

DOI: 10.1215/2834703X-10734026

Data Worlds: An Introduction

KATHERINE BODE and LAUREN M. E. GOODLAD

It might seem a strange decision to devote the inaugural issue of a journal called *Critical AI* to the topic of “data worlds.” Whereas *artificial intelligence* evokes science fiction and the bleeding edge of technology, *data* suggests something more mundane and humble: numbers, spreadsheets, or graphs, for example. According to the *Oxford English Dictionary*, *data* is a plural noun derived from the Latin for “given,” denoting items of “(chiefly numerical) information,” especially when “obtained by scientific work” and “collected together for reference, analysis, or calculation.” *OED*’s sample of usages includes Richard Price, the early statistician and moral philosopher whose *Observations on Reversionary Payments* became the textbook for a still-emerging insurance industry. An actuary *avant la lettre*, Price ([1771] 1773: 287) believed that if London’s parishes improved their record-keeping, they would “supply the necessary *data* for computing accurately the values of all *life-annuities* and reversions.” If Price reminds us that *computation* also predates the digital computer (with the original signification of *computer* referring to a person and frequently a woman),¹ he also articulates a practical need for accurate data that many machine learning (ML) researchers in our own time also express.

It is all the more striking then that *data*—the substrate for ML’s predictive patterns—seldom comes up in the public-facing discourse of the most influential players in “AI.” Consider Stanford University’s Institute for Human-Centered AI. As faculty codirectors Fei-Fei Li and John Etchemendy (n.d.) evoke a coming “Age of Artificial Intelligence,” they seek to uplift readers with the good news of a technology that can enable human plurality. AI’s “new world” will be “more profound than any other period of transformation in our history,” they write. The “creators and designers” of this revolution “must be broadly representative of humanity. . . a true diversity of thoughts—across gender, ethnicity, nationality, culture and age, as well as across disciplines.” Notably, *data* is entirely absent from this aspiration, which is all the more noteworthy given that Li built her reputation in computer vision by creating a benchmark dataset, ImageNet (as Nicolas Malevé and Katrina Sluis discuss in this special issue). In a recent call for critical histories of data, Emily Denton and colleagues (2021: 2) note that influential datasets such as ImageNet “form the critical information infrastructure underpinning ML research and development, as well as a critical base upon which algorithmic decision-making operates.” Why is it, then, that as “AI” has become the preferred nomenclature for technologies rooted in these very practices, the focus on *data* has tended to fade from public-facing discourse?

Consider a second example of this noteworthy retreat. “Responsible AI: Looking Back at 2022, and to the Future” (Croak and Gennai 2023) is a blog post coauthored by Google’s VP for “responsible AI” and Google’s director of “responsible innovation.” The authors observe that since “AI is playing a larger role in society,” this “powerful and helpful new technology,” which is “core to Google products. . . must be developed thoughtfully and responsibly.” *Data* comes up in the context of the company’s “AI principles,” which include (in a sentence bereft of active subjects) the “resources” and “tools” provided for “monitoring products’ responsible AI maturity” such as “updates” on “data transparency.” *Prediction* surfaces in the humanizing claim that “people look to AI to address global issues ranging from disease detection to natural disaster prediction.” Enshrouding *data* in a nimbus of reassuring nouns (“transparency,” “maturity,” “ethics,” “principles”), and aligning *prediction* with salvation from disease and disaster, “Responsible AI” upholds a “helpful” technology managed by dutiful corporate stewards. What the authors decline to mention is that appropriating customers’ data for predictive analytics and the sale of ads has been Google’s core business model since the company’s IPO in 2004.² Indeed, Google pioneered this lucrative mode of “surveillance capitalism” (Zuboff 2019; cf. Zuboff 2015; Foster and McChesney 2014; Doctorow 2020), which Facebook embraced in 2008

when it hired Sheryl Sandberg, the former head of Google’s advertising division (Frenkel and Kang 2021).

Compare these recent examples of elite “AI” discourse to a McKinsey Global Institute report on big data from about a decade ago (Manyika et al. 2011). Whereas data hardly dares to speak its name in the “AI” messaging of Stanford and Google, McKinsey’s analysts can hardly contain their ardor:³

The amount of data in our world has been exploding. Companies capture trillions of bytes of information about their customers, suppliers, and operations, and millions of networked sensors are being embedded in the physical world in devices such as mobile phones and automobiles, sensing, creating, and communicating data. Multimedia and individuals with smartphones and on social network sites will continue to fuel exponential growth. Big data—large pools of data that can be captured, communicated, aggregated, stored, and analyzed—is now part of every sector and function of the global economy. . . . It is increasingly the case that much of modern economic activity, innovation, and growth simply couldn’t take place without data.

Notably, while McKinsey’s analysts dramatize the superfluity of *data’s* quantity and impact, they say almost nothing about *artificial intelligence*. Describing ML as an important approach to data analytics, they define it as a “subspecialty of computer science (within a field historically called ‘artificial intelligence’).” A contemporaneous article in the *Harvard Business Review* (McAfee and Brynjolfsson 2012) does not even mention “AI.” Arguing that big data has ushered in a “management revolution,” the authors look to Google’s director of research for an explanation: “We don’t have better algorithms,” he tells them; “we just have more data” (McAfee and Brynjolfsson 2012).

This narrative changed with the turn to “deep learning” (DL) after 2012. Li’s ImageNet (fourteen million images downloaded from the internet and labeled by on-demand workers through Amazon’s Mechanical Turk) precipitated this paradigm shift by providing a data-rich benchmark suitable for demonstrating the potential of “deep” “neural networks” for object recognition.⁴ Loosely modeled on the structure of the human brain, the idea of artificial *neural networks* dates back to the 1940s, but its realization awaited advances in the speed and power of computer processing as well as sufficient caches of data.⁵ The specified *learning* attributed to DL (as in all ML) entails the system’s ability to optimize performance on a data-driven task by updating the weights in an elaborate network of calculations. That is so whether the predictions in question concern the next move in a game of Go, the probable behavior of a hurricane, an individual’s creditworthiness, or (as in OpenAI’s much-discussed ChatGPT) a likely sequence of words in response to a user’s prompt. What is “deep” in DL thus connotes the multiple layers in a pretrained model through which new inputs pass before delivering one or more predictive outputs.

As both DL enthusiasts (such as pioneer Terrence J. Sejnowski) and critics (like AI Now Institute cofounder and Signal president Meredith Whittaker) agree, it was above all “big data” that drove the DL “shock wave” of the 2010s (Sejnowski 2018: 10, 127–28). A much-cited quotation attributed to Google’s CEO in 2010 estimated that every two days, more information is generated through digital devices than the five exabytes created from “the dawn of civilization” to the dawn of the twenty-first century—and “the pace is increasing” (Seagate Technology n.d.). Thus, as Whittaker (2021: 51–52) puts it, the machine vision “breakthrough” of 2012 was not the neural design itself but rather what “large-scale data and computational resources” enabled that architecture to do. “The year 2012,” she writes, “showed the commercial potential” of data-driven ML, and “the power of AI” as a “marketing hook. . . . *Tech companies quickly (re)branded machine learning and other data-dependent approaches as AI, framing them as the product of breakthrough scientific innovation*” (emphasis added).

In our own time, as Big Tech faces “mounting regulatory pressure,” the need to create such “tech-positive narratives” (51) helps to explain both the boosting of “AI” and the correlative muzzling of straight talk about data that we have observed. Long associated with science fiction, and imbued with mysterious anthropomorphic endowments, “AI” provides both an ideal “marketing hook” and a distraction from a worrisome corporate concentration of power and resources—including data.⁶ The focus on “AI” thus encourages an unwitting public to allow a “handful of tech firms” to define research agendas and “enclose knowledge” about data-driven systems “behind corporate secrecy”

(51–52). As Whittaker draws sobering parallels between AI’s political economy and the Cold War–era military-industrial complex, she urges knowledge workers in industry and academia to organize. Only by developing lasting solidarities—“muscles of care and mutual accountability”—will researchers and their allies empower themselves to “name the dynamics of coercion and capture more safely” (55). “Coercion and capture” may recall the well-known story of Google suppressing the research of its ethics division. But it also speaks to Whittaker’s own experiences with Google and with New York University, the latter of which defunded the AI Now institute in 2021.⁷

If this is hardly an uplifting glimpse into the dominant political economy of “AI” in our time, it points to the pressing need for interdisciplinary collaboration, and it explains why the coeditors of this special issue chose to focalize data. Our initial collaboration took the form of a yearlong workshop series, sponsored by the National Endowment for the Humanities and a Rutgers Global seed grant. Setting aside AI per se, we first organized a series, “The Ethics of Data Curation,” to explore the curatorial processes and practices through which data is (or might be) mobilized for ML broadly conceived. We envisioned the successor, “Data Ontologies,” as an opportunity to focus on the attendant implications of data’s entanglements in sociomaterial relations, being, and worlds. Yet, as we and our workshop participants soon discovered, questions of *data curation* (including the ethical consequences of what is done to and through data) inevitably raise questions about what data *is*. That is, since *data* becomes what it is in relation to these very practices, distinctions between data *creation* and data *curation* tend to blur. Moreover, knowledge *about* data is invariably complicated by the profound ways in which data (as a presumed object of knowledge) participates in forming the situations and structures in and through which data can be known (whether practically, theoretically, or both).

