "Humans Hiding behind the Chatbots."

24. Huet.

25. See Sadowski, "Potemkin AI."

26. Taylor, "Automation Charade."

27. Taylor.

28. Gray and Suri, *Ghost Work.*

29. Standage, *Turk,* 23.

30. Standage, 23.

31. See, e.g., Aytes, "Return of the Crowds," 80.

32. Irani, "Difference and Dependence among Digital Workers," 225.

33. Pontin, "Artificial Intelligence."

34. Menabrea and Lovelace, "Sketch of the Analytical Engine."

35. Babbage, *On the Economy of Machinery and Manufactures,* 39–43.

36. Babbage evidently acquired an interest in quality-control processes while trying (vainly) to establish a reliable supply chain for the components of his calculating engines.

37. Schaffer, "Babbage's Calculating Engines and the Factory System," 280.

38. Taylor, *People's Platform,* 42.

39. Katz and Krueger, "Rise and Nature of Alternative Work Arrangements."

40. Rehmann, "Taylorism and Fordism in the Stockyards," 26.

41. Braverman, *Labor and Monopoly Capital,* 56, 67; Specht, *Red Meat Republic.*

42. Taylor, *Principles of Scientific Management.*

43. Marx, *Poverty of Philosophy,* 22.

44. Qiu, Gregg, and Crawford, "Circuits of Labour"; Qiu, *Goodbye iSlave.*

45. Markoff, "Skilled Work, without the Worker."

46. Guendelsberger, *On the Clock,* 22.

47. Greenhouse, "McDonald's Workers File Wage Suits."

48. Greenhouse.

49. Mayhew and Quinlan, "Fordism in the Fast Food Industry."

50. Ajunwa, Crawford, and Schultz, "Limitless Worker Surveillance."

51. Mikel, "WeWork Just Made a Disturbing Acquisition."

52. Mahdawi, "Domino's 'Pizza Checker' Is Just the Beginning."

53. Wajcman, "How Silicon Valley Sets Time."

54. Wajcman, 1277.

55. Gora, Herzog, and Tripathi, "Clock Synchronization."

56. Eglash, "Broken Metaphor," 361.

57. Kemeny and Kurtz, "Dartmouth Timesharing," 223.

58. Eglash, "Broken Metaphor," 364.

59. Brewer, "Spanner, TrueTime."

60. Corbett et al., "Spanner," 14, cited in House, "Synchronizing Uncertainty," 124.

61. Galison, *Einstein's Clocks, Poincaré's Maps,* 104.

62. Galison, 112.

63. Colligan and Linley, "Media, Technology, and Literature," 246.

64. Carey, "Technology and Ideology."

65. Carey, 13.

66. This contrasts with what Foucault called the "microphysics of power" to describe how institutions and apparatuses create particular logics and forms of validity. Foucault, *Discipline and Punish,* 26.

67. Spargo, *Syndicalism, Industrial Unionism, and Socialism.*

68. Personal conversation with the author at an Amazon fulfillment center tour, Robbinsville, N.J., October 8, 2019.

69. Muse, “Organizing Tech.”
70. Abdi Muse, personal conversation with the author, October 2, 2019.
71. Gurley, “60 Amazon Workers Walked Out.”
72. Muse quoted in *Organizing Tech.*
73. Desai quoted in *Organizing Tech.*
74. Estreicher and Owens, “Labor Board Wrongly Rejects Employee Access to Company Email.”
75. This observation comes from conversations with various labor organizers, tech workers, and researchers, including Astra Taylor, Dan Greene, Bo Daley, and Meredith Whittaker.
76. Kerr, “Tech Workers Protest in SF.”

3
Data

1. National Institute of Standards and Technology (NIST), “Special Database 32—Multiple Encounter Dataset (MEDS).”
2. Russell, *Open Standards and the Digital Age.*
3. Researchers at NIST (then the National Bureau of Standards, NBS) began working on the first version of the FBI’s Automated Fingerprint Identification System in the late 1960s. See Garris and Wilson, “NIST Biometrics Evaluations and Developments,” 1.
4. Garris and Wilson, 1.
5. Garris and Wilson, 12.
6. Sekula, “Body and the Archive,” 17.
7. Sekula, 18–19.
8. Sekula, 17.
9. See, e.g., Grother et al., “2017 IARPA Face Recognition Prize Challenge (FRPC).”
10. See, e.g., Ever AI, “Ever AI Leads All US Companies.”
11. Founds et al., “NIST Special Database 32.”
12. Curry et al., “NIST Special Database 32 Multiple Encounter Dataset I (MEDS-I),” 8.
13. See, e.g., Jaton, “We Get the Algorithms of Our Ground Truths.”
14. Nilsson, *Quest for Artificial Intelligence,* 398.
15. “ImageNet Large Scale Visual Recognition Competition (ILSVRC).”
16. In the late 1970s, Ryszard Michalski wrote an algorithm based on symbolic variables and logical rules. This language was popular in the 1980s and 1990s, but as the rules of decision-making and qualification became more complex, the language became less usable. At the same moment, the potential of using large training sets triggered a shift from this conceptual clus-

tering to contemporary machine learning approaches. Michalski, "Pattern Recognition as Rule-Guided Inductive Inference."

