KATE CRAWFORD

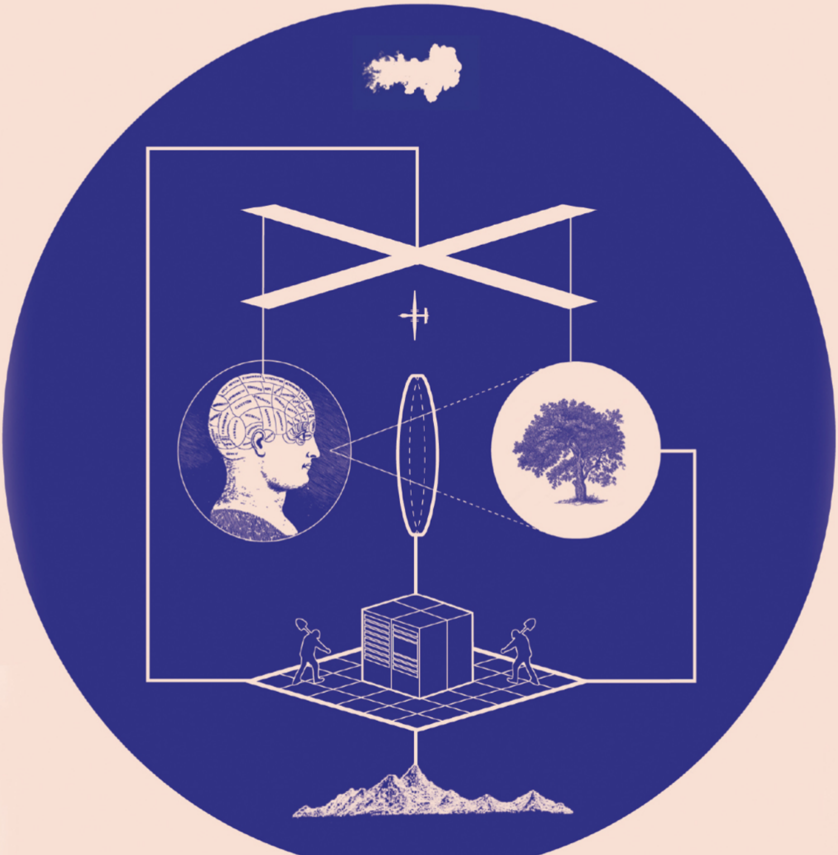

ATLAS OF AI

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

Contents

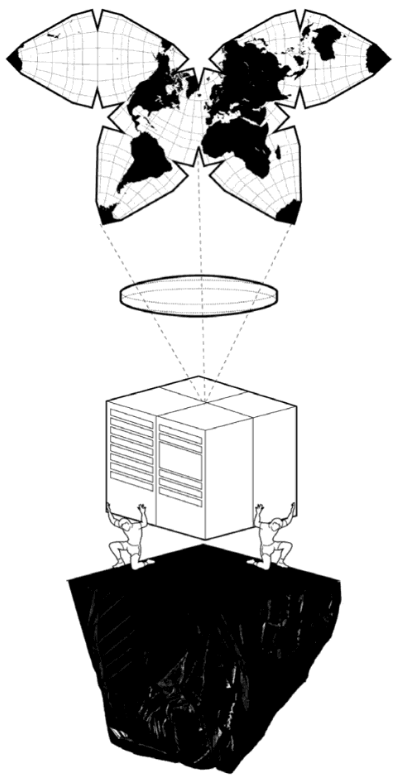

Introduction

The Smartest Horse in the World

At the end of the nineteenth century, Europe was captivated by a horse called Hans. "Clever Hans" was nothing less than a marvel: he could solve math problems, tell time, identify days on a calendar, differentiate musical tones, and spell out words and sentences. People flocked to watch the German stallion tap out answers to complex problems with his hoof and consistently arrive at the right answer. "What is two plus three?" Hans would diligently tap his hoof on the ground five times. "What day of the week is it?" The horse would then tap his hoof to indicate each letter on a purpose-built letter board and spell out the correct answer. Hans even mastered more complex questions, such as, "I have a number in mind. I subtract nine and have three as a remainder. What is the number?" By 1904, Clever Hans was an international celebrity, with the *New York Times* championing him as "Berlin's Wonderful Horse; He Can Do Almost Everything but Talk."[1]

Hans's trainer, a retired math teacher named Wilhelm von Osten, had long been fascinated by animal intelligence.