If these observations confirm the lessons of an influential ontological turn in the humanities and social sciences that has emphasized the entwinement of being and knowing, they likewise recall the insights of historical materialism as a mainspring for cultural studies.⁸ As Raymond Williams (1989: 151–52) put it in “The Future of Cultural Studies,” “You cannot understand an intellectual or artistic project without also understanding its formation.” Thus, the “crucial theoretical intervention” of cultural studies was to demonstrate how *formation* is never merely the “context” or “background” for such projects. By refusing “to give priority either [to] the project or the formation,” Williams wrote, cultural studies set out instead to grasp that “*common* disposition of energy and direction” that both were materializing. In doing so, it enabled a position from which “to understand existing and possible formations” and to activate “certain projects towards the future.” Put this way, our workshops reminded us that big data and DL are not inert backgrounds for data curation (any more than “AI” and its discourse are mere contexts for today’s data worlds). Rather, in researching the formation of a given political economy, or probing the possibilities of a countervailing practice, one should look for that “*common* disposition of energy and direction” on which future projects could depend.

More than a year later, as our collaboration gives rise to this special issue, we recognize *critical AI studies*—understood as a field that brings the humanities, social sciences, and arts into dialogue with AI-adjacent technoscientific disciplines—as the bearer of these synthetic insights. We perceive the potential to theorize, historicize, and open pathways for transforming the structures of power that have consolidated around data-driven technologies at their most unaccountable. As critical theorists, we locate this intellectual project in an array of crises that so-called postcritical positions are powerless to address: not only the deeply undemocratic political economy and social harms of “AI” broadly conceived, but also (and in tandem) climate precarity, political polarization, racism, misogyny, police brutality, rising autocracy and plutocracy, and the systematic underfunding of the public infrastructures (or commons) necessary to democratic thriving.⁹ As interdisciplinary scholars, we favor collaborative opportunities (including the testing of new research, alternative technologies, teaching experiments, policy ideas, and community partnerships) that cross an unproductive “two cultures” division that continues to impede critical dialogue. In both regards, we are fortunate (indeed, honored) to edit and introduce a special issue that:

- opens with an interview in which Sasha Costanza-Chock—media scholar, designer, head of research for OneProject.org, and longstanding member of the Design Justice Network—discusses practices of design justice with interlocutors from the fields of justice and technoscience (Katherine Henne), decolonial natural language processing (Sabelo Mhlambi), and distributed information processing (Anand Sawarte);
- continues with a discussion of how Li’s paradigmatic curatorial pipeline for ImageNet activated a “latent photographic theory” that extended an instrumental realism dating back to the scientific positivism of the nineteenth century and that remains discernible in large image models such as DALL-E 2 (coauthored by Nicolas Malevé, a visual artist and computer scientist, and Katrina Sluis, a scholar of photography and media arts);
- proceeds to a “provocation” for “thick” and multisensory approaches to studying data capitalism (coauthored by anthropologist Caroline Schuster and digital curation scholar Kristen Schuster), conceived partly as a response to Marion Fourcade and Kieran Healy’s (2017) influential framing in “Seeing Like a Market” and seeking to resist the rhetoric of seamless data formation and monetized implementation through attention to the diverse actors, relations, and situations that constitute technological systems;
- continues with an article in which Christopher Newfield, a scholar of literature and numeracy studies, interprets Brian Cantwell Smith’s *The Promise of Artificial Intelligence* as a call for “epistemic equality,” building on the latter’s cogent account of data-driven machine “reckoning” in order to contrast Smith’s schema for a robustly ethical “AI” to the problematic “two cultures” division between STEM and humanities disciplines that C. P. Snow inadvertently ushered in during the Cold War era, and which continues to shape technology discourse in our own time;
- and closes with a manifesto by Sam Lavigne, an artist and technology educator, who articulates “Scrapism” as a way of resisting the privatization of the internet by deploying the very same practice private businesses use—scraping the web—to reopen data silos and repurpose them for public projects.

The issue also includes a vibrant array of interdisciplinary book reviews.

Although we are eager to open up this conversation, we want to be clear that the project of critically engaging data practices and worlds is more than even this interdisciplinary assemblage of outside-the-box thinkers can hope to compass in a single cluster of essays. Thus, we uphold their work (and offer our own) not in the spirit of a *fait accompli*—a doorway to a historic condition that one opens and shuts—but rather as an invitation to a space of dissensus that, we hope, may expand our readers’ lived relations to data and data worlds. To quote Jacques Rancière from a 2021 interview on the COVID-19 pandemic: “All we can do” is to “try to create and enlarge spaces of non-consent.” The “challenge” is “to maintain dissensus, maintain a distance. What can this distance produce in the future? I don’t know. But even these figures of distance are a way of living differently in the world we challenge” (Dejean and Lalanne 2021).

On Theorizing Data

We turn now to thoughts gathered during and since our workshops on the question of why *data* is so challenging to theorize—explaining why we coeditors believe that it is often more productive to focalize *datafication* (defined as the practices and relations through which “data” is constituted and made legible) and *data worlds* (the ontoepistemological conditions from which those practices and relations emerge and to which they give rise). As we have already observed, in ordinary language, *data* denotes “items of (chiefly numerical) information,” especially when “obtained by scientific work.” According to the authors of *Data Feminism* (D’Ignazio and Klein 2020), *data* in this familiar sense emerged in the early modern period to confer rhetorical authority on information that served the powerful—a restriction on who and what counts in data formation that feminist and intersectional data science must perpetually contest.

While “data” has continued to convey objectivity across many centuries, the conditions, of course, have dramatically changed. When Richard Price urged London’s parishes to expand their record-keeping, his goal was to spur the data creation necessary for the actuarial projects of a modernizing economy. More than two centuries later, what McKinsey’s analysts call “big data” no longer depends exclusively on records or “scientific work” and instead “explodes” through “millions of networked sensors” that “capture trillions of bytes of information.” As McKinsey’s report celebrates the creation of data and points to its harvesting for “economic activity, innovation, and growth,” one perceives the entrenchment of a feedback loop. Data’s pervasiveness and multifariousness tend to reinforce data’s profitability and vice versa.

In its multifariousness and pervasiveness, data includes not only numbers but also, in Sarah Ciston’s (2023) words, “emails, a collection of scanned manuscripts, the steps you walked to the train, the pose of a dancer, or the breach of a whale.” *Data*, in her broad definition, is “values assigned to any ‘thing.’” Jathan Sadowski’s (2019: 2) more explicitly technocentric definition makes *data* “a recorded abstraction of the world created and valorised by people using technology.” As he goes on to clarify, “not all data is the same, nor is it used the same way” (5). While Ciston accentuates a real heterogeneity, Sadowski’s emphasis on valorization, like Catherine D’Ignazio and Lauren F. Klein’s stress on objectivity, reminds us that *how* data is used (including by whom and for what purpose) is inseparable from what data turns out to be. Computer scientist Rediet Abebe and economist Maximilian Kasy (2021) have coined the term “means of prediction” to highlight the stakes of such queries. Since data-driven ML systems are designed to “maximize some objective,” they urge, researchers must take care to ask: “Whose goals count?”¹⁰

Notably, at least one prevalent expression of data’s multifariousness sets out to obviate that line of research. The discourse of *data positivism*, as we will call it, holds that unprecedented scale has transformed data from a mere constituent in the knowledge-making enterprise to a unique epistemic substrate. The most audacious example of this position is Chris Anderson’s (2008) often-discussed article in *Wired*, which heralds the so-called Petabyte Age as “the end of theory” and the “scientific method.” “We can throw the numbers into the biggest computing clusters the world has ever seen and let statistical algorithms find patterns where science cannot.”¹¹ Anderson here riffs on a contemporaneous discourse that Google’s computer scientists (Halevy, Norvig, and Pereira 2009) soon dubbed *the unreasonable effectiveness of data*.¹² Born of research that Google alone could advance at hyperscale, the “unreasonable effectiveness of data,” in the words of data journalist Meredith Broussard (qtd. in Chen 2018; cf. Broussard 2018: 118–19), articulates the belief that “if I have a big enough dataset, I don’t actually need to be an expert on the topic that the data covers. I can draw conclusions simply from the vastness of the data.” Notably, the meaning of *unreasonable* in this context echoes a 1960 essay on the unanticipated correspondence between mathematics and physics, where it speaks to seemingly “miraculous” insights that researchers do not “understand.” *Unreasonable*, that is, denotes the human scientist’s stipulated inability to fathom (or reason about) the grounds of an adduced “effectiveness.”¹³ Whereas Google’s claims for the powers of data-derived pattern-finding are thus hyperbolic, Anderson’s data positivism borders on a theism or sublime. In effect guiding his readers from data *science* to data *omniscience*, his essay positions predictive algorithms as the oracular successors to causal inquiry and testable hypotheses: “This is a world where massive amounts of data and applied mathematics replace every other tool that might be brought to bear.”