17. Bush, "As We May Think."

18. Light, "When Computers Were Women"; Hicks, *Programmed Inequality.*

19. As described in Russell and Norvig, *Artificial Intelligence,* 546.

20. Li, "Divination Engines," 143.

21. Li, 144.

22. Brown and Mercer, "Oh, Yes, Everything's Right on Schedule, Fred."

23. Lem, "First Sally (A), or Trurl's Electronic Bard," 199.

24. Lem, 199.

25. Brown and Mercer, "Oh, Yes, Everything's Right on Schedule, Fred."

26. Marcus, Marcinkiewicz, and Santorini, "Building a Large Annotated Corpus of English."

27. Klimt and Yang, "Enron Corpus."

28. Wood, Massey, and Brownell, "FERC Order Directing Release of Information," 12.

29. Heller, "What the Enron Emails Say about Us."

30. Baker et al., "Research Developments and Directions in Speech Recognition."

31. I have participated in early work to address this gap. See, e.g., Gebru et al., "Datasheets for Datasets." Other researchers have also sought to address this problem for AI models; see Mitchell et al., "Model Cards for Model Reporting"; Raji and Buolamwini, "Actionable Auditing."

32. Phillips, Rauss, and Der, "FERET (Face Recognition Technology) Recognition Algorithm Development and Test Results," 9.

33. Phillips, Rauss, and Der, 61.

34. Phillips, Rauss, and Der, 12.

35. See Aslam, "Facebook by the Numbers (2019)"; and "Advertising on Twitter."

36. Fei-Fei Li, as quoted in Gershgorn, "Data That Transformed AI Research."

37. Deng et al., "ImageNet."

38. Gershgorn, "Data That Transformed AI Research."

39. Gershgorn.

40. Markoff, "Seeking a Better Way to Find Web Images."

41. Hernandez, "CU Colorado Springs Students Secretly Photographed."

42. Zhang et al., "Multi-Target, Multi-Camera Tracking by Hierarchical Clustering."

43. Sheridan, "Duke Study Recorded Thousands of Students' Faces."

44. Harvey and LaPlace, "Brainwash Dataset."

45. Locker, "Microsoft, Duke, and Stanford Quietly Delete Databases."

46. Murgia and Harlow, "Who's Using Your Face?" When the *Financial*

Times exposed the contents of this dataset, Microsoft removed the set from the internet, and a spokesperson for Microsoft claimed simply that it was removed "because the research challenge is over." Locker, "Microsoft, Duke, and Stanford Quietly Delete Databases."

47. Franceschi-Bicchierai, "Redditor Cracks Anonymous Data Trove."

48. Tockar, "Riding with the Stars."

49. Crawford and Schultz, "Big Data and Due Process."

50. Franceschi-Bicchierai, "Redditor Cracks Anonymous Data Trove."

51. Nilsson, *Quest for Artificial Intelligence,* 495.

52. And, as Geoff Bowker famously reminds us, "Raw data is both an oxymoron and a bad idea; to the contrary, data should be cooked with care." Bowker, *Memory Practices in the Sciences,* 184–85.

53. Fourcade and Healy, "Seeing Like a Market," 13, emphasis added.

54. Meyer and Jepperson, "'Actors' of Modern Society."

55. Gitelman, *"Raw Data" Is an Oxymoron,* 3.

56. Many scholars have looked closely at the work these metaphors do. Media studies professors Cornelius Puschmann and Jean Burgess analyzed the common data metaphors and noted two widespread categories: data "as a natural force to be controlled and [data] as a resource to be consumed." Puschmann and Burgess, "Big Data, Big Questions," abstract. Researchers Tim Hwang and Karen Levy suggest that describing data as "the new oil" carries connotations of being costly to acquire but also suggests the possibility of "big payoffs for those with the means to extract it." Hwang and Levy, "'The Cloud' and Other Dangerous Metaphors."

57. Stark and Hoffmann, "Data Is the New What?"

58. Media scholars Nick Couldry and Ulises Mejías call this "data colonialism," which is steeped in the historical, predatory practices of colonialism but married to (and obscured by) contemporary computing methods. However, as other scholars have shown, this terminology is double-edged because it can occlude the real and ongoing harms of colonialism. Couldry and Mejías, "Data Colonialism"; Couldry and Mejías, *Costs of Connection;* Segura and Waisbord, "Between Data Capitalism and Data Citizenship."

59. They refer to this form of capital as "ubercapital." Fourcade and Healy, "Seeing Like a Market," 19.

60. Sadowski, "When Data Is Capital," 8.

61. Sadowski, 9.

62. Here I'm drawing from a history of human subjects review and large-scale data studies coauthored with Jake Metcalf. See Metcalf and Crawford, "Where Are Human Subjects in Big Data Research?"

63. "Federal Policy for the Protection of Human Subjects."

64. See Metcalf and Crawford, "Where Are Human Subjects in Big Data Research?"

65. Seo et al., "Partially Generative Neural Networks." Jeffrey Brantingham, one of the authors, is also a co-founder of the controversial predictive policing company PredPol. See Winston and Burrington, "A Pioneer in Predictive Policing."

66. "CalGang Criminal Intelligence System."

