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 fundamental ways in which humans are embodied, relational, and set within wider ecologies. The second myth is that intelligence is something that exists independently, as though it were natural and distinct from social, cultural, historical, and political forces. In fact, the concept of intelligence has done inordinate harm over centuries and has been used to justify relations of domination from slavery to eugenics.[7]

These mythologies are particularly strong in the field of artificial intelligence, where the belief that human intelligence can be formalized and reproduced by machines has been axiomatic since the mid-twentieth century. Just as Hans's intelligence was considered to be like that of a human, fostered carefully like a child in elementary school, so AI systems have repeatedly been described as simple but humanlike forms of intelligence. In 1950, Alan Turing predicted that "at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted."[8] The mathematician John von Neumann claimed in 1958 that the human nervous system is "prima facie digital."[9] MIT professor Marvin Minsky once responded to the question of whether machines could think by saying, "Of course machines can think; we can think and we are 'meat machines.'"[10] But not everyone was convinced. Joseph Weizenbaum, early AI inventor and creator of the first chatbot program, known as ELIZA, believed that the idea of humans as mere information processing systems is far too simplistic a notion of intelligence and that it drove the "perverse grand fantasy" that AI scientists could create a machine that learns "as a child does."[11]

This has been one of the core disputes in the history of artificial intelligence. In 1961, MIT hosted a landmark lecture series titled "Management and the Computer of the Future." A stellar lineup of computer scientists participated, including

Grace Hopper, J. C. R. Licklider, Marvin Minsky, Allen Newell, Herbert Simon, and Norbert Wiener, to discuss the rapid advances being made in digital computing. At its conclusion, John McCarthy boldly argued that the differences between human and machine tasks were illusory. There were simply some complicated human tasks that would take more time to be formalized and solved by machines.[12]

But philosophy professor Hubert Dreyfus argued back, concerned that the assembled engineers "do not even consider the possibility that the brain might process information in an entirely different way than a computer."[13] In his later work *What Computers Can't Do,* Dreyfus pointed out that human intelligence and expertise rely heavily on many unconscious and subconscious processes, while computers require all processes and data to be explicit and formalized.[14] As a result, less formal aspects of intelligence must be abstracted, eliminated, or approximated for computers, leaving them unable to process information about situations as humans do.

Much in AI has changed since the 1960s, including a shift from symbolic systems to the more recent wave of hype about machine learning techniques. In many ways, the early fights over what AI can do have been forgotten and the skepticism has melted away. Since the mid-2000s, AI has rapidly expanded as a field in academia and as an industry. Now a small number of powerful technology corporations deploy AI systems at a planetary scale, and their systems are once again hailed as comparable or even superior to human intelligence.

Yet the story of Clever Hans also reminds us how narrowly we consider or recognize intelligence. Hans was taught to mimic tasks within a very constrained range: add, subtract, and spell words. This reflects a limited perspective of what horses or humans can do. Hans was already performing remarkable feats of interspecies communication, public performance, 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 well-visited and lesser-known landscapes of computation: the pits of mines, the long corridors of energy-devouring data centers, skull archives, image databases, and the fluorescent-lit hangars of delivery warehouses. These sites are included not just to illustrate the material construction of AI and its ideologies but also to "illuminate the unavoidably subjective and political aspects of mapping, and to provide alternatives to hegemonic, authoritative—and often naturalized and reified—approaches," as media scholar Shannon Mattern writes.[27]

Models for understanding and holding systems accountable have long rested on ideals of transparency. As I've written with the media scholar Mike Ananny, being able to *see* a system is sometimes equated with being able to know how it works and how to govern it.[28] But this tendency has serious limitations. In the case of AI, there is no singular black box to open, no secret to expose, but a multitude of interlaced systems of power. Complete transparency, then, is an impossible goal. Rather, we gain a better understanding of AI's role in the world by engaging with its material architectures, contextual environments, and prevailing politics and by tracing how they are connected.

My thinking in this book has been informed by the disciplines of science and technology studies, law, and political philosophy and from my experience working in both academia and an industrial AI research lab for almost a decade. Over those years, many generous colleagues and communities have changed the way I see the world: mapping is always a collective exercise, and this is no exception.[29] I'm grateful to the scholars who created new ways to understand sociotechnical systems, including Geoffrey Bowker, Benjamin Bratton, Wendy Chun, Lorraine Daston, Peter Galison, Ian Hacking, Stuart Hall, Donald MacKenzie, Achille Mbembé, Alondra Nelson, Susan Leigh Star, and Lucy Suchman, among many others. This book benefited from many in-person conversations and reading the recent work by authors studying the politics of technology, including Mark Andrejevic, Ruha Benjamin, Meredith Broussard, Simone Browne, Julie Cohen, Sasha Costanza-Chock, Virginia Eubanks, Tarleton Gillespie, Mar Hicks, Tung-Hui Hu, Yuk Hui, Safiya Umoja Noble, and Astra Taylor.

As with any book, this one emerges from a specific lived experience that imposes limitations. As someone who has lived and worked in the United States for the past decade, my focus skews toward the AI industry in Western centers of power. But my aim is not to create a complete global atlas—the very idea invokes capture and colonial control. Instead, any author's view can be only partial, based on local observations and interpretations, in what environmental geographer Samantha Saville calls a "humble geography" that acknowledges one's specific perspectives rather than claiming objectivity or mastery.[30]

Just as there are many ways to make an atlas, so there are many possible futures for how AI will be used in the world. The expanding reach of AI systems may seem inevitable, but this is contestable and incomplete. The underlying visions of the AI field do not come into being autonomously but instead have been constructed from a particular set of beliefs and perspectives. The chief designers of the contemporary atlas of AI are a small and homogenous group of people, based in a handful of cities, working in an industry that is currently the wealthiest in the world. Like medieval European *mappae mundi,* which illustrated religious and classical concepts as much as coordinates, the maps made by the AI industry are political interventions, as opposed to neutral reflections of the world. This book is made against the spirit of colonial mapping logics, and it embraces different stories, locations, and knowledge bases to better understand the role of AI in the world.

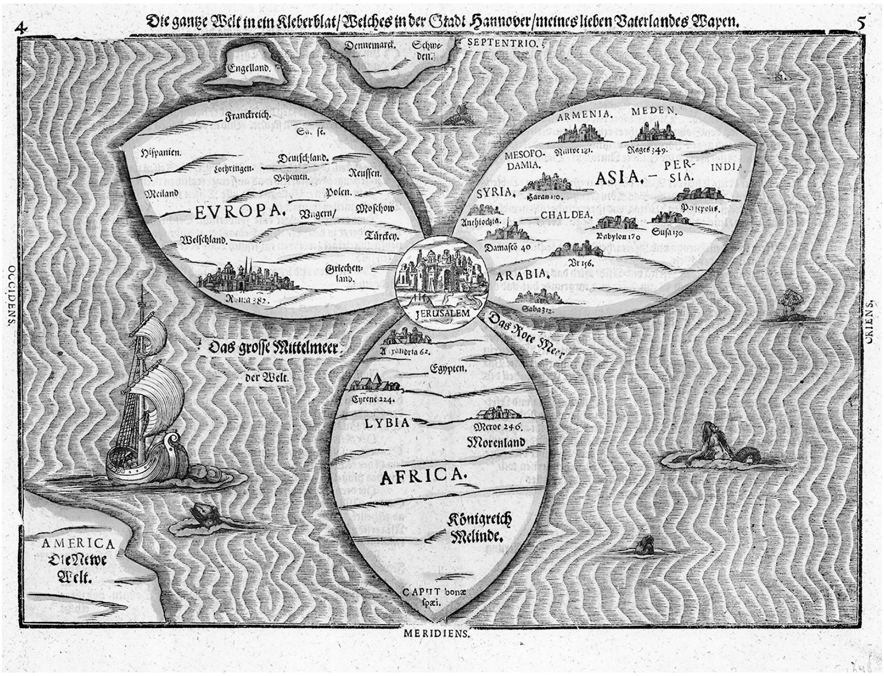

Heinrich Bünting's *mappa mundi,* known as *The Bünting Clover Leaf Map,* which symbolizes the Christian Trinity, with the city of Jerusalem at the center of the world. From *Itinerarium Sacrae Scripturae* (Magdeburg, 1581)