Thus, for those who share Anderson’s (2008) belief that “correlation is enough,” causal questions are obsolete: “Who knows why people do what they do?” he writes. “The point is they do it, and we can track and measure it with unprecedented fidelity.” It does not matter for Anderson—as it does for AI researcher Judea Pearl and his coauthor Dana Mackenzie (2018: 6)—that data cannot tell you *why*; that it is, to that extent, “profoundly dumb.” “With enough data,” Anderson asserts, “the numbers speak for themselves.” Nor does Anderson seem aware that the datasets to which he attributes this “unprecedented fidelity” are frequently shown to be riddled with errors and biases of various kinds. As mathematician Cathy O’Neil (2016: 204–5) has argued, when big data processes underrepresent difference, they “codify the past,” injecting yesterday’s prejudice into tomorrow’s decision-making.¹⁴ Moreover, claims for the revolutionary potential of correlation date back to the late nineteenth

century's eugenicist pseudosciences. The "urtext" for such projects, according to Kate Crawford (2021: 93), is the criminal mug shot: an example of how images freighted with history join an "aggregate mass" of data points to "drive a broader system." The eugenicist roots of correlation matter, explains Wendy Hui Kyong Chun (2021), because, then as now, correlation aims to predetermine the future. Thus, behind the "unreasonable effectiveness of data," we contend, is an unaccountable positivism that perpetuates discrimination while legitimating a political economy that fails to ask: "Whose goals count?"¹⁵

One way to challenge data positivism is to shift the theoretical gaze from *data* (a product) to *datafication* (the ongoing formation through which that product is created and deployed). Consider that while *data*, as Ciston puts it, is "values assigned to any 'thing,'" the converse is equally true: that which has no assigned "values" is not—or is not (yet) *recognizable as*—data. Thus, when Google's Schmidt contrasts the exiguous "five exabytes" collected since "the dawn of civilization" to the data multitudes of his own time, he is, in effect, pointing to the surveillant practices that his company pioneered. Nonetheless, to observe that a past civilization was not "datafied"¹⁶ is merely to show that it did not privilege the kinds of record-keeping that a calculation like Schmidt's might conceivably register. It does not mean that the culture in question did not assign "values" to myriad "things" that could, in the present day, be extensively datafied. That is why researchers like José van Dijck, Poell, and de Waal (2018: 33) define *datafication* in terms of the extraordinary investment in techniques that "render into data many aspects of the world that have never been quantified before."¹⁷ Put differently, *datafication* entails the processes or practices that render a vast multiplicity of objects, locations, activities, and conditions legible *as* data.

For example, according to Rafael Alvarado and Paul Humphreys (2017: 736), the emergence of massive online platforms that mediate billions of "person-to-person communications" has given rise to new and "historically unique" forms of *entextualization*: a data-creating process through which "ephemeral discourse is materialized into persistent media forms." When users converse on social media, for instance, these "entextualized" communications may "influence social life beyond the original context of utterance." In a comparable vein, D'Ignazio and Klein (2020: 12–13) describe how "corporations, often aided by academic researchers, are currently scrambling to see what behaviors—both online and off—remain to be turned into data and then monetized." "Nothing," they add, "is outside of [this] *datafication*." These observations demonstrate how the data-driven political economy has propagated (if not wholly normalized) surveillant practices that would have been deemed exceptional in the pre-internet era.¹⁸

That said, to notice these structures is by no means to assume that datafication has (or soon will) become *universalized*.¹⁹ Indeed, it is important to contest the misguided notion that huge datasets are effective totalities that imbue large models with earth-spanning comprehensiveness or democratizing effects. Such faulty assumptions at minimum forget that about a third of the world's population has no internet access whatsoever. Thus, while *data worlds* are demonstrably global phenomena, the processes through which they become so are ongoing, heterogeneous, and unevenly developed. The name "Google" derives from a neologism for an unfathomably large number. But the datafication that "exploded" alongside the rise of the search engine has produced huge epistemological gaps and zones of exclusion (e.g., Leurs and Shepherd 2017: 125; cf. Noble 2018; Benjamin 2019; Eubanks 2017). Indeed, according to one team of researchers (Tapo, Coulibaly et al. 2020: 23), "the vast majority of the world's languages—*representing billions of native speakers worldwide*"—lacks sufficient data to support state-of-the-art performance on common tasks in natural language processing (emphasis added). Moreover, this lack of data typically "reflects a broader paucity" of digital resources. Hence, the more that *datafication* appears to approach a universal condition, the more that very process entrenches exclusions, inequities, and silences.

Data and Capitalism

While data's multifariousness and pervasiveness thus challenge researchers to eschew a false positivism or universalism, the issue of data's *profitability* complicates the work of theory in different ways. Many commentators have likened data to *oil*—a fuel that runs the predictive machinery of the new

digital economy (e.g., Srnicek 2016: 40; Sejnowski 2018: 3). But such framing of data as an extractable natural resource obscures the complex and contingent processes through which data is created and (often) appropriated for private use.²⁰ The notion of data as a *commodity* is also limited, not least because not all data is commodified. On these grounds, Sadowski (2019) suggests, it is more instructive to regard data as *a form of capital*—a topic we will now explore.

Probably the best-known account of the new datacentric political economy is Zuboff's (2019) *Age of Surveillance Capitalism: The Fight for a Human Future and the New Frontier of Power*, which centers on the profitable tracking of behavioral data for the purpose of targeted advertising and other content. Zuboff's coinage of *behavioral surplus* to theorize this lucrative business model (introduced by Google but adopted by Facebook, Microsoft, and many others) draws on an analogy to the Marxist notion of *surplus labor*. Roughly speaking, she argues that surveillance capitalism's power to "extract" user data in excess of what digital platforms require for efficient function constitutes a "digital dispossession" that "forms the basis of a wholly new class of market exchange" (99). Scholars have criticized this thesis on the grounds that (despite what advertisers may believe) the assumption that microtargeting is an effective "means of behavioral modification" (203) lacks sufficient evidence. Moreover, Zuboff's notion of a behavioral "surplus" sits awkwardly on its Marxist foundations. According to Marx, *surplus labor* is the value of workers' labor power over and above the means of their own subsistence, a surplus that capitalists appropriate for themselves. Through a leaky analogy, Zuboff makes individual behavioral data equivalent to labor power even though the exploitation in question differs from the exploitation of workers deprived of surplus value.²¹ Her analysis thus tends to proffer emotive appeals: "Forget the cliché that 'if it's free, you are the product.' You are not the product; you are the abandoned carcass. The 'product' derives from the surplus that is ripped from your life" (377). Lee Vinsel (2021) has argued that Zuboff engages in *criti-hype*, defined as a scholar's unwitting rehearsal of inflationary marketing claims made on behalf of "emerging technologies" for the intended purpose of criticizing them. However, as we will suggest, *The Age of Surveillance Capitalism* more successfully lays the ground for a Marxist theory of *primitive accumulation* (99).