67. Libby, "Scathing Audit Bolsters Critics' Fears."

68. Hutson, "Artificial Intelligence Could Identify Gang Crimes."

69. Hoffmann, "Data Violence and How Bad Engineering Choices Can Damage Society."

70. Weizenbaum, *Computer Power and Human Reason,* 266.

71. Weizenbaum, 275–76.

72. Weizenbaum, 276.

73. For more on the history of extraction of data and insights from marginalized communities, see Costanza-Chock, *Design Justice;* and D'Ignazio and Klein, *Data Feminism.*

74. Revell, "Google DeepMind's NHS Data Deal 'Failed to Comply.'"

75. "Royal Free–Google DeepMind Trial Failed to Comply."

4
Classification

1. Fabian, *Skull Collectors.*

2. Gould, *Mismeasure of Man,* 83.

3. Kolbert, "There's No Scientific Basis for Race."

4. Keel, "Religion, Polygenism and the Early Science of Human Origins."

5. Thomas, *Skull Wars.*

6. Thomas, 85.

7. Kendi, "History of Race and Racism in America."

8. Gould, *Mismeasure of Man,* 88.

9. Mitchell, "Fault in His Seeds."

10. Horowitz, "Why Brain Size Doesn't Correlate with Intelligence."

11. Mitchell, "Fault in His Seeds."

12. Gould, *Mismeasure of Man,* 58.

13. West, "Genealogy of Modern Racism," 91.

14. Bouche and Rivard, "America's Hidden History."

15. Bowker and Star, *Sorting Things Out,* 319.

16. Bowker and Star, 319.

17. Nedlund, "Apple Card Is Accused of Gender Bias"; Angwin et al., "Machine Bias"; Angwin et al., "Dozens of Companies Are Using Facebook to Exclude."

18. Dougherty, "Google Photos Mistakenly Labels Black People 'Gorillas'"; Perez, "Microsoft Silences Its New A.I. Bot Tay"; McMillan, "It's Not

You, It's It"; Sloane, "Online Ads for High-Paying Jobs Are Targeting Men More Than Women."

19. See Benjamin, *Race after Technology;* and Noble, *Algorithms of Oppression.*

20. Greene, "Science May Have Cured Biased AI"; Natarajan, "Amazon and NSF Collaborate to Accelerate Fairness in AI Research."

21. Dastin, "Amazon Scraps Secret AI Recruiting Tool."

22. Dastin.

23. This is part of a larger trend toward automating aspects of hiring. For a detailed account, see Ajunwa and Greene, "Platforms at Work."

24. There are several superb accounts of the history of inequality and discrimination in computation. These are a few that have informed my thinking on these issues: Hicks, *Programmed Inequality;* McIlwain, *Black Software;* Light, "When Computers Were Women"; and Ensmenger, *Computer Boys Take Over.*

25. Cetina, *Epistemic Cultures,* 3.

26. Merler et al., "Diversity in Faces."

27. Buolamwini and Gebru, "Gender Shades"; Raj et al. "Saving Face."

28. Merler et al., "Diversity in Faces."

29. "YFCC100M Core Dataset."

30. Merler et al., "Diversity in Faces," 1.

31. There are many excellent books on these issues, but in particular, see Roberts, *Fatal Invention,* 18–41; and Nelson, *Social Life of DNA,* 43. See also Tishkoff and Kidd, "Implications of Biogeography."

32. Browne, "Digital Epidermalization," 135.

33. Benthall and Haynes, "Racial Categories in Machine Learning."

34. Mitchell, "Need for Biases in Learning Generalizations."

35. Dietterich and Kong, "Machine Learning Bias, Statistical Bias."

36. Domingos, "Useful Things to Know about Machine Learning."

37. *Maddox v. State,* 32 Ga. 5S7, 79 Am. Dec. 307; *Pierson v. State,* 18 Tex. App. 55S; *Hinkle v. State,* 94 Ga. 595, 21 S. E. 601.

38. Tversky and Kahneman, "Judgment under Uncertainty."

39. Greenwald and Krieger, "Implicit Bias," 951.

40. Fellbaum, *WordNet,* xviii. Below I am drawing on research into ImageNet conducted with Trevor Paglen. See Crawford and Paglen, "Excavating AI."