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 around the entire planet. All these things are part of what artificial intelligence is—a two-word phrase onto which is mapped a complex set of expectations, ideologies, desires, and fears.

AI can seem like a spectral force—as disembodied computation—but these systems are anything but abstract. They are physical infrastructures that are reshaping the Earth, while simultaneously shifting how the world is seen and understood.

It's important for us to contend with these many aspects of artificial intelligence—its malleability, its messiness, and its spatial and temporal reach. The promiscuity of AI as a term, its openness to being reconfigured, also means that it can be put to use in a range of ways: it can refer to everything from consumer devices like the Amazon Echo to nameless backend processing systems, from narrow technical papers to the biggest industrial companies in the world. But this has its usefulness, too. The breadth of the term "artificial intelligence" gives us license to consider all these elements and how they are deeply imbricated: from the politics of intelligence to the mass harvesting of data; from the industrial concentration of the tech sector to geopolitical military power; from the deracinated environment to ongoing forms of discrimination.

The task is to remain sensitive to the terrain and to watch the shifting and plastic meanings of the term "artificial intelligence"—like a container into which various things are placed and then removed—because that, too, is part of the story.

Simply put, artificial intelligence is now a player in the shaping of knowledge, communication, and power. These reconfigurations are occurring at the level of epistemology, principles of justice, social organization, political expression, culture, understandings of human bodies, subjectivities, and identities: what we are and what we can be. But we can go further. Artificial intelligence, in the process of remapping and intervening in the world, is politics by other means—although rarely acknowledged as such. These politics are driven by the Great Houses of AI, which consist of the half-dozen or so companies that dominate large-scale planetary computation.

Many social institutions are now influenced by these tools and methods, which shape what they value and how decisions are made while creating a complex series of downstream effects. The intensification of technocratic power has been under way for a long time, but the process has now accelerated. In part this is due to the concentration of industrial capital at a time of economic austerity and outsourcing, including the defunding of social welfare systems and institutions that once acted as a check on market power. This is why we must contend with AI as a political, economic, cultural, and scientific force. As Alondra Nelson, Thuy Linh Tu, and Alicia Headlam Hines observe, "Contests around technology are always linked to larger struggles for economic mobility, political maneuvering, and community building."[36]

We are at a critical juncture, one that requires us to ask hard questions about the way AI is produced and adopted. We need to ask: What is AI? What forms of politics does it propagate? Whose interests does it serve, and who bears the greatest risk of harm? And where should the use of AI be constrained? These questions will not have easy answers. But neither is this an irresolvable situation or a point of no return—dystopian forms of thinking can paralyze us from taking action and prevent urgently needed interventions.[37] As Ursula Franklin writes, "The viability of technology, like democracy, depends in the end on the practice of justice and on the enforcement of limits to power."[38]

This book argues that addressing the foundational problems of AI and planetary computation requires connecting issues of power and justice: from epistemology to labor rights, 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.

6
State

I'm sitting in front of an air-gapped laptop on the tenth floor of a warehouse building in New York. On the screen is a software program normally used for digital forensics, a tool for investigating evidence and validating information held on hard drives. I'm here to research an archive that contains some of the most specific details about how machine learning began to be used in the intelligence sector, as led by some of the wealthiest governments in the world. This is the Snowden archive: all the documents, PowerPoint presentations, internal memos, newsletters, and technical manuals that former NSA contractor and whistleblower Edward Snowden leaked in 2013. Each page is marked with a header noting different forms of classification. TOP SECRET // SI // ORCON // NOFORN.[1] Each is a warning and a designation.

The filmmaker Laura Poitras first gave me access to this archive in 2014. It was overwhelming to read: the archive held well over a decade of intelligence thinking and communication, including internal documents of the National Security Agency in the United States and the Government Communication Headquarters in the United Kingdom, and the international network of the Five Eyes.[2] This knowledge was strictly off-limits to those without high-level clearance. It was part of the "classified empire" of information, once estimated to be growing five times faster than publicly accessible knowledge but now is anyone's guess.[3] The Snowden archive captures the years when the collection of data metastasized: when phones, browsers, social media platforms, and email all became data sources for the state. The documents reveal how the intelligence community contributed to the development of many of the techniques we now refer to as artificial intelligence.

The Snowden archive reveals a parallel AI sector, one developed in secrecy. The methods share many similarities, but there are striking differences in terms of the reach, the objectives, and the result. Gone are any rhetorical constructs justifying extraction and capture: every software system is simply described as something to be owned, to be defeated; all data platforms are fair game, and very little is designated as protected. One NSA PowerPoint deck outlines TREASUREMAP, a program designed to build a near real-time, interactive map of the internet.[4] It claims to track the location and owner of any connected computer, mobile device, or router: "Map the entire internet—any device, anywhere, all the time," the slide boasts. A few slides on "TREASUREMAP as an Enabler" offers up a layer-cake image of signals analysis. Above the geographical layer and the network layer is the "cyber persona layer"—quaintly represented on the slide by jellybean-era iMacs and Nokia feature phones—and then the "persona layer" of personal connections. This is meant to depict all people who use connected devices around the world, in a "300,000-foot view of the internet." It also looks remarkably like the work of social network mapping and manipulation companies like Cambridge Analytica.