Von Osten had tried and failed to teach kittens and bear cubs cardinal numbers, but it wasn't until he started working with his own horse that he had success. He first taught Hans to count by holding the animal's leg, showing him a number, and then tapping on the hoof the correct number of times. Soon Hans responded by accurately tapping out simple sums. Next von Osten introduced a chalkboard with the alphabet spelled out, so Hans could tap a number for each letter on the board. After two years of training, von Osten was astounded by the animal's strong grasp of advanced intellectual concepts. So he took Hans on the road as proof that animals could reason. Hans became the viral sensation of the belle époque.

But many people were skeptical, and the German board of education launched an investigative commission to test Von Osten's scientific claims. The Hans Commission was led by the psychologist and philosopher Carl Stumpf and his assistant Oskar Pfungst, and it included a circus manager, a retired schoolteacher, a zoologist, a veterinarian, and a cavalry officer. Yet after extensive questioning of Hans, both with his trainer present and without, the horse maintained his record of correct answers, and the commission could find no evidence of deception. As Pfungst later wrote, Hans performed in front of "thousands of spectators, horse-fanciers, trick-trainers of first rank, and not one of them during the course of many months' observations are able to discover any kind of regular signal" between the questioner and the horse.[2]

The commission found that the methods Hans had been taught were more like "teaching children in elementary schools" than animal training and were "worthy of scientific examination."[3] But Strumpf and Pfungst still had doubts. One finding in particular troubled them: when the questioner did not know the answer or was standing far away, Hans rarely gave the correct answer. This led Pfungst and Strumpf to con-

Wilhelm von Osten and Clever Hans

sider whether some sort of unintentional signal had been providing Hans with the answers.

As Pfungst would describe in his 1911 book, their intuition was right: the questioner's posture, breathing, and facial expression would subtly change around the moment Hans reached the right answer, prompting Hans to stop there.[4] Pfungst later tested this hypothesis on human subjects and confirmed his result. What fascinated him most about this discovery was that questioners were generally unaware that they were providing pointers to the horse. The solution to the Clever Hans riddle, Pfungst wrote, was the unconscious direction from the horse's questioners.[5] The horse was trained to produce the results his owner wanted to see, but audiences felt that this was not the extraordinary intelligence they had imagined.

The story of Clever Hans is compelling from many angles: the relationship between desire, illusion, and action, the business of spectacles, how we anthropomorphize the nonhuman, how biases emerge, and the politics of intelligence. Hans inspired a term in psychology for a particular type of conceptual trap, the Clever Hans Effect or observer-expectancy effect, to describe the influence of experimenters' unintentional cues on their subjects. The relationship between Hans and von Osten points to the complex mechanisms by which biases find their ways into systems and how people become entangled with the phenomena they study. The story of Hans is now used in machine learning as a cautionary reminder that you can't always be sure of what a model has learned from the data it has been given.[6] Even a system that appears to perform spectacularly in training can make terrible predictions when presented with novel data in the world.

This opens a central question of this book: How is intelligence "made," and what traps can that create? At first glance, the story of Clever Hans is a story of how one man constructed intelligence by training a horse to follow cues and emulate humanlike cognition. But at another level, we see that the practice of making intelligence was considerably broader. The endeavor required validation from multiple institutions, including academia, schools, science, the public, and the military. Then there was the market for von Osten and his remarkable horse—emotional and economic investments that drove the tours, the newspaper stories, and the lectures. Bureaucratic authorities were assembled to measure and test the horse's abilities. A constellation of financial, cultural, and scientific interests had a part to play in the construction of Hans's intelligence and a stake in whether it was truly remarkable.

We can see two distinct mythologies at work. The first myth is that nonhuman systems (be it computers or horses) are analogues for human minds. This perspective assumes that with sufficient training, or enough resources, humanlike intelligence can be created from scratch, without addressing the mance, and considerable patience, yet these were not recognized as intelligence. As author and engineer Ellen Ullman puts it, this belief that the mind is like a computer, and vice versa, has "infected decades of thinking in the computer and cognitive sciences," creating a kind of original sin for the field.[15] It is the ideology of Cartesian dualism in artificial intelligence: where AI is narrowly understood as disembodied intelligence, removed from any relation to the material world.