41. Fellbaum, xix.

42. Nelson and Kucera, *Brown Corpus Manual.*

43. Borges, "The Analytical Language of John Wilkins."

44. These are some of the categories that have now been deleted entirely from ImageNet as of October 1, 2020.

45. See Keyes, "Misgendering Machines."

46. Drescher, "Out of DSM."
47. See Bayer, *Homosexuality and American Psychiatry.*
48. Keyes, "Misgendering Machines."
49. Hacking, "Making Up People," 23.
50. Bowker and Star, *Sorting Things Out,* 196.
51. This is drawn from Lakoff, *Women, Fire, and Dangerous Things.*
52. ImageNet Roulette was one of the outputs of a multiyear research collaboration between the artist Trevor Paglen and me, in which we studied the underlying logic of multiple benchmark training sets in AI. ImageNet Roulette, led by Paglen and produced by Leif Ryge, was an app that allowed people to interact with a neural net trained on the "person" category of ImageNet. People could upload images of themselves—or news images or historical photographs—to see how ImageNet would label them. People could also see how many of the labels are bizarre, racist, misogynist, and otherwise problematic. The app was designed to show people these concerning labels while warning them in advance of the potential results. All uploaded image data were immediately deleted on processing. See Crawford and Paglen, "Excavating AI."
53. Yang et al., "Towards Fairer Datasets," paragraph 4.2.
54. Yang et al., paragraph 4.3.
55. Markoff, "Seeking a Better Way to Find Web Images."
56. Browne, *Dark Matters,* 114.
57. Scheuerman et al., "How We've Taught Algorithms to See Identity."
58. UTKFace Large Scale Face Dataset, https://susanqq.github.io/UTKFace.
59. Bowker and Star, *Sorting Things Out,* 197.
60. Bowker and Star, 198.
61. Edwards and Hecht, "History and the Technopolitics of Identity," 627.
62. Haraway, *Modest_Witness@Second_Millennium,* 234.
63. Stark, "Facial Recognition Is the Plutonium of AI," 53.
64. In order of the examples, see Wang and Kosinski, "Deep Neural Networks Are More Accurate than Humans"; Wu and Zhang, "Automated Inference on Criminality Using Face Images"; and Angwin et al., "Machine Bias."
65. Agüera y Arcas, Mitchell, and Todorov, "Physiognomy's New Clothes."
66. Nielsen, *Disability History of the United States;* Kafer, *Feminist, Queer, Crip;* Siebers, *Disability Theory.*
67. Whittaker et al., "Disability, Bias, and AI."
68. Hacking, "Kinds of People," 289.
69. Bowker and Star, *Sorting Things Out,* 31.
70. Bowker and Star, 6.
71. Eco, Infinity of Lists.
72. Douglass, "West India Emancipation."

5
Affect

1. Particular thanks to Alex Campolo, who was my research assistant and interlocutor for this chapter, and for his research into Ekman and the history of emotions.

2. "Emotion Detection and Recognition"; Schwartz, "Don't Look Now."

3. Ohtake, "Psychologist Paul Ekman Delights at Exploratorium."

4. Ekman, *Emotions Revealed,* 7.

5. For an overview of researchers who have found flaws in the claim that emotional expressions are universal and can be predicted by AI, see Heaven, "Why Faces Don't Always Tell the Truth."

6. Barrett et al., "Emotional Expressions Reconsidered."

7. Nilsson, "How AI Helps Recruiters."

8. Sánchez-Monedero and Dencik, "Datafication of the Workplace," 48; Harwell, "Face-Scanning Algorithm."

9. Byford, "Apple Buys Emotient."

10. Molnar, Robbins, and Pierson, "Cutting Edge."

11. Picard, "Affective Computing Group."

12. "Affectiva Human Perception AI Analyzes Complex Human States."

13. Schwartz, "Don't Look Now."

14. See, e.g., Nilsson, "How AI Helps Recruiters."

15. "Face: An AI Service That Analyzes Faces in Images."

16. "Amazon Rekognition Improves Face Analysis"; "Amazon Rekognition—Video and Image."

17. Barrett et al., "Emotional Expressions Reconsidered," 1.

18. Sedgwick, Frank, and Alexander, *Shame and Its Sisters,* 258.

19. Tomkins, *Affect Imagery Consciousness.*

20. Tomkins.

21. Leys, *Ascent of Affect,* 18.

22. Tomkins, *Affect Imagery Consciousness,* 23.

23. Tomkins, 23.

24. Tomkins, 23.

25. For Ruth Leys, this "radical dissociation between feeling and cognition" is the major reason for its attractiveness to theorists in the humanities, most notably Eve Kosofsky Sedgwick, who wants to revalorize our experiences of error or confusion into new forms of freedom. Leys, *Ascent of Affect,* 35; Sedgwick, *Touching Feeling.*