Let us begin by comparing Zuboff's study to Fourcade and Healy's (2017) influential article "Seeing Like a Market," which Schuster and Schuster discuss in this issue. As researchers in economic sociology, Fourcade and Healy focus on how a new "data imperative" has enabled neoliberal markets "to 'see' in a new way" and to "teach us to see ourselves in that way" (10). In making that case, they adapt Pierre Bourdieu's usage of *symbolic capital* (14) to index the "intangible forms" of capital that individuals accrue "from their social position," such as education, professional attainment, and/or lifestyle.²² The new market imperative to "see" personal data, Fourcade and Healy suggest, creates a novel form of symbolic capital through algorithmic "scoring, grading and ranking" (14).²³ To this extent, Fourcade and Healy offer a Bourdieu-inflected alternative to Zuboff's "behavioral surplus." Both theories point to a lucrative process of "sorting and slotting people into categories and ranks" (14). But rather than allege a wholly new means of behavioral modification, Fourcade and Healy perceive the amplification of existing socioeconomic and moral hierarchies. When predictive models pair "a wealthy customer with a quality credit card, or a heavy drinker with a bad insurance plan," these algorithmically mediated determinations are "perceived to be natural" because much of "economic life is *already structured this way*" (17, emphasis added). Moreover, the outcomes are rooted in "cues about ourselves that we volunteered, or erratically left behind, or that were extracted from us in various parts of the digital infrastructure" (17). For the disadvantaged or marginalized, the "seeing" market thus amounts to new ways of "subsum[ing] unlucky circumstance and uncaring structure into morally evaluable behavior" (25).

In their essay in this special issue, Schuster and Schuster engage Fourcade and Healy from the vantage point of feminist and queer theories of embodiment as well as critical data studies. They offer "thick description" as a means of attending to multisensory complexity in the effort to interrupt the technocentric affirmation of a "seamless and inevitable journey" from data creation to "monetizable domain knowledge and useful services." "Fine-grained empirical analysis that pays attention to embodied human behaviors beyond 'seeing,'" they contend, "can reveal alternative meanings and

values attached to data, algorithmic processes, predictive modeling, and practices of ‘AI’ writ large.” Schuster and Schuster thus invite us to supplement Fourcade and Healy’s analysis.

Imagine, for example, an analysis of the “seeing market” that documents the frequent failures—functional as well as social—of automated decision-making systems that erroneously predict fraud, credit risk, or facial matching with individuals sought by law enforcement.²⁴ These still underreported failures, which disproportionately affect poor people and people of color, occur when systems are trained on datasets that reduce tangled relations, multisensory bodies, and complex social situations to available data points. An economic sociology alert to these devastating blunders could open new research pathways in which, for example, “Seeing Like a Market” confronts “Failing Like a Decision-Making System,” “Discriminating Like a Dataset,” or “Incarcerating Like a Proprietary Algorithm.” At minimum, a critical AI studies in this vein could interrogate the “uncaring structure” that imposes “morally evaluable behavior” on those least able to combat the devastating effects of dysfunctional decision-making technologies. Such systems, as Williams (1989: 151–52) reminds us, are never merely “context” or “background.” Probing their failures can thus motivate genealogies of the “means of prediction” that vividly demonstrate whose goals “count”; research partnerships with affected parties premised on design justice principles; cross-disciplinary spaces of Rancièrian dissensus; or, in Fredric Jameson’s (2009: 146) dialectical terms, opportunities to “change the valences” by “adjust[ing] the lens of thought” on the problem “in such a way” as to lead to “new and wholly unexpected directions.”

Although Fourcade and Healy stop short of a valence-changing lens on the seeing market, their appeal to Bourdieuvian categories is nonetheless suggestive. In the most extensive Marxist theorization so far, “When Data Is Capital,” Sadowski (2019) joins Fourcade and Healy in invoking Bourdieu’s modes of extra-economic capital. As in Bourdieu’s social, cultural, and symbolic formulations, Sadowski’s *data capital* supports *economic capital*—that is, money—without necessarily converting into the latter. Instead, data capital is institutionalized in infrastructures for the “collecting, storing, and processing” (4) of data. In this infrastructural form, the impact of data on human subjects may manipulate behavior (Zuboff), naturalize moral hierarchies (Fourcade and Healy), or both. But Sadowski’s case for “data as capital” extends beyond subjective effects by theorizing the mechanisms through which this new political economy is “driven by the logic of perpetual (data) capital accumulation and circulation” (2).

To explain, he quotes *Capital Volume I*, reminding us that capitalism’s immediate aim is never to create “use-value” or even to realize “profit,” but rather to propel the “unceasing movement of profit-making” (4). Here he draws on Marx’s well-known description of capitalism as a system that favors the endless circulation of *exchange-value* (in the form of liquid capital or money) over the socially beneficial distribution of *use-value* (in the form of commodities that can be consumed to satisfy social needs). This endemic privileging of exchange over use derives from a logic of *accumulation* (Marx [1867] 1990: 254–55) which further prescribes that “ceaseless” circulation augments capital’s value. Building on these insights, Sadowski (2019: 4) notes that, like money, data “can be continually captured and circulated” through a “logic of capital accumulation.” When present-day firms amass data, they regard these activities as “an intrinsic motivation” or “driving force” in the same way that capitalism in general regards the ceaseless accumulation of capital. The circulation of liquid capital that Marx ([1867] 1990: 253) described as “an end in itself” thus extends to data, fueling the expansion of data-accumulating infrastructures and practices.

Before developing this important thesis, we turn to the most recent effort to theorize data’s new political economy: Orit Halpern and colleagues’ (2022) notion of *surplus data*. Written as the introduction to a coedited special issue on that topic in *Critical Inquiry*, the essay lays down a proposition with which we largely agree: “Data has become a source of capital dynamics, a means of governance and control of populations, a mode of the administration of territory, in short, a new structural condition” (199). At the same time, the authors conceive data (as we also do) as a “material agent” (198): *data*, they write, does not merely describe or represent the world; rather, it constitutes a force in “creating the world it would describe” (203). Although we concur, we regard the latter position as a characteristic insight of the ontological turn—part of an ongoing effort to theorize worlds in light of

more-than-human practices and nonhuman objects such as data. By contrast, Halpern and colleagues argue that whereas data “was once a stable, recorded point for static reference” or “an abstraction from the ‘real’ world,” it has only recently become a creative force. Indeed, this putative transition from passive to active constitutes “surplus” data: “Data does more than it was intended to and often produces greater effects than measurement of a world that stands externally to it. Put simply, in our era, data is not simply descriptive or analytical but actively constructive.”²⁵

To be sure, dramatic changes in the scale and application of data analytics do indeed constitute a “new structural condition”—so much so that, as we have shown, the biggest players prefer mystified talk of “responsible AI” to legible discussion of data processing. However, for that very reason, we think it misleading to frame this condition as the product of data’s (allegedly) “new” ontoepistemological capabilities while downplaying the politicoeconomic concentration of power and resources. Such assumptions invite the hyperbolic claims of data positivism; indeed, Halpern et al. (197; 199–200) point both to Anderson’s *Wired* essay and Google’s “unreasonable effectiveness of data” as key evidence for data’s new “efficacy and impact on the world.” While they do not celebrate this positivism, their theorization explicitly requires it. “Do we have the terms and concepts to account for [the] ‘unreasonable effectiveness’ of large data sets?” they ask. “Or has data become so complex in its constantly growing and changing surplus that conventional critique is now untenable, unable to keep up as data claims to produce more and more insight?” For Halpern et al. the answer is largely yes.

Although there is much worth discussing in this thought-provoking special issue, we focus on the claim that “surplus data” has become “a structural condition of capital” (199). Halpern et al. propose their key term as an alternative to Zuboff’s *behavioral surplus*, and, like Zuboff, model their formulation on Marx’s *surplus labor*. They write, “it is not behavior that is in surplus, even if behavior recorded bolsters surplus value” (200). Rather, “what Google has achieved is the transformation of a *finite*, if extremely large, resource into a seemingly endless source of value through the recombination and discovery of new relations and patterns in *the same data set*. *Finitude* becomes a flexible frontier through new techniques of data analytics. . . . There is always a little bit more to obtain from an existing data set” (200–201; emphasis added). Here Halpern and colleagues reject Zuboff’s thesis only to propose another leaky analogy for surplus labor. For, despite careful parsing of Marx (“the secret of surplus labor. . . is nothing more than extension of the working day beyond the time required to replace the value put into production” [201]), the quoted passage is neither a correlative for a longer working day nor, more importantly, a lucid account of how data works (for Google or anyone else). For one thing, it is difficult to see how (absent a global catastrophe) data can be characterized as “finite” in any practical sense. As Halpern and colleagues are well aware, new data flows copiously through the myriad channels that powered its “explosion” twenty years ago. Hence, while datasets do yield multiple “insights,” get reused for varied purposes, and retain some value far into the future, this “flexibility” by no means obviates the need—and still the less the drive—to amass *new* data aggressively. Indeed, the premium on fresh data streams is so considerable that when Apple installed new privacy features on its cell phones, their move cost Meta (Facebook) billions in lost revenue, despite the social media company’s colossal stock of “existing data sets” (e.g., O’Flaherty 2022).