What Is AI? Neither Artificial nor Intelligent

Let's ask the deceptively simple question, What is artificial intelligence? If you ask someone in the street, they might mention Apple's Siri, Amazon's cloud service, Tesla's cars, or Google's search algorithm. If you ask experts in deep learning, they might give you a technical response about how neural nets are organized into dozens of layers that receive labeled data, are assigned weights and thresholds, and can classify data in ways that cannot yet be fully explained.[16] In 1978, when discussing expert systems, Professor Donald Michie described AI as knowledge refining, where "a reliability and competence of codification can be produced which far surpasses the highest level that the unaided human expert has ever, perhaps even could ever, attain."[17] In one of the most popular textbooks on the subject, Stuart Russell and Peter Norvig state that AI is the attempt to understand and build intelligent entities. "Intelligence is concerned mainly with rational action," they claim. "Ideally, an intelligent agent takes the best possible action in a situation."[18]

Each way of defining artificial intelligence is doing work, setting a frame for how it will be understood, measured, valued, and governed. If AI is defined by consumer brands for corporate infrastructure, then marketing and advertising have predetermined the horizon. If AI systems are seen as more reliable or rational than any human expert, able to take the "best possible action," then it suggests that they should be trusted to make high-stakes decisions in health, education, and criminal justice. When specific algorithmic techniques are the sole focus, it suggests that only continual technical progress matters, with no consideration of the computational cost of those approaches and their far-reaching impacts on a planet under strain.

In contrast, in this book I argue that AI is neither *artificial* nor *intelligent.* Rather, artificial intelligence is both embodied and material, made from natural resources, fuel, human labor, infrastructures, logistics, histories, and classifications. AI systems are not autonomous, rational, or able to discern anything without extensive, computationally intensive training with large datasets or predefined rules and rewards. In fact, artificial intelligence as we know it depends entirely on a much wider set of political and social structures. And due to the capital required to build AI at scale and the ways of seeing that it optimizes AI systems are ultimately designed to serve existing dominant interests. In this sense, artificial intelligence is a registry of power.

In this book we'll explore how artificial intelligence is made, in the widest sense, and the economic, political, cultural, and historical forces that shape it. Once we connect AI within these broader structures and social systems, we can escape the notion that artificial intelligence is a purely technical domain. At a fundamental level, AI is technical and social practices, institutions and infrastructures, politics and culture. Computational reason and embodied work are deeply interlinked: AI systems both reflect and produce social relations and understandings of the world.

It's worth noting that the term "artificial intelligence" can create discomfort in the computer science community. The phrase has moved in and out of fashion over the decades and is used more in marketing than by researchers. "Machine learning" is more commonly used in the technical literature. Yet the nomenclature of AI is often embraced during funding application season, when venture capitalists come bearing checkbooks, or when researchers are seeking press attention for a new scientific result. As a result, the term is both used and rejected in ways that keep its meaning in flux. For my purposes, I use AI to talk about the massive industrial formation that includes politics, labor, culture, and capital. When I refer to machine learning, I'm speaking of a range of technical approaches (which are, in fact, social and infrastructural as well, although rarely spoken about as such).

But there are significant reasons *why* the field has been focused so much on the technical—algorithmic breakthroughs, incremental product improvements, and greater convenience. The structures of power at the intersection of technology, capital, and governance are well served by this narrow, abstracted analysis. To understand how AI is fundamentally political, we need to go beyond neural nets and statistical pattern recognition to instead ask *what* is being optimized, and *for whom*, and *who* gets to decide. Then we can trace the implications of those choices.

Seeing AI Like an Atlas

How can an atlas help us to understand how artificial intelligence is made? An atlas is an unusual type of book. It is a collection of disparate parts, with maps that vary in resolution from a satellite view of the planet to a zoomed-in detail of an archipelago. When you open an atlas, you may be seeking specific information about a particular place—or perhaps you are wandering, following your curiosity, and finding unexpected pathways and new perspectives. As historian of science Lorraine Daston observes, all scientific atlases seek to school the eye, to focus the observer's attention on particular telling details and significant characteristics.[19] An atlas presents you with a particular viewpoint of the world, with the imprimatur of science—scales and ratios, latitudes and longitudes—and a sense of form and consistency.