26. Tomkins, *Affect Imagery Consciousness,* 204.

27. Tomkins, 206; Darwin, *Expression of the Emotions;* Duchenne (de Boulogne), *Mécanisme de la physionomie humaine.*

28. Tomkins, 243, quoted in Leys, *Ascent of Affect,* 32.

29. Tomkins, *Affect Imagery Consciousness,* 216.
30. Ekman, *Nonverbal Messages,* 45.
31. Tuschling, "Age of Affective Computing," 186.
32. Ekman, *Nonverbal Messages,* 45.
33. Ekman, 46.
34. Ekman, 46.
35. Ekman, 46.
36. Ekman, 46.
37. Ekman, 46.
38. Ekman and Rosenberg, *What the Face Reveals,* 375.
39. Tomkins and McCarter, "What and Where Are the Primary Affects?"
40. Russell, "Is There Universal Recognition of Emotion from Facial Expression?" 116.
41. Leys, *Ascent of Affect,* 93.
42. Ekman and Rosenberg, *What the Face Reveals,* 377.
43. Ekman, Sorenson, and Friesen, "Pan-Cultural Elements in Facial Displays of Emotion," 86, 87.
44. Ekman and Friesen, "Constants across Cultures in the Face and Emotion," 128.
45. Aristotle, *Categories,* 70b8–13, 527.
46. Aristotle, 805a, 27–30, 87.
47. It would be difficult to overstate the influence of this work, which has since fallen into disrepute: by 1810 it went through sixteen German and twenty English editions. Graham, "Lavater's Physiognomy in England," 561.
48. Gray, *About Face,* 342.
49. Courtine and Haroche, *Histoire du visage,* 132.
50. Ekman, "Duchenne and Facial Expression of Emotion."
51. Duchenne (de Boulogne), *Mécanisme de la physionomie humaine.*
52. Clarac, Massion, and Smith, "Duchenne, Charcot and Babinski," 362–63.
53. Delaporte, *Anatomy of the Passions,* 33.
54. Delaporte, 48–51.
55. Daston and Galison, *Objectivity.*
56. Darwin, *Expression of the Emotions in Man and Animals,* 12, 307.
57. Leys, *Ascent of Affect,* 85; Russell, "Universal Recognition of Emotion," 114.
58. Ekman and Friesen, "Nonverbal Leakage and Clues to Deception," 93.
59. Pontin, "Lie Detection."
60. Ekman and Friesen, "Nonverbal Leakage and Clues to Deception," 94. In a footnote, Ekman and Friesen explained: "Our own research and the evidence from the neurophysiology of visual perception strongly suggest

that micro-expressions that are as short as one motion-picture frame (1/50 of a second) can be perceived. That these micro-expressions are not usually seen must depend upon their being embedded in other expressions which distract attention, their infrequency, or some learned perceptual habit of ignoring fast facial expressions."

61. Ekman, Sorenson, and Friesen, "Pan-Cultural Elements in Facial Displays of Emotion," 87.

62. Ekman, Friesen, and Tomkins, "Facial Affect Scoring Technique," 40.

63. Ekman, *Nonverbal Messages,* 97.

64. Ekman, 102.

65. Ekman and Rosenberg, *What the Face Reveals.*

66. Ekman, *Nonverbal Messages,* 105.

67. Ekman, 169.

68. Eckman, 106; Aleksander, *Artificial Vision for Robots.*

69. "Magic from Invention."

70. Bledsoe, "Model Method in Facial Recognition."

71. Molnar, Robbins, and Pierson, "Cutting Edge."

72. Kanade, *Computer Recognition of Human Faces.*

73. Kanade, 16.

74. Kanade, Cohn, and Tian, "Comprehensive Database for Facial Expression Analysis," 6.

75. See Kanade, Cohn, and Tian; Lyons et al., "Coding Facial Expressions with Gabor Wavelets"; and Goeleven et al., "Karolinska Directed Emotional Faces."

76. Lucey et al., "Extended Cohn-Kanade Dataset (CK+)."

77. McDuff et al., "Affectiva-MIT Facial Expression Dataset (AM-FED)."

78. McDuff et al.

79. Ekman and Friesen, *Facial Action Coding System (FACS).*

80. Foreman, "Conversation with: Paul Ekman"; Taylor, "2009 Time 100"; Paul Ekman Group.

81. Weinberger, "Airport Security," 413.

82. Halsey, "House Member Questions $900 Million TSA 'SPOT' Screening Program."

83. Ekman, "Life's Pursuit"; Ekman, *Nonverbal Messages,* 79–81.

84. Mead, "Review of *Darwin and Facial Expression,*" 209.

85. Tomkins, *Affect Imagery Consciousness,* 216.

86. Mead, "Review of *Darwin and Facial Expression,*" 212. See also Fridlund, "Behavioral Ecology View of Facial Displays." Ekman later conceded to many of Mead's points. See Ekman, "Argument for Basic Emotions"; Ekman, *Emotions Revealed;* and Ekman, "What Scientists Who Study Emotion Agree About." Ekman also had his defenders. See Cowen et al., "Mapping the Pas-

sions"; and Elfenbein and Ambady, "Universality and Cultural Specificity of Emotion Recognition."

87. Fernández-Dols and Russell, *Science of Facial Expression,* 4.

88. Gendron and Barrett, *Facing the Past,* 30.

89. Vincent, "AI 'Emotion Recognition' Can't Be Trusted.'" Disability studies scholars have also noted that assumptions about how biology and bodies function can also raise concerns around bias, especially when automated through technology. See Whittaker et al., "Disability, Bias, and AI."

90. Izard, "Many Meanings/Aspects of Emotion."

91. Leys, *Ascent of Affect,* 22.

92. Leys, 92.

93. Leys, 94.

94. Leys, 94.

95. Barrett, "Are Emotions Natural Kinds?" 28.

96. Barrett, 30.

97. See, e.g., Barrett et al., "Emotional Expressions Reconsidered."

98. Barrett et al., 40.

99. Kappas, "Smile When You Read This," 39, emphasis added.

100. Kappas, 40.

101. Barrett et al., 46.

102. Barrett et al., 47–48.

103. Barrett et al., 47, emphasis added.

104. Apelbaum, "One Thousand and One Nights."

105. See, e.g., Hoft, "Facial, Speech and Virtual Polygraph Analysis."

106. Rhue, "Racial Influence on Automated Perceptions of Emotions."

107. Barrett et al., "Emotional Expressions Reconsidered,"48.

108. See, e.g., Connor, "Chinese School Uses Facial Recognition"; and Du and Maki, "AI Cameras That Can Spot Shoplifters."