By zeroing in on the value-adding powers of algorithmic “optimization” (“extending data life beyond. . . its initial gathering” [Halpern et al. 2022: 201]), as distinct from the need for continuous data accumulation, Halpern et al. portray “surplus data” as a technoscientific feat. Their scant attention to the social and infrastructural requirements of this new structural condition thus marginalizes political action. In fact, strong privacy measures and the enforcement of antitrust laws have the potential to curb data surveillance in accord with popular will. It follows that the “surplus” in question springs as much from the size and underregulation of tech monopolies, which are determined to accumulate more and more data, as from the “insights” of predictive analytics—which, in actuality, vary in quality. Likewise, robust consumer and civil rights protections (buttressed by emerging practices of algorithmic auditing) can significantly redress the faulty decision-making systems behind so much real-world harm.²⁶

As Crawford (2021: 202–3) laments, data-driven systems are too often perceived as both magically unknowable and absolute in their predictive pattern-finding—an “enchanted determinism” that

“obscures power and closes off informed public discussion, critical scrutiny, or outright rejection.” A critical AI studies that refuses the theisms of “unreasonable effectiveness”; probes patterns of inequity; tracks monopoly power; and explores community collaborations, policy, legislation, and alternative technologies can help to ignite these discussions and enlarge popular dissent. Such projects, we contend, could inspire a valence-changing dialectics: on one side, dismissing a technolibertarian mirage that portrays the uptake of surveillant products as “democratization,” while on the other, building alternatives (e.g., design justice collaborations, interdisciplinary partnerships, and the enactment of a new data commons). Before saying more, we turn to some of the ontological thinking helpful to addressing such data worlds.

Data Ontologies

The topic of “Data Worlds” is, of course, explicitly ontological. While it is now common to speak of an “ontological turn” in the social sciences and humanities, what that points to in practice is an increasing willingness to set aside the presuppositions of any anthropocentric ontology. Rather than focus on what humans specifically contribute to the world, or on a world presumed to be uniquely available for human apprehension or use, the ontological turn has made theorists more alive to the limitations of any perspective that begins with what Eduardo Viveiros de Castro (2015: 4) has described as a “Culture/Nature distinction,” according to which “persons and things (also, humans and non-humans)” and “language meanings and extra-linguistic reality (concepts and objects)” are irrevocably divided. As frequent collaborators, we coauthors differ in the ontological emphases we find most generative for our research. But we share the belief that critical AI studies can benefit both from a nuanced *critical realism* that recognizes differences in the way humans and nonhumans participate in data practices and worlds, and a more decidedly “flat” ontology that challenges human-centered practices and states of being including intentionality, individuality, and self-determination.²⁷ To understand these varying commitments and their potential to shape critical accounts of technology, we might imagine a spectrum, with human-centered models of the world at one end and, to borrow Arturo Escobar’s (2017) term, a “pluriverse” at the other.

An influential argument at the human-centered end of this spectrum is Bender et al.’s (2021) “On the Dangers of Stochastic Parrots: Can Large Language Models Be Too Big?” (hereafter “SP”), which became a touchstone in broader debates over “AI” and ethics, due partly to Google’s aggressive efforts to silence it (e.g., Hao 2020). These efforts backfired spectacularly, as the ensuing publicity massively increased the audience for this article’s account of the dangers of large language models (LLMs), while suggesting the unequal racial and gender dynamics at play in Google’s “ethics” division. What interests us here, however, are the article’s (largely implicit) ontological commitments. That is, while we endorse the article’s important assessment of LLMs’ dangers along with its reasonable recommendations for redress,²⁸ we hope to demonstrate what arguments in this vein might gain from dialogue with humanist (and posthumanist) ontological perspectives.

To begin, the article’s core argument—that LLMs “are not performing natural language understanding”—rests on an a priori distinction between humans and language models elaborated through discussion of the communicative contexts for human language. The reason why LLMs “do not have access to meaning,” according to the authors, is that human communication is premised on individuals who “share common ground,” have “communicative intents,” and “model each other’s mental states as they communicate”—embodied intersubjective capacities that LLMs categorically lack. More specifically, since natural languages are “systems of signs” that pair “form and meaning,” and since the training data for LLMs “is only form,” these systems “do not have access to meaning” (615).²⁹ This human-centered formulation facilitates the article’s robust discussion of social harms and its critique of the anthropomorphic hyping of “AI.” But it also limits the authors’ ability to consider more-than-human practices and worlds.

A clear example is the use of nonhuman animals as analogies for language models, including the titular parrot in SP, and the octopus (“O”) in an earlier article on which SP builds: Emily M. Bender and Alexander Koller’s (2020) “Climbing Towards NLU.” While the “parrot” provides a familiar metaphor for how LLMs repeat linguistic forms “mindlessly,” “mechanically,” or “by rote,”³⁰ the

octopus metaphor combines an updated variation on John Searle's (1980) "Chinese room" with elements of the so-called Turing test (Turing 1950). To explain how LLMs can produce coherent-seeming language sequences without understanding, Bender and Koller (2020: 5188) introduce O, "a hyper-intelligent deep-sea octopus" who "is very good at detecting statistical patterns." By observing the communications via undersea telegraph of two humans (A and B) on separate islands, O uses statistical modeling to predict what words each will use.³¹ But while O knows the form of words, he is unable to forge "the relation" between "form and something external to language"—which is to say, *meaning* as Bender and Koller define it. In failing this test, O demonstrates that he is not human and that he "lacks the ability to connect [his] utterances to the world."

This is a good story; but it rehearses the tautological truth that (metaphorical) parrots and (fictional statistics-loving) octopi, like the LLMs they figure, cannot engage in "human-human communication," a stance that considers these nonhumans only in terms of their lack of access to human practices and ways of knowing. Absent from this playful setup is any curiosity about this animal's state of being—the same anthropocentric perspective that transforms the highly intelligent parrot into a metaphor for lack of understanding. Also absent is the possibility of O sharing the telegraph with other octopi who connect these signals to a world they inhabit together. Indeed, because LLMs (which in practice are trained and run on clusters of computers) are analogized to isolated animals, relationality in this ontology becomes an exclusively human feature. Once adduced as metaphors for mechanical and virtual entities, figurative animals lose the ontological status of sentient creatures that animals (figurative or otherwise) would usually inhabit.

We understand that the depiction of LLMs as tools contests a hyperbolic rhetoric that portrays proprietary "AI" systems as the triumphant humanlike creations of founders and engineers who, like Mary Shelley's modern Prometheus, imagine themselves as would-be Olympians stealing fire from the gods. But the same mentality that elevates self-appointed supermen also justifies the instrumentalizing treatment of human workers, nonhuman entities, and the planet itself. As we have seen, this very attitude prompts Stanford's institute to proffer "human-centered AI" as the answer to AI's problems even while obfuscating what that very technology actually is and does. The point is not that LLMs are sentient or that it is unethical to conceive them as tools. Rather, the point is to expand the ontological horizons through which researchers explore the world-shaping effects of data-driven analytics beyond the agonistic contest between hubristic anthropomorphisms and the "human-centered" agenda that purports to combat them.

As many readers of this journal will know, LLMs (the topic of *Critical AI's* next issue) are an especially controversial subject, not only because they trigger the "ELIZA effect,"³² but also because they are the principal technology behind the current "arms race" between the small number of companies now vying to commercialize and dominate "generative AI" at breakneck speed (Roose 2023). In the months and years ahead, as interdisciplinary researchers build on the important critiques of Bender and coauthors among others, they will encounter the long and ongoing histories of the many nonhuman practices and media that have shaped human communication—and which merit attention in their own right. Attention to the social, technical, cultural, and economic materiality of these diverse formations will enable nuanced (if always difficult, messy, and contestable) distinctions between technological practices that favor environmental stewardship, social justice, and democratic stability, and those that deliver power and profits to a small class of investors and corporate entities. Moreover, a singular focus on the "human" may unintentionally obscure the historical contingencies of social discrimination. Racism, sexism, ableism, and transphobia, after all, are not only attitudes or "biases" that humans either exhibit or resist; they are part of the entangling processes that produce some subjects as human and others as less than human.