Yet an atlas is as much an act of creativity—a subjective, political, and aesthetic intervention—as it is a scientific collection. The French philosopher Georges Didi-Huberman thinks of the atlas as something that inhabits the aesthetic paradigm of the visual and the epistemic paradigm of knowledge. By implicating both, it undermines the idea that science and art are ever completely separate.[20] Instead, an atlas offers us the possibility of rereading the world, linking disparate pieces differently and "reediting and piecing it together again without thinking we are summarizing or exhausting it."[21]

Perhaps my favorite account of how a cartographic approach can be helpful comes from the physicist and technology critic Ursula Franklin: "Maps represent purposeful endeavors: they are meant to be useful, to assist the traveler and bridge the gap between the known and the as yet unknown; they are testaments of collective knowledge and insight."[22]

Maps, at their best, offer us a compendium of open pathways—shared ways of knowing—that can be mixed and combined to make new interconnections. But there are also maps of domination, those national maps where territory is carved along the fault lines of power: from the direct interventions of drawing borders across contested spaces to revealing the colonial paths of empires. By invoking an atlas, I'm suggesting that we need new ways to understand the empires of artificial intelligence. We need a theory of AI that accounts for the states and corporations that drive and dominate it, the extractive mining that leaves an imprint on the planet, the mass capture of data, and the profoundly unequal and increasingly exploitative labor practices that sustain it. These are the shifting tectonics of power in AI. A topographical approach offers different perspectives and scales, beyond the abstract promises of artificial intelligence or the latest machine learning models. The aim is to understand AI in a wider context by walking through the many different landscapes of computation and seeing how they connect.[23]

There's another way in which atlases are relevant here. The field of AI is explicitly attempting to capture the planet in a computationally legible form. This is not a metaphor so much as the industry's direct ambition. The AI industry is making and normalizing its own proprietary maps, as a centralized God's-eye view of human movement, communication, and labor. Some AI scientists have stated their desire to capture the world and to supersede other forms of knowing. AI professor Fei-Fei Li describes her ImageNet project as aiming to "map out the entire world of objects."[24] In their textbook, Russell and Norvig describe artificial intelligence as "relevant to any intellectual task; it is truly a universal field."[25] One of the founders of artificial intelligence and early experimenter in facial recognition, Woody Bledsoe, put it most bluntly: "in the long run, AI is the *only* science."[26] This is a desire not to create an atlas of the world but to be *the* atlas—the dominant way of seeing. This colonizing impulse centralizes power in the AI field: it determines how the world is measured and defined while simultaneously denying that this is an inherently political activity.

Instead of claiming universality, this book is a partial account, and by bringing you along on my investigations, I hope to show you how my views were formed. We will encounter