The Snowden documents were released in 2013, but they still read like the AI marketing brochures of today. If

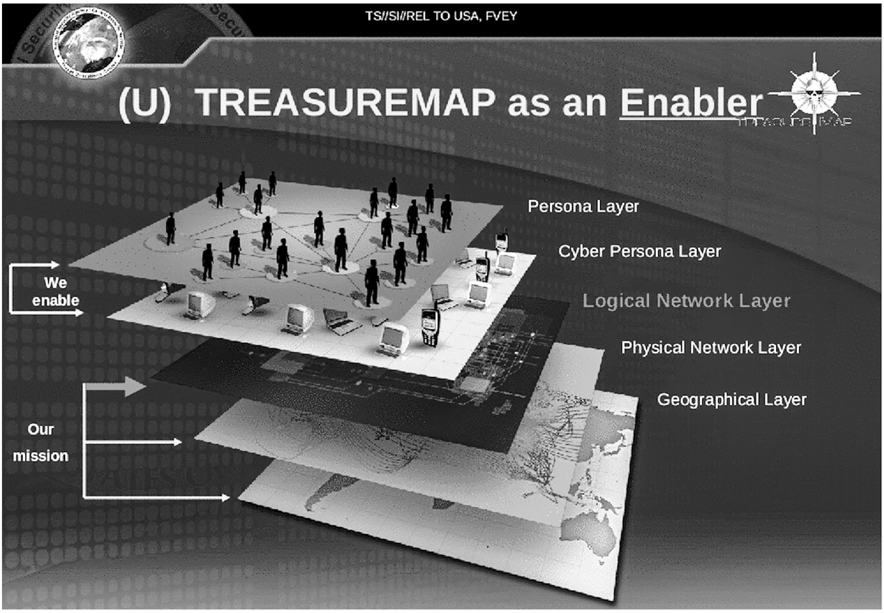

TS//SI//REL TO USA, FVEY

(U) TREASUREMAP as an Enabler

Persona Layer

Cyber Persona Layer

Logical Network Layer

Physical Network Layer

Geographical Layer

We enable

Our mission

TREASUREMAP as an Enabler. Snowden archive

TREASUREMAP was a precursor to Facebook's God's-eye network view, then the program called FOXACID is reminiscent of Amazon Ring for a home computer: recording everyday activity.[5] "If we can get the target to visit us in some sort of browser, we can probably own them," the slide explains.[6] Once individuals have been tempted to click on a spam email or visit a captured website, the NSA drops files through a browser that will permanently live on their device, quietly reporting everything they do back to base. One slide describes how analysts "deploy very targeted emails" that require "a level of guilty knowledge" about the target.[7] The restrictions on the NSA gathering that guilty knowledge (when it comes to data from American citizens, at least) are rarely discussed. One document notes that the agency was working on multiple fronts to "aggressively pursue legal authorities and a policy framework mapped more fully to the information age."[8] In other words, change the laws to fit the tools, not the other way around.

The U.S. intelligence agencies are the old guards of big data. Along with the Defense Advanced Research Projects Agency, they have been major drivers of AI research since the 1950s. As the historian of science Paul Edwards describes in *The Closed World,* military research agencies actively shaped the emerging field that would come to be known as AI from its earliest days.[9] The Office of Naval Research, for example, partly funded the first Summer Research Project on Artificial Intelligence at Dartmouth College in 1956.[10] The field of AI has always been strongly guided by military support and often military priorities, long before it was clear that AI could be practical at scale. As Edwards notes:

> As the project with the least immediate utility and the farthest-reaching ambitions, AI came to rely unusually heavily on ARPA funding. As a result, ARPA became the primary patron for the first twenty years of AI research. Former director Robert Sproull proudly concluded that "a whole generation of computer experts got their start from DARPA funding" and that "all the ideas that are going into the fifth-generation [advanced computing] project [of the mid-1980s]—artificial intelligence, parallel computing, speech understanding, natural-languages programming—ultimately started from DARPA-funded research."[11]

The military priorities of command and control, automation, and surveillance profoundly shaped what AI was to become. The tools and approaches that came out of DARPA funding have marked the field, including computer vision, automatic translation, and autonomous vehicles. But these technical methods have deeper implications. Infused into the overall logics of AI are certain kinds of classificatory thinking—from explicitly battlefield-oriented notions such as target, asset, and anomaly detection to subtler categories of high, medium, and low risk. Concepts of constant situational awareness and targeting would drive AI research for decades, creating epistemological frameworks that would inform both industry and academia.

From the point of view of the state, the turn to big data and machine learning expanded the modes of information extraction and informed a social theory of how people can be tracked and understood: *you shall know them by their metadata.* Who is texted, which locations are visited, what is read, when devices spring into action and for what reason—these molecular actions became a vision of threat identification and assessment, guilt or innocence. Harvesting and measuring large aggregates of data at a distance became the preferred way to develop alleged insights into groups and communities as well as assessments of potential targets for killing. The NSA and GCHQ are not unique—China, Russia, Israel, Syria, and many other countries have similar agencies. There are many systems of sovereign surveillance and control, a multitude of war machines that never wind down. The Snowden archive underscores how state and corporate actors collaborate in order to produce what Achille Mbembe calls "infrastructural warfare."[12]

But the relationship between national militaries and the AI industry has expanded beyond security contexts. Technologies once only available to intelligence agencies—that were *extralegal* by design—have filtered down to the state's municipal arms: government and law enforcement agencies. While the NSA has been a focus for privacy concerns, less attention is given to the growing commercial surveillance sector, which aggressively markets its tools and platforms to police departments and public agencies. The AI industry is simultaneously challenging and reshaping the traditional role of states while also being used to shore up and expand older forms of geopolitical power. Algorithmic governance is both part of and exceeds traditional state governance. To paraphrase the theorist Benjamin Bratton, the state is taking on the armature of a machine because the machines have already taken on the roles and register of the state.[13]

Making the Third Offset

The story of the internet's creation has been centered around U.S. military and academic innovation and dominance.[14] But in the space of AI, we see that there is no pure national system. Instead, AI systems operate within a complex interwoven network of multinational and multilateral tools, infrastructures, and labor. Take, for example, a facial recognition system that was rolled out in the streets of Belgrade.[15] The director of police ordered the installation of two thousand cameras in eight hundred locations around the city to capture faces and license plates. The Serbian government signed an agreement with Chinese telecommunications giant Huawei to provide the video surveillance, 4G network support, and unified data and command centers. Such deals are common. Local systems are often hybrids, with infrastructure from China, India, the United States, and elsewhere, with porous boundaries, different security protocols, and potential data backdoors.

But the rhetoric around artificial intelligence is much starker: we are repeatedly told that we are in an AI war. The dominant objects of concern are the supernational efforts of the United States and China, with regular reminders that China has stated its commitment to be the global leader in AI.[16] The data practices of China's leading tech companies, including Alibaba, Huawei, Tencent, and ByteDance, are often framed as direct Chinese state policy and thus seen as inherently more threatening than U.S. private actors such as Amazon and Facebook, even though the lines between state and corporate imperatives and incentives are complexly intertwined. Yet the language of war is more than just the usual articulation of xenophobia, mutual suspicion, international espionage, and network hacking. As media scholars such as Wendy Chun and Tung-Hui Hu have noted, the liberal vision of global digital citizens engaging as equals in the abstract space of networks has shifted toward a paranoid vision of defending a national cloud against the racialized enemy.[17] The specter of the foreign threat works to assert a kind of sovereign power over AI and to redraw the locus of power of tech companies (which are transnational in infrastructure and influence) back within the bounds of the nation-state.

Yet the nationalized race for technological superiority is both rhetorical and real at the same time, creating the dynamics for geopolitical competition across and within commercial and military sectors, increasingly blurring the lines between the two. The dual use of AI applications in both civilian and military domains has also produced strong incentives for close collaboration and funding.[18] In the United States, we can see how this became an explicit strategy: to seek national control and international dominance of AI in order to secure military and corporate advantage.