That anthropocentrism can inform arguments that do not unequivocally divide the ontological status of human and nonhuman becomes clear as we move further along our posited spectrum to Smith's (2019) aforementioned *Promise of Artificial Intelligence: Reckoning and Judgment*. Newfield's essay in this special issue elaborates Smith's careful contrast between the failures of first-wave AI's symbolic formalisms and the comparative strengths of second-wave AI's *reckoning* (data-driven calculations).³³ However, neither of these approaches to AI has as yet been able to synthesize

judgment: a nuanced engagement with the world that represents human intelligence at its most ethical, and that, according to Smith, machines may perhaps achieve in the future.

Smith's relevance to the ontological discussion of an emerging critical AI studies, as we understand it, entails four key points.

1. First-wave AI failed because of the "vaguely Cartesian assumptions" of its human developers (8). In extrapolating logical rules and taxonomies from their own ordered perceptions and conceptual frameworks, these researchers failed to account for the mediation of human biology and culture in shaping their understanding of the world and mistakenly assumed that the world itself is transparent and discrete.
2. It follows that the data-driven "reckoning" of second-wave AI is useful in part because it abandons these human illusions. Contemporary ML systems "register" the world at a "subconceptual" level—through a wealth of "data obtained directly from low-level sensors—visual pixels, haptic signals, and so on" (58). Their "representations," wrought through the statistical modeling of large datasets, are relatively free from human conceptual baggage.
3. Although the detailed "registrations" of the world that ML systems thus generate can be epistemologically powerful, there is, for Smith, no foreseeable throughline wherein data-driven pattern-finding achieves *judgment*, a standard of "intelligence" that is exacting even for humans. Hence, unlike *Wired's* "end of theory" or Google's "unreasonable effectiveness of data," Smith's enthusiasm for the rewards of data-driven ML involves neither theistic nor positivistic absolutes.
4. Smith's rigorous standard of judgment rests on appreciating the "difference between appearance and reality"—through a metacognitive capacity to reflect on "internal representations" and hold them to account (110). Such conscious reflection enables the judging subject to "take the existential and ethical commitment of being bound by and deferring to the "world as world" (xiii).

Notably, Smith's embrace of data-driven "reckoning" contrasts with SP's focus on the dangers and harms of LLMs. At the same time, however, Smith's judging subject has much in common with the communicative human individual of SP (e.g., intentionality, situatedness, and the ability to "register" the world or communicate about it in a meaningful way). By making *judgment* the normative standard for ethical decision-making, Smith explicitly centers human subjectivities. His epistemological framework, moreover, is grounded in what the philosopher and physicist Karen Barad (2007: 46) calls *representationalism*: "belief in the ontological distinction" between *subjects* (those who represent), *knowledge* (the content of what those subjects represent), and *the world* (the purported referents of those representations).³⁴ This "representationalism" complicates Smith's break from Cartesian assumptions in that the strong distinction between subject and world is itself a kind of Cartesian premise. With little attention to the material or social conditions that promote judgment—a high moral standard that few people will develop as a matter of course—Smith's criteria are robustly ethical, but not especially political.

At the same time, Smith's ontology would look very different if it explored the ethical potential of more distributed and materialized assemblages, or if it centered the entanglement of human and nonhuman infrastructures. Such an approach might allow for the complex embeddedness of ML systems as a form of embodiment and enworlding from which something like judgment might possibly develop.³⁵ "Judgment" so conceived might not be the same as Smith's human and subject-centered ideal, but it would allow for how norms are materialized in systems in which humans and machines jointly participate.

As our posited ontological spectrum moves toward more-than-human approaches, we turn to Wendy Hui Kyong Chun's (2021) political critique of "AI," *Discriminating Data: Correlation, Neighborhoods, and the New Politics of Recognition*. Chun's book, reviewed in this issue by James Smithies, works through cogent discussion of material conditions. One of the book's highlights is its attention to

the long and ongoing legacy of eugenics. As Chun shows, statistical practices are never merely mechanical or neutrally epistemological. Rather, they embed histories of discrimination and prejudice “not simply at the level of data, but also at the levels of procedure, prediction, and logic” (16). Although Chun’s approach to data-driven systems tends to decenter human subjects, her materialist ontology retains a strong emphasis on ethical and political practices. “Big data is arguably the bastard child of psychoanalysis and eugenics,” she writes; nonetheless, “data analysis” can “foster ways to inhabit our world less destructively” (26).

In the present special issue, Lavigne’s “Scrapism: A Manifesto” works in a comparable vein, exploring the material conditions of meaning production in systems that include humans while paying significant attention to the ontoepistemological effects of nonhuman objects and practices. Lavigne explores the political implications of a data-saturated world and how they are shaped and enabled by the material conditions of the web. HTML, for instance, the protocol that links the web together, was fundamental to the utopian claims of the internet as a decentralized, democratic, and open network. Yet, the same open framework accommodated centralizing search engines and privatizing data scraping. That said, Lavigne’s *scrapism* looks beyond a politics that merely lays out this condition: his “counterpractice of web scraping” uses the data infrastructures created for private profit for “artistic, emotional, political and critical ends.” In this way, Lavigne reimagines the technoutopianism of the early web as a political, cultural, and infrastructural commons that interrupts and redeploys the material enclosures and power asymmetries of the bureaucratic and capitalist web.

“The Photographic Pipeline of Machine Vision: or Machine Vision’s Latent Photography Theory,” Malevé and Sluis’s account in this issue of the ImageNet dataset (and their concluding discussion of new “generative” models), is comparably materialist. By arguing that a theory of photography is “latent” but nonetheless tangibly embedded in the tools, practices, and discourses of machine vision research, Malevé and Sluis decenter the intentions of machine vision researchers; they explore how diverse technocultural formations—photographic snapshots, social media platforms, search engines, and algorithmic gig work—have shaped what and how photography came to mean in the crucial years that gave rise to the DL paradigm. Curation as they conceive it is a medley of more-than-human techniques and affordances—a pattern that appears to persist (and may even be radicalized) in the emerging political economy of proprietary image models.

For Sasha Costanza-Chock, communities of practice are an agency for enactment of design justice principles. Design by its very nature, she emphasizes, involves sociotechnical systems that humans never wholly anticipate, organize, or control. Hence, in the interview published in this issue, Costanza-Chock highlights the inadequacy of “good intentions,” as when designers set out to create “AI for good.” Such projects assume that human designers can predict and control every outcome; but in a multiplistic world, community-led approaches offer the best possibility for ensuring that design process and practices work on behalf of affected people and places.

Of course, approaches at this more-than-human end of the spectrum have their own set of challenges, including the complications of concrete political practices. Put simply, it is easier to localize responsibility and to envision political, legal, and/or institutional redress when one conceives harms as the relatively discrete products of a given subject, entity, or institutional formation, as opposed to accentuating distributed agencies across sociotechnical assemblages. In *Cloud Ethics*, Louise Amoore (2020: 58) illustrates the political conundrum of a strong turn toward distributed agency when she argues (in a discussion of algorithmic systems) that there is “no meaningfully unified locus of control.” Instead, in “every singular action of an apparently autonomous system” there “resides a multiplicity of human and algorithmic judgments, assumptions, thresholds, and probabilities” (64). Amoore’s point is not to obviate the ethical imperative to redress injustice. Rather, for scholars who share her strong turn toward a distributed and more-than-human ontology, the reluctance to localize blame or theorize causal connections stems from the effort to center the complex particularities of and relations among communities, technologies, and other agencies (which might be governments, corporations, regulatory bodies, universities, venture capitalists, or the “seeing” market).

We might imagine Amoore's ontological approach as the dialectical counterpart to Smith's subject-centered judgment: a very different but still exacting ideal that Amoore premises on the normative embrace of "opacity, partiality, and illegibility." In *Cloud Ethics*, this radical epistemic humility pertains not only to ML systems but to "all forms of giving an account," human as well as algorithmic (8). Thus, whereas Smith's (2019: 147) judging subject authorizes the "ontological commitments, and epistemic practices that allow" humans to "go to bat for the world as world," Amoore's (2020: 8) ontological stakes are the "ungrounded politics of all forms of ethical relations." Illegibility and ungroundedness are certainly challenging positions to defend when much of the controversy about data-driven ML concerns the vast scale and undocumented condition of many datasets. But Amoore's commitment to tracing particular algorithmic relations nonetheless has something in common with the principled refusal, as in Bender and colleagues, to turn "scaling up" into an end in itself when doing so obscures differences that matter.