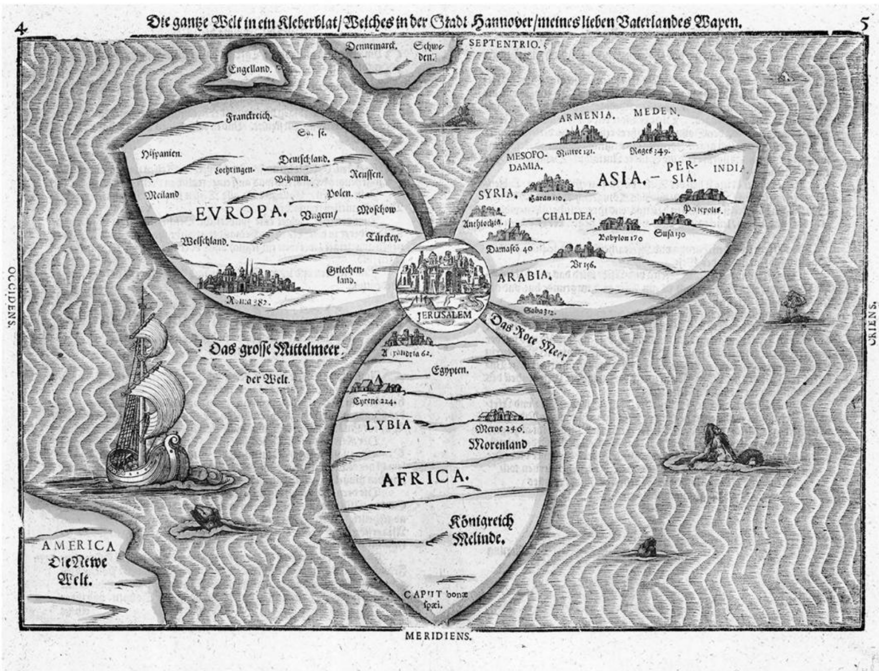

Heinrich Bünting's *mappa mundi*, known as *The Bünting Clover Leaf Map*, which symbolizes the Christian Trinity, with the city of Jerusalem at the center of the world. From *Itinerarium Sacrae Scripturae* (Magdeburg, 1581)

Topographies of Computation

How, at this moment in the twenty-first century, is AI conceptualized and constructed? What is at stake in the turn to artificial intelligence, and what kinds of politics are contained in the way these systems map and interpret the world? What are the social and material consequences of including AI and related algorithmic systems into the decision-making systems of social institutions like education and health care, finance, government operations, workplace interactions and hiring, communication systems, and the justice system? This book is not a story about code and algorithms or the latest thinking in computer vision or natural language processing or reinforcement learning. Many other books do that. Neither is it an ethnographic account of a single community and the effects of AI on their experience of work or housing or medicine—although we certainly need more of those.

Instead, this is an expanded view of artificial intelligence as an *extractive industry.* The creation of contemporary AI systems depends on exploiting energy and mineral resources from the planet, cheap labor, and data at scale. To observe this in action, we will go on a series of journeys to places that reveal the makings of AI.

In chapter 1, we begin in the lithium mines of Nevada, one of the many sites of mineral extraction needed to power contemporary computation. Mining is where we see the extractive politics of AI at their most literal. The tech sector's demand for rare earth minerals, oil, and coal is vast, but the true costs of this extraction is never borne by the industry itself. On the software side, building models for natural language processing and computer vision is enormously energy hungry, and the competition to produce faster and more efficient models has driven computationally greedy methods that expand AI's carbon footprint. From the last trees in Malaysia that were harvested to produce latex for the first transatlantic undersea cables to the giant artificial lake of toxic residues in Inner Mongolia, we trace the environmental and human birthplaces of planetary computation networks and see how they continue to terraform the planet.

Chapter 2 shows how artificial intelligence is made of human labor. We look at the digital pieceworkers paid pennies on the dollar clicking on microtasks so that data systems can seem more intelligent than they are.[31] Our journey will take us inside the Amazon warehouses where employees must keep in time with the algorithmic cadences of a vast logistical empire, and we will visit the Chicago meat laborers on the disassembly lines where animal carcasses are vivisected and prepared for consumption. And we'll hear from the workers who are protesting against the way that AI systems are increasing surveillance and control for their bosses.

Labor is also a story about time. Coordinating the actions of humans with the repetitive motions of robots and line machinery has always involved a controlling of bodies in space and time.[32] From the invention of the stopwatch to Google's TrueTime, the process of time coordination is at the heart of workplace management. AI technologies both require and create the conditions for ever more granular and precise mechanisms of temporal management. Coordinating time demands increasingly detailed information about what people are doing and how and when they do it.

Chapter 3 focuses on the role of data. All publicly accessible digital material—including data that is personal or potentially damaging—is open to being harvested for training datasets that are used to produce AI models. There are gigantic datasets full of people's selfies, of hand gestures, of people driving cars, of babies crying, of newsgroup conversations from the 1990s, all to improve algorithms that perform such functions as facial recognition, language prediction, and object detection. When these collections of data are no longer seen as people's personal material but merely as *infrastructure,* the specific meaning or context of an image or a video is assumed to be irrelevant. Beyond the serious issues of privacy and ongoing surveillance capitalism, the current practices of working with data in AI raise profound ethical, methodological, and epistemological concerns.[33]

And how is all this data used? In chapter 4, we look at the practices of classification in artificial intelligence systems, what sociologist Karin Knorr Cetina calls the "epistemic machinery."[34] We see how contemporary systems use labels to predict human identity, commonly using binary gender, essentialized racial categories, and problematic assessments of character and credit worthiness. A sign will stand in for a system, a proxy will stand for the real, and a toy model will be asked to substitute for the infinite complexity of human subjectivity. By looking at how classifications are made, we see how technical schemas enforce hierarchies and magnify inequity. Machine learning presents us with a regime of normative reasoning that, when in the ascendant, takes shape as a powerful governing rationality.