The latest iteration of this strategy emerged under Ash Carter, who served as U.S. secretary of defense from 2015 to 2017. Carter played a significant role in bringing Silicon Valley into closer relationship to the military, convincing tech companies that national security and foreign policy depended on American dominance of AI.[19] He called this the Third Offset strategy. An offset is generally understood as a way of compensating for an underlying military disadvantage by changing the conditions, or as former secretary of defense Harold Brown stated in 1981, "Technology can be a force multiplier, a resource that can be used to help offset numerical advantages of an adversary. Superior technology is one very effective way to balance military capabilities other than by matching an adversary tank-for-tank or soldier-for-soldier."[20]

The First Offset is commonly understood as the use of nuclear weapons in the 1950s.[21] The Second was the expansion of covert, logistical, and conventional weapons in the 1970s and 1980s. The Third, according to Carter, should be a combination of AI, computational warfare, and robots.[22] But unlike the NSA, which already had robust surveillance capabilities, the U.S. military lacked the AI resources, expertise, and infrastructure of America's leading technology companies.[23] In 2014, Deputy Defense Secretary Robert Work outlined the Third Offset as an attempt to "exploit all the advances in artificial intelligence and autonomy."[24]

To build AI war machines, the Department of Defense would need gigantic extractive infrastructures. Yet in order to gain access to highly paid engineering labor and sophisticated development platforms, partnering with industry was necessary. The NSA had paved the way with systems like PRISM, both working with and secretly infiltrating telecommunications and technology companies.[25] But these more covert approaches faced renewed political pushback after the Snowden disclosures. Congress passed the USA Freedom Act in 2015, which introduced some limitations on the NSA's access to real-time data from Silicon Valley. Yet the possibility for a larger military-industrial complex around data and AI remained tantalizingly close. Silicon Valley had already built and monetized the logics and infrastructures of AI required to drive a new offset. But first the tech sector had to be convinced that partnering on creating the infrastructure of warfare would be worth it without alienating their employees and deepening public mistrust.

Enter Project Maven

In April 2017, the Department of Defense published a memo announcing the Algorithmic Warfare Cross-Functional Team, code-named Project Maven.[26] "The Department of Defense must integrate artificial intelligence and machine learning more effectively across operations to maintain advantages over increasingly capable adversaries and competitors," wrote the deputy defense secretary.[27] The goal of the program was to get the best possible algorithmic systems into the battlefield quickly, even when they were just 80 percent complete.[28] It was part of a much bigger plan, the Joint Enterprise Defense Infrastructure cloud project—or JEDI—an enormous redesign of the entire IT infrastructure of the Defense Department, from the Pentagon to field-level support. Project Maven was a small piece of this larger picture, and the aim was to create an AI system that would allow analysts to select a target and then see every existing clip of drone footage that featured the same person or vehicle.[29] Ultimately, the Defense Department wanted an automated search engine of drone videos to detect and track enemy combatants.

The technical platforms and machine learning skills needed for Project Maven were centered in the commercial tech sector. The Defense Department decided to pay tech companies to analyze military data collected from satellites and battlefield drones in places where U.S. domestic privacy laws did not apply. This would align military and U.S. tech sector financial interests around AI without directly triggering constitutional privacy tripwires, as the National Security Agency

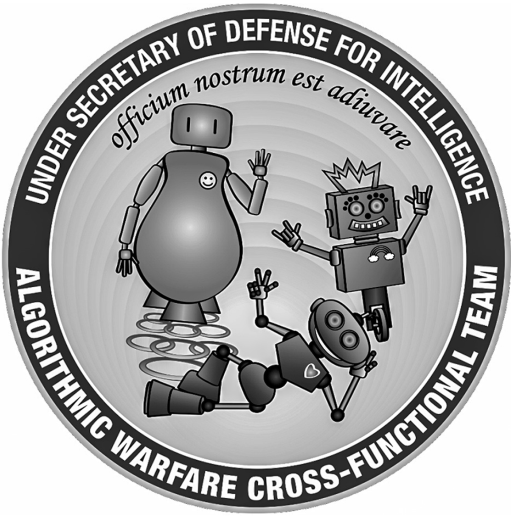
UNDER SECRETARY OF DEFENSE FOR INTELLIGENCE
officium nostrum est adiuvare
ALGORITHMIC WARFARE CROSS-FUNCTIONAL TEAM

The official seal of the Algorithmic Warfare Cross-Functional Team, code-named Project Maven. The Latin motto translates as "Our job is to help." Produced by U.S. Department of Defense

had done. A bidding war began among the technology companies that wanted the Maven contract, including Amazon, Microsoft, and Google.

The first Project Maven contract went to Google. Under the agreement, the Pentagon would use Google's TensorFlow AI infrastructure to comb through drone footage and detect objects and individuals as they moved between locations.[30] Fei-Fei Li, then chief scientist of AI/ML at Google, was already an expert in building object recognition datasets, given her experience creating ImageNet and using satellite data to detect and analyze cars.[31] But she was adamant that the project should be kept secret. "Avoid at ALL COSTS any mention or implication of AI," Li wrote in an email to Google colleagues that was later leaked. "Weaponized AI is probably one of the most sensitized topics of AI—if not THE most. This is red meat to the media to find all ways to damage Google."[32]

But in 2018, Google employees discovered the extent of the company's role in the project. They were furious that their work was being used for warfare purposes, especially after it became known that Project Maven's image identification goals included objects such as vehicles, buildings, and humans.[33] More than 3,100 employees signed a letter of protest stating that Google should not be in the business of war and demanded that the contract be canceled.[34] Under increasing pressure, Google officially ended its work on Project Maven and withdrew from the competition for the Pentagon's ten-billion-dollar JEDI contract. In October that year, Microsoft's president, Brad Smith, announced in a blog post that "we believe in the strong defense of the United States and we want the people who defend it to have access to the nation's best technology, including from Microsoft."[35] The contract ultimately went to Microsoft, which outbid Amazon.[36]

Shortly after the internal uprising, Google released its Artificial Intelligence Principles, which included a section on "AI applications we will not pursue."[37] These included making "weapons or other technologies whose principal purpose or implementation is to cause or directly facilitate injury to people," as well as "technologies that gather or use information for surveillance violating internationally accepted norms."[38] While the turn to AI ethics quelled some internal and external concerns, the enforceability and parameters of ethical restraint were left unclear.[39]

In response, former Google CEO Eric Schmidt characterized the pushback over Project Maven as "a general concern in the tech community of somehow the military-industrial complex using our stuff to kill people incorrectly, if you will."[40] This shift, from the debate over whether to use AI in warfare at all to a debate over whether AI could help to "kill people correctly," was quite strategic.[41] It moved the focus away from the foundational ethics of AI as a military technology toward questions of precision and technical accuracy. But Lucy Suchman argues that the problems with automated warfare go far beyond whether the killing was accurate or "correct."[42] Particularly in the case of object detection, Suchman asks, who is building the training sets and using what data, and how are things labeled as an imminent threat? What kinds of classificatory taxonomies are used to decide what constitutes sufficiently abnormal activity to trigger a legal drone attack? And why should we condone attaching life or death consequences to these unstable and inherently political classifications?[43]

The Maven episode, as well as the AI principles that emerge, points to the deep schisms in the AI industry about the relationship between the military and civilian spheres. The AI war, both real and imagined, instills a politics of fear and insecurity that creates a climate that is used to stifle internal dissent and promote unquestioning support for a nationalist agenda.[44] After the fallout from Maven faded, Google's chief legal officer, Kent Walker, said that the company was pursuing higher security certifications in order to work more closely with the Defense Department. "I want to be clear," he said. "We are a proud American company."[45] Articulating patriotism as policy, tech companies are increasingly expressing strong alignment with the interests of the nation-state, even as their platforms and capacities exceed traditional state governance.