6
State

1. NOFORN stands for Not Releasable to Foreign Nationals. "Use of the 'Not Releasable to Foreign Nationals' (NOFORN) Caveat."

2. The Five Eyes is a global intelligence alliance comprising Australia, Canada, New Zealand, the United Kingdom, and the United States. "Five Eyes Intelligence Oversight and Review Council."

3. Galison, "Removing Knowledge," 229.

4. Risen and Poitras, "N.S.A. Report Outlined Goals for More Power"; Müller-Maguhn et al., "The NSA Breach of Telekom and Other German Firms."

5. FOXACID is software developed by the Office of Tailored Access Operations, now Computer Network Operations, a cyberwarfare intelligence-gathering unit of the NSA.

6. Schneier, "Attacking Tor." Document available at "NSA Phishing Tactics and Man in the Middle Attacks."

7. Swinhoe, "What Is Spear Phishing?"

8. "Strategy for Surveillance Powers."

9. Edwards, *Closed World.*

10. Edwards.

11. Edwards, 198.

12. Mbembé, *Necropolitics,* 82.

13. Bratton, *Stack,* 151.

14. For an excellent account of the history of the internet in the United States, see Abbate, *Inventing the Internet.*

15. SHARE Foundation, "Serbian Government Is Implementing Unlawful Video Surveillance."

16. Department of International Cooperation Ministry of Science and Technology, "Next Generation Artificial Intelligence Development Plan."

17. Chun, *Control and Freedom;* Hu, *Prehistory of the Cloud,* 87–88.

18. Cave and ÓhÉigeartaigh, "AI Race for Strategic Advantage."

19. Markoff, "Pentagon Turns to Silicon Valley for Edge."

20. Brown, *Department of Defense Annual Report.*

21. Martinage, "Toward a New Offset Strategy," 5–16.

22. Carter, "Remarks on 'the Path to an Innovative Future for Defense'"; Pellerin, "Deputy Secretary."

23. The origins of U.S. military offsets can be traced back to December 1952, when the Soviet Union had almost ten times more conventional military divisions than the United States. President Dwight Eisenhower turned to nuclear deterrence as a way to "offset" these odds. The strategy involved not only the threat of the retaliatory power of the U.S. nuclear forces but also accelerating the growth of the U.S. weapons stockpile, as well as developing long-range jet bombers, the hydrogen bomb, and eventually intercontinental ballistic missiles. It also included increased reliance on espionage, sabotage, and covert operations. In the 1970s and 1980s, U.S. military strategy turned to computational advances in analytics and logistics, building on the influence of such military architects as Robert McNamara in search of military supremacy. This Second Offset could be seen in military engagements like Operation Desert Storm during the Gulf War in 1991, where reconnaissance, suppression of enemy defenses, and precision-guided munitions dominated how the United States not only fought the war but thought and spoke about it. Yet as Russia and China began to adopt these capacities and deploy digi-

tal networks for warfare, anxiety grew to reestablish a new kind of strategic advantage. See McNamara and Blight, *Wilson's Ghost.*

24. Pellerin, "Deputy Secretary."

25. Gellman and Poitras, "U.S., British Intelligence Mining Data."

26. Deputy Secretary of Defense to Secretaries of the Military Departments et al.

27. Deputy Secretary of Defense to Secretaries of the Military Departments et al.

28. Michel, *Eyes in the Sky,* 134.

29. Michel, 135.

30. Cameron and Conger, "Google Is Helping the Pentagon Build AI for Drones."

31. For example, Gebru et al., "Fine-Grained Car Detection for Visual Census Estimation."

32. Fang, "Leaked Emails Show Google Expected Lucrative Military Drone AI Work."

33. Bergen, "Pentagon Drone Program Is Using Google AI."

34. Shane and Wakabayashi, "'Business of War.'"

35. Smith, "Technology and the US Military."

36. When the JEDI contract was ultimately awarded to Microsoft, Brad Smith, the president of Microsoft, explained that the reason that Microsoft won the contract was that it was seen "not just as a sales opportunity, but really, a very large-scale engineering project." Stewart and Carlson, "President of Microsoft Says It Took Its Bid."

37. Pichai, "AI at Google."

38. Pichai. Project Maven was subsequently picked up by Anduril Industries, a secretive tech startup founded by Oculus Rift's Palmer Luckey. Fang, "Defense Tech Startup."

39. Whittaker et al., *AI Now Report 2018.*

40. Schmidt quoted in Scharre et al., "Eric Schmidt Keynote Address."

41. As Suchman notes, "'Killing people correctly' under the laws of war requires adherence to the Principle of Distinction and the identification of an imminent threat." Suchman, "Algorithmic Warfare and the Reinvention of Accuracy," n. 18.

42. Suchman.

43. Suchman.

44. Hagendorff, "Ethics of AI Ethics."

45. Brustein and Bergen, "Google Wants to Do Business with the Military."

46. For more on why municipalities should more carefully assess the risks of algorithmic platforms, see Green, *Smart Enough City.*

47. Thiel, "Good for Google, Bad for America."

48. Steinberger, "Does Palantir See Too Much?"

49. Weigel, "Palantir goes to the Frankfurt School."

50. Dilanian, "US Special Operations Forces Are Clamoring to Use Software."

51. "War against Immigrants."