From here, we travel to the hill towns of Papua New Guinea to explore the history of affect recognition, the idea that facial expressions hold the key to revealing a person's inner emotional state. Chapter 5 considers the claim of the psychologist Paul Ekman that there are a small set of universal emotional states which can be read directly from the face. Tech companies are now deploying this idea in affect recognition systems, as part of an industry predicted to be worth more than seventeen billion dollars.[35] But there is considerable scientific controversy around emotion detection, which is at best incomplete and at worst misleading. Despite the unstable premise, these tools are being rapidly implemented into hiring, education, and policing systems.

In chapter 6 we look at the ways in which AI systems are used as a tool of state power. The military past and present of artificial intelligence have shaped the practices of surveillance, data extraction, and risk assessment we see today. The deep interconnections between the tech sector and the military are now being reined in to fit a strong nationalist agenda. Meanwhile, extralegal tools used by the intelligence community have now dispersed, moving from the military world into the commercial technology sector, to be used in classrooms, police stations, workplaces, and unemployment offices. The military logics that have shaped AI systems are now part of the workings of municipal government, and they are further skewing the relation between states and subjects.

The concluding chapter assesses how artificial intelligence functions as a structure of power that combines infrastructure, capital, and labor. From the Uber driver being nudged to the undocumented immigrant being tracked to the public housing tenants contending with facial recognition systems in their homes, AI systems are built with the logics of capital, policing, and militarization—and this combination further widens the existing asymmetries of power. These ways of seeing depend on the twin moves of abstraction and extraction: abstracting away the material conditions of their making while extracting more information and resources from those least able to resist.

But these logics can be challenged, just as systems that perpetuate oppression can be rejected. As conditions on Earth change, calls for data protection, labor rights, climate justice, and racial equity should be heard together. When these interconnected movements for justice inform how we understand artificial intelligence, different conceptions of planetary politics become possible.

Extraction, Power, and Politics

Artificial intelligence, then, is an idea, an infrastructure, an industry, a form of exercising power, and a way of seeing; it's also a manifestation of highly organized capital backed by vast systems of extraction and logistics, with supply chains that wrap resource extraction to data protections, racial inequity to climate change. To do that, we need to expand our understanding of what is under way in the empires of AI, to see what is at stake, and to make better collective decisions about what should come next.

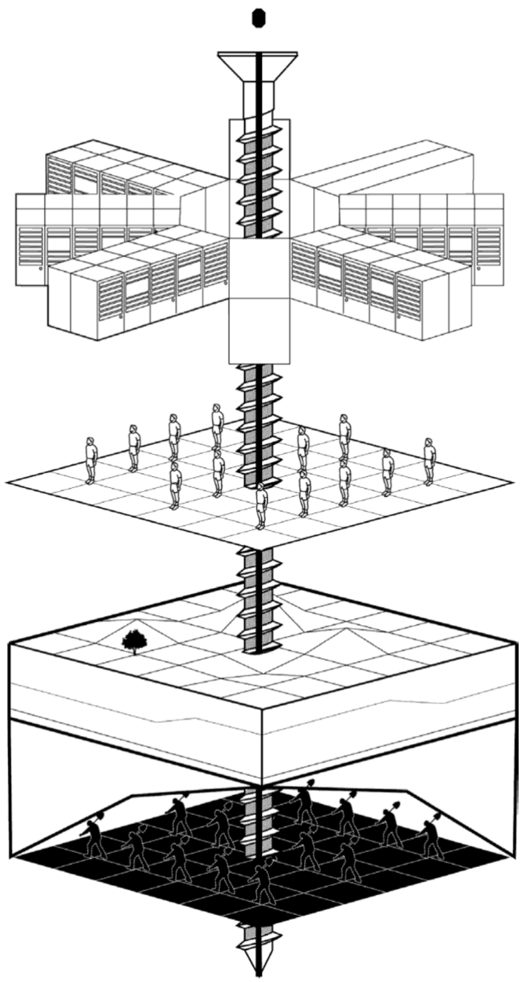

1
Earth

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 Franthe 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 cenfree." 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 clifrom 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 CO_2 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 ing 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.