The Outsourced State

The relationship between the state and the AI industry goes well beyond national militaries. The technologies once reserved for war zones and espionage are now used at the local level of government, from welfare agencies to law enforcement. This shift has been propelled by outsourcing key functions of the state to technology contractors. On the surface, this does not seem very different than the usual outsourcing of government functions to the private sector through companies such as Lockheed Martin or Halliburton. But now militarized forms of pattern detection and threat assessment are moving at scale into municipal-level services and institutions.[46] A significant example of this phenomenon is the company named after the magical seeing stones in *Lord of the Rings:* Palantir.

Palantir was established in 2004, cofounded by PayPal billionaire Peter Thiel, who was also an adviser and financial supporter of President Trump. Thiel would later argue in an opinion piece that AI is first and foremost a military technology: "Forget the sci-fi fantasy; what is powerful about actually existing AI is its application to relatively mundane tasks like computer vision and data analysis. Though less uncanny than Frankenstein's monster, these tools are nevertheless valuable to any army—to gain an intelligence advantage, for example. . . . No doubt machine learning tools have civilian uses, too."[47]

While Thiel recognizes the nonmilitary uses of machine learning, he particularly believes in the *in-between space:* where commercial companies produce military-styled tools to be provided to anyone who would like to gain an intelligence advantage and is willing to pay for it. Both he and Palantir's CEO, Alex Karp, describe Palantir as "patriotic," with Karp accusing other technology companies that refuse to work with the military agencies as "borderline craven."[48] In an insightful essay, the writer Moira Weigel studied Karp's university dissertation, which reveals his early intellectual interest in aggression and a belief that "the desire to commit violence is a constant founding fact of human life."[49] Karp's thesis was titled "Aggression in the Life World."

Palantir's original clients were federal military and intelligence agencies, including the Defense Department, National Security Agency, FBI, and CIA.[50] As revealed in an investigation by Mijente, after Trump took the presidency, Palantir's contracts with U.S. agencies totaled more than a billion dollars.[51] But Palantir did not style itself as a typical defense contractor in the mold of Lockheed Martin. It adopted the character of the Silicon Valley start-up, based in Palo Alto and predominantly staffed by young engineers, and it was backed by In-Q-Tel, the venture capital firm funded by the CIA. Beyond its initial intelligence agency clients, Palantir began to work with hedge funds, banks, and corporations like Walmart.[52] But its DNA was shaped working for, and within, the defense community. It deployed the same approaches seen in the Snowden documents, including extracting data across devices and infiltrating networks in order to track and evaluate people and assets. Palantir quickly became a preferred outsourced surveillance provider, including designing the databases and management software to drive the mechanics of deportation for Immigration and Customs Enforcement (ICE).[53]

Palantir's business model is based on a mix of data analysis and pattern detection using machine learning, combined with more generic consulting. Palantir sends engineers into a company, who extract a wide variety of data—emails, call logs, social media, when employees enter and leave buildings, when they book plane tickets, everything the company is prepared to share—then look for patterns and give advice on what to do next. One common approach is to search for current or potential so-called bad actors, disgruntled employees who may leak information or defraud the company. The underlying worldview built into Palantir's tools is reminiscent of the NSA: collect everything, then look for anomalies in the data. However, while the NSA's tools are built to surveil and target enemies of the state, in either conventional or covert warfare, Palantir's approach has been directed against civilians. As described in a major investigation by Bloomberg in 2018, Palantir is "an intelligence platform designed for the global War on Terror" that is now "weaponized against ordinary Americans at home": "Palantir cut its teeth working for the Pentagon and the CIA in Afghanistan and Iraq. . . . The U.S. Department of Health and Human Services uses Palantir to detect Medicare fraud. The FBI uses it in criminal probes. The Department of Homeland Security deploys it to screen air travelers and keep tabs on immigrants."[54]

Soon, keeping tabs on undocumented workers evolved into capturing and deporting people at schools and places of work. In furtherance of this objective, Palantir produced a phone app called FALCON, which functions as a vast dragnet, gathering data from multiple law enforcement and public databases that list people's immigration histories, family relationships, employment information, and school details. In 2018, ICE agents used FALCON to guide their raid of almost a hundred 7-Elevens across the United States in what was called "the largest operation against a single employer in the Trump era."[55]

Despite Palantir's efforts to maintain secrecy about what it builds or how its systems work, its patent applications give us some insight into the company's approach to AI for deportation. In an application innocuously entitled *Database systems and user interfaces for dynamic and interactive mobile image analysis and identification,* Palantir brags about the app's ability to photograph people in short-time-frame encounters and, regardless of whether they are under suspicion or not, to run their image against all available databases. In essence, the system uses facial recognition and back-end processing to create a framework on which to base any arrest or deportation.

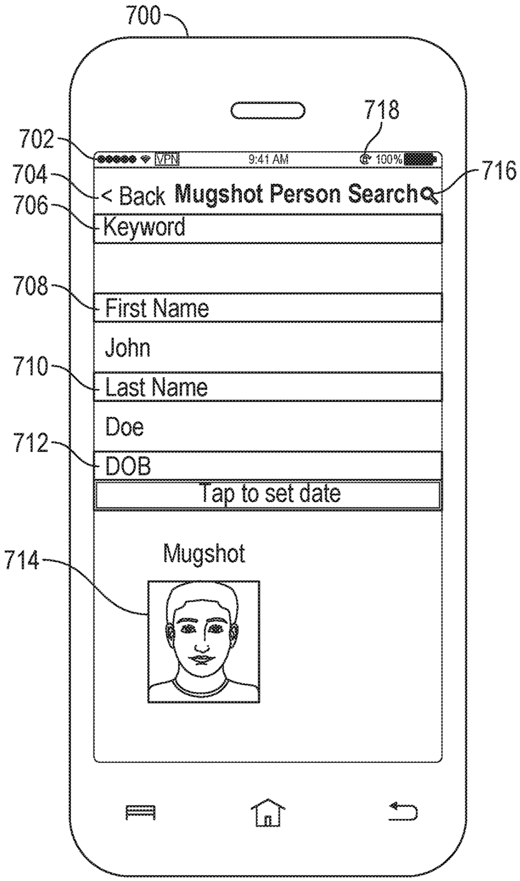

< Back Mugshot Person Search

Keyword

First Name

John

Last Name

Doe

DOB

Tap to set date

Mugshot

An image from Palantir's patent US10339416B2.
Courtesy U.S. Patent and Trademark Office

While Palantir's systems have structural similarities to those at the NSA, they have devolved to a local community level, to be sold to supermarket chains and local law enforcement alike. This represents a shift away from traditional policing toward the goals more associated with military intelligence infrastructures. As law professor Andrew Ferguson explains, "We are moving to a state where prosecutors and police are going to say 'the algorithm told me to do it, so I did, I had no idea what I was doing.' And this will be happening at a widespread level with very little oversight."[56]

The sociologist Sarah Brayne was one of the first scholars to observe directly how Palantir's data platforms are used in situ, specifically by the Los Angeles Police Department. After more than two years of riding along with police on patrols, watching them at their desks, and conducting multiple interviews, Brayne concluded that in some domains these tools merely amplify prior police practices but that in other ways they are transforming the process of surveillance entirely. In short, police are turning into intelligence agents:

> The shift from traditional to big data surveillance is associated with a migration of law enforcement operations toward intelligence activities. The basic distinction between law enforcement and intelligence is as follows: law enforcement typically becomes involved once a criminal incident has occurred. Legally, the police cannot undertake a search and gather personal information until there is probable cause. Intelligence, by contrast, is fundamentally predictive. Intelligence activities involve gathering data; identifying suspicious patterns, locations, activity, and individuals; and preemptively intervening based on the intelligence acquired.[57]

Although everyone is subject to these types of surveillance, some people are more likely to be subjected to it than others: immigrants, the undocumented, the poor, and communities of color. As Brayne observed in her study, the use of Palantir's software reproduces inequality, making those in predominantly poor, Black, and Latinx neighborhoods subject to even greater surveillance. Palantir's point system lends an aura of objectivity: it's "just math," in the words of one police officer. But it creates a reinforcing loop of logic.[58] Brayne writes:

> Despite the stated intent of the point system to avoid legally contestable bias in police practices, it hides both intentional and unintentional bias in policing and creates a self-perpetuating cycle: if individuals have a high point value, they are under heightened surveillance and therefore have a greater likelihood of being stopped, further increasing their point value. Such practices hinder the ability of individuals already in the criminal justice system from being further drawn into the surveillance net, while obscuring the role of enforcement in shaping risk scores.[59]

The machine learning approaches of Palantir and its ilk can lead to a feedback loop, where those included in a criminal justice database are more likely to be surveilled and thus more likely to have more information about them included, which justifies further police scrutiny.[60] Inequity is not only deepened but tech-washed, justified by the systems that appear immune to error yet are, in fact, intensifying the problems of overpolicing and racially biased surveillance.[61] The intelligence models that began in national government agencies have now become part of the policing of local neighborhoods. The NSA-ification of police departments exacerbates historical inequality and radically transforms and expands the practices of police work.

Despite the massive expansion of government contracts for AI systems, little attention has been given to the question of whether private vendors of these technologies should be legally accountable for the harms produced when governments use their systems. Given how often governments are turning to contractors to provide the algorithmic architectures for state decision-making, be it policing or welfare systems, there is a case that technology contractors like Palantir should be liable for discrimination and other violations. Currently most states attempt to disclaim any responsibility for problems created by the AI systems they procure, with the argument that "we cannot be responsible for something we don't understand." This means that commercial algorithmic systems are contributing to the process of government decision making without meaningful mechanisms of accountability. With the legal scholar Jason Schultz, I've argued that developers of AI systems that directly influence government decisions should be found to be state actors for purposes of constitutional liability in certain contexts.[62] That is, they could be found legally liable for harms in the same way that states can be. Until then, vendors and contractors have little incentive to ensure that their systems aren't reinforcing historical harms or creating entirely new ones.[63]

Another example of this phenomenon is Vigilant Solutions, established in 2005. The company works on the basis of a single premise: take surveillance tools that might require judicial oversight if operated by governments and turn them into a thriving private enterprise outside constitutional privacy limits. Vigilant began its venture in multiple cities across the United States by installing automatic license-plate recognition (ALPR) cameras, placing them everywhere from cars to light poles, parking lots to apartment buildings. This array of networked cameras photographs every passing car, storing license plate images in a massive perpetual database. Vigilant then sells access to that database to the police, private investigators, banks, insurance companies, and others who want access to it. If police officers want to track a car across the entire state and mark every place it has been, Vigilant can show them. Likewise, if a bank wanted to repossess a car, Vigilant could reveal where it was, for a price.

California-based Vigilant markets itself as "one of those trusted crime fighting tools to help law enforcement develop leads and solve crimes faster," and it has partnered with a range of governments in Texas, California, and Georgia to provide their police with a suite of ALPR systems to use on patrol, along with access to Vigilant's database.[64] In return, the local governments provide Vigilant with records of outstanding arrest warrants and overdue court fees. Any license plates flagged to match those associated with outstanding fines in the database are fed into police officers' mobile systems, altering them to pull these drivers over. Drivers are then given two options: pay the outstanding fine on the spot or be arrested. On top of taking a 25 percent surcharge, Vigilant keeps records of every license plate reading, extracting that data to add to its massive databases.

Vigilant signed a significant contract with ICE that gave the agency access to five billion records of license plates gathered by private businesses, as well as 1.5 billion data points contributed by eighty local law enforcement agencies across the United States—including information on where people live and work. That data can stem from informal arrangements between local police and ICE and may already violate state data-sharing laws. ICE's own privacy policy limits data collection near "sensitive locations" like schools, churches, and protests. But in this case, ICE doesn't collect the data or maintain the database—the agency simply buys access to Vigilant's systems, which has far fewer restrictions. This is a de facto privatization of public surveillance, a blurring between private contractors and state entities, and it creates opaque forms of data harvesting that live outside of traditional protective guidelines.[65]

Vigilant has since expanded its "crime-solving" toolkit beyond license plate readers to include ones that claim to recognize faces. In doing so, Vigilant seeks to render human faces as the equivalent of license plates and then feed them back into the policing ecology.[66] Like a network of private detectives, Vigilant creates a God's-eye view of America's interlaced roads and highways, along with everyone who travels along them, while remaining beyond any meaningful form of regulation or accountability.[67]

If we move from the police cruiser to the front porch, we see yet another location where the differences between public and private sector data practices are eroding. A new generation of social media crime-reporting apps like Neighbors, Citizen, and Nextdoor allow users to get alerts about local incidents reported in real time, then discuss them, as well as broadcast, share, and tag security camera footage. Neighbors, which is made by Amazon and relies on its Ring doorbell cameras, defines itself as the "new neighborhood watch" and classifies footage into categories like Crime, Suspicious, or Stranger. Videos are often shared with police.[68] In these residential surveillance ecosystems, the logics of TREASUREMAP and FOXACID conjoin, but connected to the home, the street, and every place in between.

For Amazon, each new Ring device sold helps build yet more large-scale training datasets inside and outside the home, with classificatory logics of normal and anomalous behavior aligned with the battlefield logics of allies and enemies. One example is a feature where users can report stolen Amazon packages. According to one journalistic investigation, many of the posts featured racist commentary, and video posts disproportionately depicted people of color as potential thieves.[69] Beyond reporting crime, Ring is also used to report Amazon employees who are seen as underperforming, such as being insufficiently careful with packages—creating a new layer of worker surveillance and retribution.[70]

To complete its public-private infrastructure of surveillance, Amazon has been aggressively marketing the Ring system to police departments, giving them discounts and offering a portal that allows police to see where Ring cameras are located in the local area and to contact homeowners directly to request footage informally without a warrant.[71] Amazon has negotiated Ring video-sharing partnerships with more than six hundred police departments.[72]

In one case, Amazon negotiated a memorandum of understanding with a police department in Florida, discovered through a public records request filed by journalist Caroline Haskins, which showed that police were incentivized to promote the Neighbors app and for every qualifying download they would receive credits toward free Ring cameras.[73] The result was a "self-perpetuating surveillance network: more people download Neighbors, more people get Ring, surveillance footage proliferates, and police can request whatever they want," Haskins writes.[74] Surveillance capacities that were once ruled over by courts are now on offer in Apple's App Store and promoted by local street cops. As media scholar Tung-Hui Hu observes, by using such apps, we "become freelancers for the state's security apparatus."[75]

Hu describes how targeting—a quintessential militaristic term—in all its forms should be considered together as one interconnected system of power—from targeted advertising to targeting suspicious neighbors to targeting drones. "We cannot merely consider one form of targeting in isolation from the other; conjoined in the sovereignty of data, they call on us to understand power in the age of the cloud differently."[76] The ways of seeing that were once the sole province of intelligence agencies have been granulated and dispersed throughout many social systems—embedded in workplaces, homes, and cars—and promoted by technology companies that live in the cross-hatched spaces that overlap the commercial and military AI sectors.