Aware of the political challenges of a more-than-human framework, Costanza-Chock shifts from an ontological discussion of *creating* particular realities to an ethico-onto-epistemology that asks what can and should be done when *participating in, inhabiting, and negotiating plural realities with varied and particular sociotechnical effects*. This normative approach enables Costanza-Chock to uphold the practices necessary for political change in diverse worlds. Citing Escobar's *Designs for the Pluriverse*, she recalls us to the one-world system that settler colonialism tried to impose on Indigenous peoples. But Indigenous people continue to cultivate their own worlds (spaces, relations, ways of knowing, and knowledge systems). Indigenous studies scholars thus urge scholars of the "ontological turn" to recognize the agent ontologies in Indigenous traditions and scholarship and to adjust their citation practices in this context (Rosiek, Snyder, and Pratt 2020).

At both ends of the spectrum, a practice of critical AI studies should bear in mind how emerging technologies, no matter how cutting-edge, take shape from within long and ongoing (but also uneven) histories. Although many academics continue to focalize predominantly Western perspectives and contexts, in actuality, multiple histories constitute the pluriverse. Accepting the equal reality of different ways of knowing and being allows for a pluriverse that recognizes the contingency of any singular proposition on relationalities of human/nonhuman, nature/culture, and writings/worlds.³⁶

Conclusion: Data Worlds/Data Commons

We began this introduction to "Data Worlds" with the observation that the elite players now shaping public perception through discourses of "responsible" and "human-centered" AI are keen to downplay how the profitable accumulation of unconsented, undocumented, and nontransparent data has subsidized a worrisome concentration of power and resources. As we complete this essay, the *Financial Times*—in an article titled "Risk of 'Industrial Capture' Looms over AI Revolution" (Murgia 2023)—reports that a "handful of individuals and corporations now control much of the resources and knowledge in the sector—and will ultimately shape its impact on our collective future." Meanwhile, the same companies aggressively promoting undertested "generative" technologies are laying off their teams for "ethical" and "responsible" AI (e.g. de Vynck and Oremus 2023). As Big Tech responds to so-called activist investors by firing these and many other workers (e.g., Teal 2023), we discern little evidence that this downsizing trend will reduce the lip service paid to ethical discourses.

Alert to this problematic political economy, we have suggested that a focus on *datafication*—the processes that render objects, practices, and conditions legible *as* data—offers an approach to ongoing technocultural developments that helps to elude the pitfalls of data positivism, data universalism, and unintentional criti-hype. The more that *datafication* appears to stand for an "effective" plenitude, we have argued, the greater the need to ask: "Whose goals count?" As humanists, social scientists, artists, and community partners increasingly join technologists and policymakers in the study of so-called AI, we have urged them to think about their ontological commitments (implicit or explicit) while connecting with the research and practice of others. Inspired by Whittaker's call for "muscles of care and mutual accountability," we look to creative solidarities, design justice principles, and critical dialectics to inspire and strengthen communities of practice—intellectual, political, cultural, and technological—that forge pathways to new knowledge and enlarged spaces of dissensus.

To be sure, plenty of evidence suggests that the reified data worlds now celebrated as “AI” are dysfunctional, antidemocratic, and unscientific. As Chun (2021: 243) writes, “Machine learning programs are not only trained on selected, discriminatory, and ‘dirty’ data; they are also verified as true only if they reproduce these data” (cf. O’Neil 2016; Noble 2018). To avoid enabling datasets that “foreclos[e] the future,” Chun (2021: 254) recommends that ML systems be understood as “spaces for political action.” In particular, by treating climate models as the template for how ML might be used in the public interest, the technology’s impact could shift from authorizing predictable futures to helping to “change the future” (26).³⁷

If there is any other single aspiration that strikes us as formative for energizing these interdisciplinary projects of (in Williams’s words) activating “certain projects toward the future,” it is the notion of a *data commons*. At one level the term denotes the need for public access to the datasets on which “AI” systems are trained (along with clear information on environmental footprint and labor practices). At another, the *commons* marks a capacious and open-ended alternative to the data worlds now enclosed by—and deployed in the interests of—a concentrated elite. The notion of the commons takes us back to the prehistory of capitalism, which, according to Marx, emerged through a centuries-long process of *primitive accumulation* through which customary use rights were rescinded and common land enclosed throughout the British Isles. The result was to proletarianize the mass of laborers, creating a vast underclass of wage workers whose surplus value could be readily expropriated by the owners of the means of production. From “this original sin,” writes Marx ([1867] 1990: 873), “dates the poverty of the great majority” and the wealth of a few.³⁸

One of the most interesting passages in Zuboff’s (2019: 98) book looks back to this history to offer primitive accumulation as the template for Google’s monetization of its users’ “*personal information*.”³⁹ Like early-modern capitalists, Google took objects and practices that had “live[d] outside the market sphere” and “declare[d] their new life as market commodities” (98). Zuboff quotes Google cofounder Larry Page, who proclaimed (as if unaware of dystopic and grandiose implications) that “everything you’ve ever heard or seen or experienced will become searchable” (98). As we know, Zuboff ties this high-tech expropriation to behavioral manipulation. But theorized differently, today’s “primitive accumulation” hews closer to what Lavigne describes as the enclosure of collectively wrought databases behind digital paywalls. Indeed, Alvarado and Humphreys (2017: 739) explicitly compare data scraping for proprietary models to “the successive periods of agricultural enclosure. . . when common land was enclosed by private landowners,” while Kyle Booten (2023) describes the promotion of advanced LLMs as an effort to “proletarianize” human writing.⁴⁰ Most recently, Costanza-Chock (2023) argues that since “generative AI” is “trained upon vast datasets of centuries of human creative and intellectual work,” the systems should “belong to the *commons*, to all humanity, rather than to a handful of powerful for-profit corporations” (emphasis added). Proposals to rein in this privatization of the commons will fail, she adds, “unless they challenge the underlying appropriation of the aggregated fruits of intellectual labor” in “the service of Capital.”

To be clear, Costanza-Chock does not advise waiting for mass movements to seize the means of prediction. While the unconsented “use of the fruits of human creative and intellectual labor” calls upon people to create a new commons, potential “stopgap” measures in the form of copyright lawsuits are, she notes, already under way. And while two laws now under consideration (the Algorithmic Accountability Act in the United States and the Artificial Intelligence Act in the European Union) are discernibly products of a neoliberal paradigm, they can nevertheless reduce harms in advance of that more “radical imagination that we need.”⁴¹

In a comparable vein, literary studies—from Williams’s classic study of enclosure in *The Country and the City* to Carolyn Lesjak’s *The Afterlife of Enclosure*—have explored how both primitive accumulation and its aftermath are always culturally mediated. Lesjak likens agricultural enclosure to a kind of *slow violence*. Originally conceived by Rob Nixon (2011: 2) to address the dilemma of climate change, the term describes “a violence that occurs gradually and out of sight, a violence of delayed destruction that is dispersed across time and space, an attritional violence that is typically not viewed as violence at all.” The abstractive process of “reducing and turning things into data” has itself been likened to a form of “violence” that turns people “into a form that can fit into sensors, servers, and

processors” (Sadowski 2020: 59).⁴² Nonetheless, as Lesjak (2021: 3) writes, “the historical trauma that was enclosure” was endowed with an enduring afterlife—a “utopian spirit” of the commons that continues to inflect the radical imagination.

Inspired by Lesjak’s idea, we sign off with the conviction that the enclosed data worlds described in this introduction are comparably imbued with an afterlife of the commons. The “utopian spirit” we have in mind is not a naive technolibertarianism that celebrates spaces of nonregulation or upholds “AI” as the solution to every problem (cf. Noble and Roberts 2019). Rather than rein in harm or animate plurality, naive notions of freedom abet the instrumentalization of power, the atomization of people and communities, and the self-universalization of narrow progress. As Costanza-Chock (2023) puts it, “If we continue to allow the profit motive” to “shape” and “harness ‘AI systems’”—whether through automated decision-making, “generative AI,” or futuristic mythologies of human-level AI—“they will contribute to injustice and, ultimately, ecological collapse.” If instead “we allow ourselves to dream larger, and imagine how to place AI systems forever in the commons, with shared governance and shared goals of a just transition to a regenerative economy, then we might be able to survive as a species and live well, in right relations with Earth” (Costanza-Chock 2023).

We hope that even those readers who disagree have found ways of “living differently” in these virtual pages, and will find even more in the articles and reviews that follow.