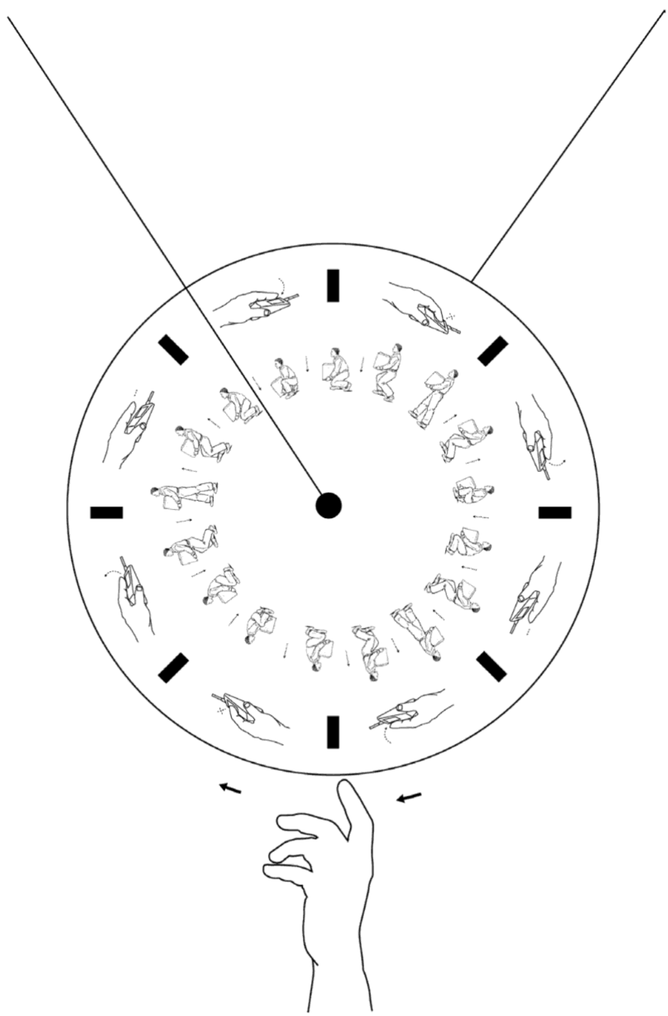

2
Labor

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 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 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 "Amaof 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) 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.

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.

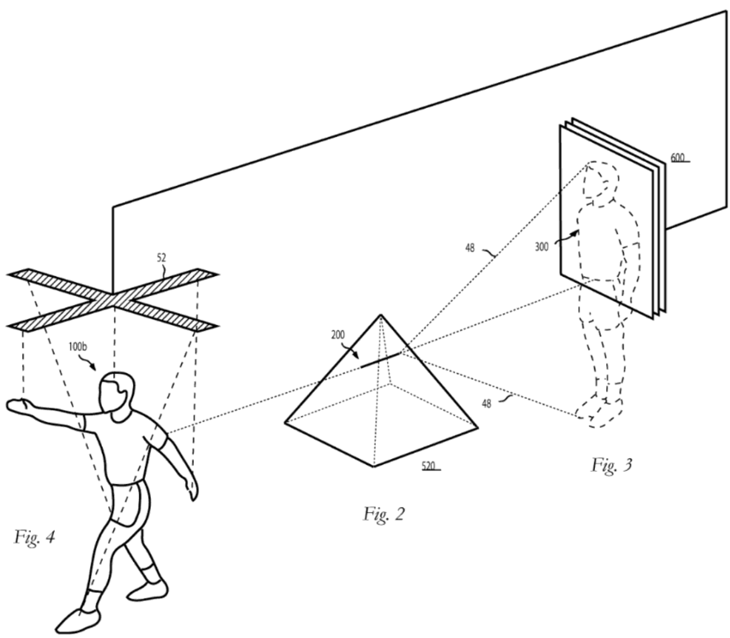

Fig. 4

Fig. 2

Fig. 3

mated fingerprint recognition and has developed methods to assess the quality of fingerprint scanners and imaging systems.[3] After the terrorist attacks of September 11, 2001, NIST became part of the national response to create biometric standards to verify and track people entering the United States.[4] This was a turning point for research on facial recognition; it widened out from a focus on law enforcement to controlling people crossing national borders.[5]

The mug shot images themselves are devastating. Some people have visible wounds, bruising, and black eyes; some are distressed and crying. Others stare blankly back at the camera. Special Dataset 32 contains thousands of photographs of deceased people with multiple arrests, as they endured repeated encounters with the criminal justice system. The people in the mug shot datasets are presented as data points; there are no stories, contexts, or names. Because mug shots are taken at the time of arrest, it's not clear if these people were charged, acquitted, or imprisoned. They are all presented alike.