From Terrorist Credit Scores to Social Credit Scores

Underlying the military logics of targeting is the idea of the *signature.* Toward the end of President George W. Bush's second term, the CIA argued that it should be able to launch drone attacks based solely on an individual's observed "pattern of behavior" or "signature."[77] Whereas a "personality strike" involves targeting a specific individual, a "signature strike" is when a person is killed due to their metadata signature; in other words, their identity is not known but data suggests that they might be a terrorist.[78] As the Snowden documents showed, during the Obama years, the National Security Agency's global metadata surveillance program would geolocate a SIM card or handset of a suspect, and then the U.S. military would conduct drone strikes to kill the individual in possession of the device.[79] "We kill people based on metadata," said General Michael Hayden, former director of the NSA and the CIA.[80] The NSA's Geo Cell division was reported to use more colorful language: "We track 'em, you whack 'em."[81]

Signature strikes may sound precise and authorized, implying a true mark of someone's identity. But in 2014, the legal organization Reprieve published a report showing that drone strikes attempting to kill 41 individuals resulted in the deaths of an estimated 1,147 people. "Drone strikes have been sold to the American public on the claim that they are 'precise.' But they are only as precise as the intelligence that feeds them," said Jennifer Gibson, who led the report.[82] But the form of the signature strike is not about precision: it is about correlation. Once a pattern is found in the data and it reaches a certain threshold, the suspicion becomes enough to take action even in the absence of definitive proof. This mode of adjudication by pattern recognition is found in many domains—most often taking the form of a score.

Consider an example from the 2015 Syrian refugee crisis. Millions of people were fleeing widespread civil war and enemy occupation in hopes of finding asylum in Europe. Refugees were risking their lives on rafts and overcrowded boats. On September 2, a three-year-old boy named Alan Kurdi drowned in the Mediterranean Sea, alongside his five-year-old brother, when their boat capsized. A photograph showing his body washed up on a beach in Turkey made international headlines as a potent symbol for the extent of the humanitarian crisis: one image standing in for the aggregate horror. But some saw this as a growing threat. It is around this time that IBM was approached about a new project. Could the company use its machine learning platform to detect the data signature of refugees who might be connected to jihadism? In short, could IBM automatically distinguish a terrorist from a refugee?

Andrew Borene, a strategic initiatives executive at IBM, described the rationale behind the program to the military publication *Defense One:* "Our worldwide team, some of the folks in Europe, were getting feedback that there were some concerns that within these asylum-seeking populations that had been starved and dejected, there were fighting-age males coming off of boats that looked awfully healthy. Was that a cause for concern in regard to ISIS and, if so, could this type of solution be helpful?"[83]

From the safe distance of their corporate offices, IBM's data scientists viewed the problem as one best addressed through data extraction and social media analysis. Setting aside the many variables that existed in the conditions of makeshift refugee camps and the dozens of assumptions used to classify terrorist behavior, IBM created an experimental "terrorist credit score" to weed out ISIS fighters from refugees. Analysts harvested a miscellany of unstructured data, from Twitter to the official list of those who had drowned alongside the many capsized boats off the shores of Greece and Turkey. They also made up a data set, modeled on the types of metadata available to border guards. From these disparate measures, they developed a hypothetical threat score: not an absolute indicator of guilt or innocence, they pointed out, but a deep "insight" into the individual, including past addresses, workplaces, and social connections.[84] Meanwhile, Syrian refugees had no knowledge that their personal data was being harvested to trial a system that might single them out as potential terrorists.

This is just one of many cases where new technical systems of state control use the bodies of refugees as test cases. These military and policing logics are now suffused with a form of financialization: socially constructed models of creditworthiness have entered into many AI systems, influencing everything from the ability to get a loan to permission to cross borders. Hundreds of such platforms are now in use around the world, from China to Venezuela to the United States, rewarding predetermined forms of social behavior and penalizing those who do not conform.[85] This "new regime of moralized social classification," in the words of sociologists Marion Fourcade and Kieran Healy, benefits the "high achievers" of the traditional economy while further disadvantaging the least privileged populations.[86] Credit scoring, in the broadest sense, has become a place where the military and commercial signatures combine.

This AI scoring logic is deeply entwined in law enforcement and border control, traditional domains of the state, but it also informs another state function: access to public benefits. As the political scientist Virginia Eubanks shows in her book *Automating Inequality,* when AI systems are deployed as part of the welfare state, they are used primarily as a way to surveil, assess, and restrict people's access to public resources rather than as a way to provide for greater support.[87]

A key example of this dynamic emerged when former Republican governor of Michigan Rick Snyder, previously the chairman of computer hardware computer Gateway, decided to implement two algorithmically driven austerity programs in an attempt to undermine the economic security of his poorest citizens under the auspices of state budget cuts. First, he directed that a matching algorithm be used to implement the state's "fugitive felon" policy, which sought automatically to disqualify individuals from food assistance based on outstanding felony warrants. Between 2012 and 2015, the new system inaccurately matched more than nineteen thousand Michigan residents and automatically disqualified each of them from food assistance.[88]

The second scheme was called the Michigan Integrated Data Automated System (MiDAS), a system built to "robo-adjudicate" and punish those it determined to be defrauding the state's unemployment insurance. MiDAS was designed to treat almost any data discrepancies or inconsistencies in an individual's record as potential evidence of illegal conduct. The system inaccurately identified more than forty thousand Michigan residents of suspected fraud. The consequences were severe: seizure of tax refunds, garnishment of wages, and imposition of civil penalties that were four times the amount people were accused of owing. Ultimately, both systems were giant financial failures, costing Michigan far more money than it saved. Those harmed were able to successfully sue the state over the systems, but not before thousands of people were affected, with many entering bankruptcy.[89]

When viewed in the overall context of state-driven AI systems, one can see the consistent logics between targeting terrorists or undocumented workers and targeting fugitive felons or suspected fraudsters. Even though food assistance and unemployment benefits were created to support the poor and to promote social and economic stability, the use of militaristic systems of command-and-control for the purposes of punishment and exclusion undermine the overall goals of the systems. In essence these systems are punitive, designed on a threat-targeting model. The motifs of scoring and risk have permeated deeply through the structures of state bureaucracy, and the automated decision systems that are imagined in those institutions drive that logic deeply into the way that communities and individuals are imagined, evaluated, scored, and served.