52. Alden, "Inside Palantir, Silicon Valley's Most Secretive Company."

53. Alden, "Inside Palantir, Silicon Valley's Most Secretive Company."

54. Waldman, Chapman, and Robertson, "Palantir Knows Everything about You."

55. Joseph, "Data Company Directly Powers Immigration Raids in Workplace"; Anzilotti, "Emails Show That ICE Uses Palantir Technology to Detain Undocumented Immigrants."

56. Andrew Ferguson, conversation with author, June 21, 2019.

57. Brayne, "Big Data Surveillance." Brayne also notes that the migration of law enforcement to intelligence was occurring even before the shift to predictive analytics, given such court decisions as Terry v. Ohio and Whren v. United States that made it easier for law enforcement to circumvent probable cause and produced a proliferation of pretext stops.

58. Richardson, Schultz, and Crawford, "Dirty Data, Bad Predictions."

59. Brayne, "Big Data Surveillance," 997.

60. Brayne, 997.

61. See, e.g., French and Browne, "Surveillance as Social Regulation."

62. Crawford and Schultz, "AI Systems as State Actors."

63. Cohen, *Between Truth and Power;* Calo and Citron, "Automated Administrative State."

64. "Vigilant Solutions"; Maass and Lipton, "What We Learned."

65. Newman, "Internal Docs Show How ICE Gets Surveillance Help."

66. England, "UK Police's Facial Recognition System."

67. Scott, *Seeing Like a State.*

68. Haskins, "How Ring Transmits Fear to American Suburbs."

69. Haskins, "Amazon's Home Security Company."

70. Haskins.

71. Haskins. "Amazon Requires Police to Shill Surveillance Cameras."

72. Haskins, "Amazon Is Coaching Cops."

73. Haskins.

74. Haskins.

75. Hu, *Prehistory of the Cloud,* 115.

76. Hu, 115.

77. Benson, "'Kill 'Em and Sort It Out Later,'" 17.

78. Hajjar, "Lawfare and Armed Conflicts," 70.

79. Scahill and Greenwald, "NSA's Secret Role in the U.S. Assassination Program."

80. Cole, "'We Kill People Based on Metadata.'"
81. Priest, "NSA Growth Fueled by Need to Target Terrorists."
82. Gibson quoted in Ackerman, "41 Men Targeted but 1,147 People Killed."
83. Tucker, "Refugee or Terrorist?"
84. Tucker.
85. O'Neil, *Weapons of Math Destruction,* 288–326.
86. Fourcade and Healy, "Seeing Like a Market."
87. Eubanks, *Automating Inequality.*
88. Richardson, Schultz, and Southerland, "Litigating Algorithms," 19.
89. Richardson, Schultz, and Southerland, 23.
90. Agre, *Computation and Human Experience,* 240.
91. Bratton, *Stack,* 140.
92. Hu, *Prehistory of the Cloud,* 89.
93. Nakashima and Warrick, "For NSA Chief, Terrorist Threat Drives Passion."
94. Document available at Maass, "Summit Fever."
95. The future of the Snowden archive itself is uncertain. In March 2019, it was announced that the *Intercept*—the publication that Glenn Greenwald established with Laura Poitras and Jeremy Scahill after they shared the Pulitzer Prize for their reporting on the Snowden materials—was no longer going to fund the Snowden archive. Tani, "Intercept Shuts Down Access to Snowden Trove."

Conclusion

1. Silver et al., "Mastering the Game of Go without Human Knowledge."
2. Silver et al., 357.
3. Full talk at the Artificial Intelligence Channel: *Demis Hassabis, DeepMind—Learning from First Principles.* See also Knight, "Alpha Zero's 'Alien' Chess Shows the Power."
4. *Demis Hassabis, DeepMind—Learning from First Principles.*
5. For more on the myths of "magic" in AI, see Elish and boyd, "Situating Methods in the Magic of Big Data and AI."
6. Meredith Broussard notes that playing games has been dangerously conflated with intelligence. She cites the programmer George V. Neville-Neil, who argues: "We have had nearly 50 years of human/computer competition in the game of chess, but does this mean that any of those computers are intelligent? No, it does not—for two reasons. The first is that chess is not a test of intelligence; it is the test of a particular skill—the skill of playing chess. If I could beat a Grandmaster at chess and yet not be able to hand you

the salt at the table when asked, would I be intelligent? The second reason is that thinking chess was a test of intelligence was based on a false cultural premise that brilliant chess players were brilliant minds, more gifted than those around them. Yes, many intelligent people excel at chess, but chess, or any other single skill, does not denote intelligence." Broussard, *Artificial Unintelligence,* 206.