The inclusion of these images in the NIST database has shifted their meaning from being used to identify individuals in systems of law enforcement to becoming the technical baseline to test commercial and academic AI systems for detecting faces. In his account of police photography, Allan Sekula has argued that mug shots are part of a tradition of technical realism that aimed to "provide a standard physiognomic gauge of the criminal."[6] There are two distinct approaches in the history of the police photograph, Sekula observes. Criminologists like Alphonse Bertillon, who invented the mug shot, saw it as a kind of biographical machine of identification, necessary to spot repeat offenders. On the other hand, Francis Galton, the statistician and founding figure of eugenics, used composite portraiture of prisoners as a way to detect a biologically determined "criminal type."[7] Galton was working within a physiognomist paradigm in which the goal was to find a generalized look that could be used to identify deep character traits from external appearances. When mug shots are used as training data, they function no longer as tools of identification but rather to fine-tune an automated form of vision. We might think of this as Galtonian formalism. They are used to detect the basic mathematical components of faces, to "reduce nature to its geometrical essence."[8]

Mug shots form part of the archive that is used to test facial-recognition algorithms. The faces in the Multiple Encounter Dataset have become standardized images, a technical substrate for comparing algorithmic accuracy. NIST, in collaboration with the Intelligence Advanced Research Projects Activity (IARPA), has run competitions with these mug shots in which researchers compete to see whose algorithm is the fastest and most accurate. Teams strive to beat one another at tasks like verifying the identity of faces or retrieving a face from a frame of surveillance video.[9] The winners celebrate these victories; they can bring fame, job offers, and industry-wide recognition.[10]

Neither the people depicted in the photographs nor their families have any say about how these images are used and likely have no idea that they are part of the test beds of AI. The subjects of the mug shots are rarely considered, and few engineers will ever look at them closely. As the NIST document describes them, they exist purely to "refine tools, techniques, and procedures for face recognition as it supports Next Generation Identification (NGI), forensic comparison, training, analysis, and face image conformance and inter-agency exchange standards."[11] The Multiple Encounter Dataset description observes that many people show signs of enduring violence, such as scars, bruises, and bandages. But the document concludes that these signs are "difficult to interpret due to the lack of ground truth for comparison with a 'clean' sample."[12] These people are not seen so much as individuals but as part of a shared technical resource—just another data component of the Facial Recognition Verification Testing program, the gold standard for the field.

I've looked at hundreds of datasets over years of research into how AI systems are built, but the NIST mug shot databases are particularly disturbing because they represent the model of what was to come. It's not just the overwhelming pathos of the images themselves. Nor is it solely the invasion of privacy they represent, since suspects and prisoners have no right to refuse being photographed. It's that the NIST databases foreshadow the emergence of a logic that has now thoroughly pervaded the tech sector: the unswerving belief that everything is data and is there for the taking. It doesn't matter where a photograph was taken or whether it reflects a moment of vulnerability or pain or if it represents a form of shaming the subject. It has become so normalized across the industry to take and use whatever is available that few stop to question the underlying politics.

Mug shots, in this sense, are the urtext of the current approach to making AI. The context—and exertion of power—that these images represent is considered irrelevant because they no longer exist as distinct things unto themselves. They are not seen to carry meanings or ethical weight as images of individual people or as representations of structural power in the carceral system. The personal, the social, and the political meanings are all imagined to be neutralized. I argue this represents a shift from *image* to *infrastructure,* where the meaning or care that might be given to the image of an individual person, or the context behind a scene, is presumed to be erased at the moment it becomes part of an aggregate mass that will drive a broader system. It is all treated as data to be run through functions, material to be ingested to improve techni-