The Tangled Haystack

I am almost at the end of a long day searching through the Snowden archive when I run across a slide that describes the planet as a "haystack of information," in which desirable intel is a needle lost somewhere among the straw. It includes a cheery clip art image of a giant haystack in a field with a blue sky overhead. This cliché of information gathering is tactical: hay is mown for the good of the farm, gleaned to produce value. This invokes a comforting pastoral imagery of data agriculture—tending the fields to further orderly extraction and production cycles. Phil Agre once observed that "technology at present is covert philosophy; the point is to make it openly philosophical."[90] The philosophy here is that data should be extracted globally and structured in order to maintain U.S. hegemony. But we've seen how these stories break down under scrutiny.

The overlapping grids of planetary computation are complex, cross-breeding corporate and state logics, exceeding traditional state border and governance limits, and they are far messier than the idea of winner takes all might imply. As Benjamin Bratton argues, "The armature of planetary-scale computation has a determining logic that is self-reinforcing if not self-fulfilling, and which through the automation of its own infrastructural operations, exceeds any national designs even if it is also used on their behalf."[91] The jingoistic idea of sovereign AI, securely contained within national borders, is a myth. AI infrastructure is already a hybrid, and as Hu argues, so is the labor force underpinning it, from factory laborers in China who make electronic components to Russian programmers providing cloud labor to Moroccan freelancers who screen content and label images.[92]

Taken together, the AI and algorithmic systems used by the state, from the military to the municipal level, reveal a covert philosophy of *en masse* infrastructural command and control via a combination of extractive data techniques, targeting logics, and surveillance. These goals have been central to the intelligence agencies for decades, but now they have spread to many other state functions, from local law enforcement to allocating benefits.[93] This is just part of the deep intermingling of state, municipal, and corporate logics through extractive planetary computation. But it is an uncomfortable bargain: states are making deals with technology companies they can't control or even fully understand, and technology companies are taking on state and extrastate functions that they are ill-suited to fulfill and for which, at some point in the future, they might be held liable.

The Snowden archive shows how far these overlapping and contradictory logics of surveillance extend. One document notes the symptoms of what an NSA employee described as an addiction to the God's-eye view that data seems to offer: "Mountaineers call this phenomenon 'summit fever'—when an 'individual becomes so fixated on reaching the summit that all else fades from consciousness.' I believe that SIGINTers, like the world-class climbers, are not immune to summit fever. It's easy enough to lose sight of the bad weather and push on relentlessly, especially after pouring lots of money, time, and resources into something."[94]

All the money and resources spent on relentless surveillance is part of a fever dream of centralized control that has come at the cost of other visions of social organization. The Snowden disclosures were a watershed moment in revealing how far a culture of extraction can go when the state and the commercial sector collaborate, but the network diagrams and PowerPoint clip art can feel quaint compared to all that has happened since.[95] The NSA's distinctive methods and tools have filtered down to classrooms, police stations, workplaces, and unemployment offices. It is the result of enormous investments, of de facto forms of privatization, and the securitization of risk and fear. The current deep entanglement of different forms of power was the hope of the Third Offset. It has warped far beyond the objective of strategic advantage in battlefield operations to encompass all those parts of everyday life that can be tracked and scored, grounded in normative definitions of how good citizens should communicate, behave, and spend. This shift brings with it a different vision of state sovereignty, modulated by corporate algorithmic governance, and it furthers the profound imbalance of power between agents of the state and the people they are meant to serve.

Conclusion
Power

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

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

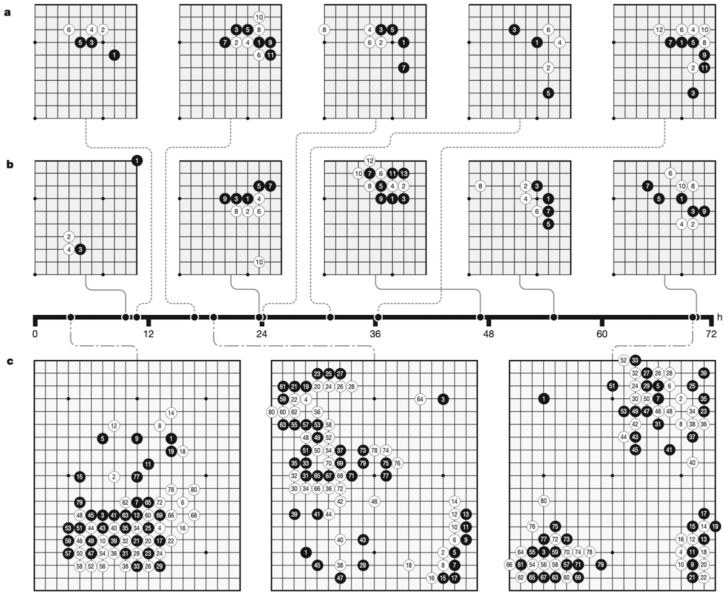

Go knowledge learned by AlphaGo Zero. Courtesy of DeepMind

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

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

Games without Frontiers

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

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

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

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

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

The Pipelines of AI

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

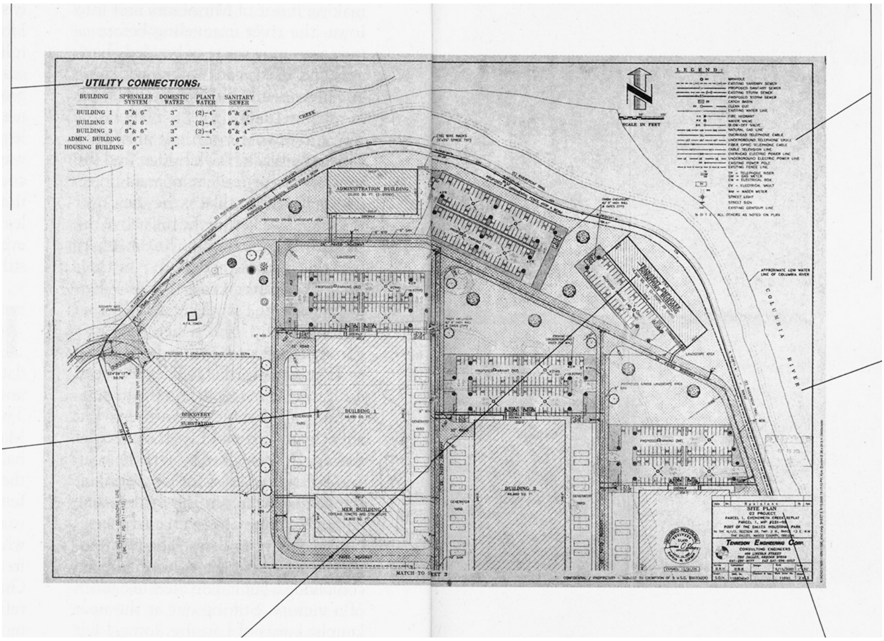

UTILITY CONNECTIONS:

BUILDING	SPRINKLER SYSTEM	DOMESTIC WATER	PLANT WATER	SANITARY SEWER
BUILDING 1	8"& 6"	3"	(2)-4"	6"& 4"
BUILDING 2	8"& 6"	3"	(2)-4"	6"& 4"
BUILDING 3	8"& 6"	3"	(2)-4"	6"& 4"
ADMIN. BUILDING	6"	3"	--	6"
HOUSING BUILDING	6"	4"	--	6"

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

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

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

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

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

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

The Map Is Not the Territory

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Toward Connected Movements for Justice

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

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

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

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

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

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

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

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

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