7. Galison, "Ontology of the Enemy."

8. Campolo and Crawford, "Enchanted Determinism."

9. Bailey, "Dimensions of Rhetoric in Conditions of Uncertainty," 30.

10. Bostrom, *Superintelligence.*

11. Bostrom.

12. Strand, "Keyword: Evil," 64–65.

13. Strand, 65.

14. Hardt and Negri, *Assembly,* 116, emphasis added.

15. Wakabayashi, "Google's Shadow Work Force."

16. Quoted in McNeil, "Two Eyes See More Than Nine," 23.

17. On the idea of data as capital, see Sadowski, "When Data Is Capital."

18. Harun Farocki discussed in Paglen, "Operational Images."

19. For a summary, see Heaven, "Why Faces Don't Always Tell the Truth."

20. Nietzsche, *Sämtliche Werke,* 11:506.

21. Wang and Kosinski, "Deep Neural Networks Are More Accurate Than Humans"; Kleinberg et al., "Human Decisions and Machine Predictions"; Crosman, "Is AI a Threat to Fair Lending?"; Seo et al., "Partially Generative Neural Networks."

22. Pugliese, "Death by Metadata."

23. Suchman, "Algorithmic Warfare and the Reinvention of Accuracy."

24. Simmons, "Rekor Software Adds License Plate Reader Technology."

25. Lorde, *Master's Tools.*

26. Schaake, "What Principles Not to Disrupt."

27. Jobin, Ienca, and Vayena, "Global Landscape of AI Ethics Guidelines."

28. Mattern, "Calculative Composition," 572.

29. For more on why AI ethics frameworks are limited in effectiveness, see Crawford et al., *AI Now 2019 Report.*

30. Mittelstadt, "Principles Alone Cannot Guarantee Ethical AI." See also Metcalf, Moss, and boyd, "Owning Ethics."

31. For recent scholarship that addresses important practical steps to do this without replicating forms of extraction and harm, see Costanza-Chock, *Design Justice.*

32. Winner, *The Whale and the Reactor,* 9.

33. Mbembé, *Critique of Black Reason,* 3.

34. Bangstad et al., "Thoughts on the Planetary."

35. Haraway, *Simians, Cyborgs, and Women,* 161.
36. Mohamed, Png, and Isaac, "Decolonial AI," 405.
37. "Race after Technology, Ruha Benjamin."
38. Bangstad et al., "Thoughts on the Planetary."

Coda

1. *Blue Origin's Mission.*
2. *Blue Origin's Mission.*
3. Powell, "Jeff Bezos Foresees a Trillion People."
4. Bezos, *Going to Space to Benefit Earth.*
5. Bezos.
6. Foer, "Jeff Bezos's Master Plan."
7. Foer.
8. "Why Asteroids."
9. Welch, "Elon Musk."
10. Cuthbertson, "Elon Musk Really Wants to 'Nuke Mars.'"
11. Rein, Tamayo, and Vokrouhlicky, "Random Walk of Cars."
12. Gates, "Bezos' Blue Origin Seeks Tax Incentives."
13. Marx, "Instead of Throwing Money at the Moon"; O'Neill, *High Frontier.*
14. "Our Mission."
15. Davis, "Gerard K. O'Neill on Space Colonies."
16. Meadows et al., *Limits to Growth.*
17. Scharmen, *Space Settlements,* 216. In recent years, scholars have suggested that the Club of Rome's models were overly optimistic, underestimating the rapid rate of extraction and resource consumption worldwide and the climate implications of greenhouse gases and industrial waste heat. See Turner, "Is Global Collapse Imminent?"
18. The case for a no-growth model that involves staying on the planet has been made by many academics in the limits to growth movement. See, e.g., Trainer, *Renewable Energy Cannot Sustain a Consumer Society.*
19. Scharmen, *Space Settlements,* 91.
20. One wonders how the Bezos mission would differ had he been inspired instead by the science fiction author Philip K. Dick, who wrote the short story "Autofac" in 1955. In it, human survivors of an apocalyptic war are left on Earth with the "autofacs"—autonomous, self-replicating factory machines. The autofacs had been tasked with producing consumer goods in prewar society but could no longer stop, consuming the planet's resources and threatening the survival of the last people left. The only way to survive was to trick the artificial intelligence machines to fight against each other

over a critical element they need for manufacturing: the rare earth element tungsten. It seems to succeed, and wild vines begin to grow throughout the factories, and farmers can return to the land. Only later do they realize that the autofacs had sought more resources deep in Earth's core and would soon launch thousands of self-replicating "seeds" to mine the rest of the galaxy. Dick, "Autofac."

21. NASA, "Outer Space Treaty of 1967."

22. U.S. Commercial Space Launch Competitiveness Act."

23. Wilson, "Top Lobbying Victories of 2015."

24. Shaer, "Asteroid Miner's Guide to the Galaxy."

25. As Mark Andrejevic writes, "The promise of technological immortality is inseparable from automation, which offers to supplant human limitations at every turn." Andrejevic, *Automated Media,* 1.

26. Reichhardt, "First Photo from Space."

27. See, e.g., Pulitzer Prize–winning journalist Wayne Biddle's account of von Braun as a war criminal who participated in the brutal treatment of slave laborers under the Nazi regime. Biddle, *Dark Side of the Moon.*

28. Grigorieff, "Mittelwerk/Mittelbau/Camp Dora."

29. Ward, *Dr. Space.*

30. Keates, "Many Places Amazon CEO Jeff Bezos Calls Home."

31. Center for Land Use Interpretation, "Figure 2 Ranch, Texas."