Notes

1. On the many female “computers” that worked for companies such as Sears Roebuck, and on the female computer programmers who originally populated the field, see Hicks (2017) and Ensmenger (2021).
2. For a more recent public communication that avoids the d-word, see Google CEO Sundar Pichai’s (2023) announcement about “AI” on the occasion of the company’s release of its new Bard chatbot. It is worth noting that at least some mainstream discourse looks past the radio silence on data accumulation, as when *Time*’s reporter refers to Google in passing as a “surveillance empire” (Perrigo 2023a).
3. Notably, Sandberg worked for McKinsey before joining Google and Facebook.
4. On Li’s “curatorial pipeline,” see Malevé and Sluis in this issue; on Amazon’s interface for on-demand gig labor, see Gray and Suri (2019). Krizhevsky, Sutskever, and Hinton’s landmark 2012 article, “ImageNet Classification with Deep Convolutional Neural Networks,” claimed to reduce the error rate for correctly classifying objects in images by 18 percent.
5. For the original article, see McCulloch and Pitts (1943). For an informative piece with a telling headline, see Hardesty (2017): “Explained: Neural Networks, Ballyhooed Artificial-Intelligence Technique Known as ‘Deep Learning’ Revises 70-Year-Old Idea.”
6. DL enthusiasts such as Sejnowski (2018) persistently liken the “learning” of neural architectures to that of “babies” (3) while claiming that AI “was made by reverse engineering brains” (ix)—anthropomorphizing claims that rest on loosely metaphorical thinking.
7. On Whittaker’s organizing efforts as a Google employee see, e.g., Scheiber and Conger (2020). As cofounder and director of the AI Now Institute, Whittaker—as recounted in Sadowski and Phan (2022: 148)—was “warned by the NYU engineering school . . . to align our views. . . . In early 2021, we were told that we didn’t align with the strategy of the school and that we needed to find another home. Then, in late 2021, NYU informed us that they were going to take all of our gift money, leaving us with the choice of litigating, which is expensive and which AI Now, as a part of NYU, can’t pay for, or of walking and trying to make up the over four million dollars that was effectively stolen by the university.” For a recent report on the continuing concentration of power in AI technology, see Kak and West (2023).
8. See, for example, Barad (2007); Bennett (2010); Clark (2016); Mol (2003); Holbraad and Pederson (2017); Adema (2021). Some readers may be surprised to find this concerted blending of insights drawn from what is often called the “new” materialisms with those of historical and dialectical materialisms that predate the latter by at least several decades. We regard such alliances as timely, intellectually productive, and in some instances socially urgent. For a cogent example of this dialectical blending in the sphere of ecocriticism, see Moore (2015); for its elaboration as a mode of critique after so-called postcritique, see Coundouriotis and Goodlad (2020). For an ontological elaboration of historical materialism, see Bhaskar and Callinicos (2003); for a complementary feminist perspective, see Haslanger (1995).
9. See also McKelvey (2023: 137), who describes the disappointing usages of AI technologies thus far and ponders whether the “tasks for a critical AI studies might be to reimagine [the technology’s] political function.” Raley and Rhee (2023), a special issue of *American Literature* devoted to critical AI studies, appeared too late for us to incorporate its findings into this discussion. Here we note our enthusiasm for the interdisciplinary approaches described in the abstract and table of contents and our plans in future issues of *Critical AI* to engage the work substantively.

10. See also Manovich (2012) on the hierarchies of power between those whose data can be collected, those who have the means to do that collecting, and those who can analyze it.
Working from a similar perspective, Crawford (2021: 114, 93–94) describes *datafication* as “the process through which information about people becomes part of an “aggregate mass” that drives a “broader system,” the ostensible justification for which is improved “technological performance.”
11. Newfield in this issue describes Anderson’s account as “postepistemological.” Notably, the same article prompted Moretti (2017: 687), a pioneer in so-called distant approaches to studying literature, to reconsider the gains of data-driven methods: “Fact is,” he wrote, “big data has produced a decline in theoretical curiosity, which, in its turn, has made our results often mediocre. ‘More’ data is not the solution here. . . . Only a resolute return to theory may change the way we work.” See also Goodlad (2020).
12. See Starmans (2016: 4) for an account of Google’s essay as the “Federalist Papers of the current data era.” According to Zuboff (2019: 188), Google’s access to the “most data” gave it an edge in the quality of their ML programs. “*Hyperscale*,” she explains, is an industry term pioneered by Google at a time when the company was considered to possess “the largest computer network on Earth.”
13. See also the original essay, Wigner (1960: 9), which concludes as follows: “The miracle of the appropriateness of the language of mathematics for the formulation of the laws of physics is a wonderful gift which we neither understand nor deserve. We should be grateful for it and hope that it will remain valid in future research and that it will extend, for better or for worse, to our pleasure, even though perhaps also to our bafflement, to wide branches of learning.”
14. Benjamin (2019) refers to this deeply entrenched racism as a new “Jim code.” See also West, Whittaker, and Crawford (2019), Chun (2021), and Whittaker (2023).
15. See also West (2020), who shows the long-standing ties between Silicon Valley and eugenicist ideology; Goodlad (2020, 2022), who ties the eugenicist origins of statistical correlation to AI’s comparable debts to the utilitarian abstraction of *homo economicus*; Mhlambi (2020), who anchors his argument in a counterdiscourse of Ubuntu; and Hu (forthcoming), who calls for a “thick constructivist” understanding of “race” in evaluation of racial discrimination in algorithm systems.
16. On the “datafied society,” see Schäfer and van Es (2017).
17. Van Dijck, Poell, and de Waal (2019) build on the earlier work of Mayer-Schönberger and Cukier (2013). Note that Newfield, Alexandrova, and John (2022) define *numeracy* in similar terms.
18. Indeed, the legacy of those older norms and regulatory frameworks deters mobile phone providers from harvesting the (now entextualizable) content of their customers’ conversations even as email providers have long helped themselves to the contents of (textualized) emails as a matter of course.
19. As Sadowski argues (2019: 2–3), with the work of boyd and Crawford (2012) partly in mind, the assumption that “everything is data” is not a “neutral observation” about the world, but a means of ordering and constructing it. “It is not a coincidence that data is treated as a universal substance right at the time when there is so much to gain for whoever can lay claim to that data and extract it from every source.”
20. Building on Gitelman (2013) as well as Thatcher, O’Sullivan, and Mahmoudi (2016), Sadowski (2019: 3) argues that “data *manufacturing*” is more apt than “data *mining*” because the framing of data as preexisting like “crude oil or raw ore . . . reinforces regimes of data accumulation,” while the “goal of transforming everything into data and the search for new sources of data and the search for new sources of data echoes imperialist modes of accumulation.”
21. Though space does not permit fuller discussion, a key context for theorizing the exploitative features of data surveillance is the violation and/or transmutation of privacy and its norms, a topic on which Zuboff resonates despite her focus on behavioral manipulation. On privacy norms, see, for example, Nissenbaum (2019) on the “food chain” for “data-ravenous” ML. On the questionable effectiveness of microtargeting see, for instance, Resnick (2018). For a book-length alternative to Zuboff’s approach to surveillance capitalism, see Doctorow (2020). Burgess et al. (2022: chap. 1) discuss Zuboff in light of a variation on Vinsel’s (2021) criti-hype that they call “Big Critique.” While we appreciate their attention to everyday cultures and share their concerns about critical hyperbole, we question “Big Critique”—a potential overcorrection given the importance of the solidarities Whittaker urges.
22. In Bourdieu and Wacquant’s (2013: 297) own words, “Any difference [between people] that is recognized, accepted as legitimate, functions by that very fact as a symbolic capital providing a profit of distinction. Symbolic capital, together with the forms of profit and power it warrants, exists only in the relationship between distinct and distinctive properties, such as the body proper, language, clothing, interior furnishings (each of which receives its value from its position in the system of corresponding properties).”
23. One example is credit scores. Fourcade and Healey choose the Nietzsche-inspired prefix “uber” to denote their version of a new mode of symbolic capital because they conceive “ubercapital” as a “metageneralized or transcendent” kind of capital that functions as an “index of superiority” identifying those who “who stand above the world and others in it” (14).
24. See Raji et al. (2022) on the vast underattention to the simple fact that “many deployed algorithmic products do not work.” For specific analysis of problematic decision-making systems that optimize “predictive accuracy,” see Wang et al. (2022). On the state of Michigan’s deployment of a dysfunctional system that falsely accused more than thirty thousand applicants for unemployment benefits of fraud, see Charette (2018).

25. As an example of data's agency that long predates the turn to data-driven DL, consider Foucault's *Discipline and Punish*. In Foucault's ([1977] 1995: 190) account of the disciplinary "innovations" that the centralizing of "individual data into cumulative systems" made possible in the early modern era, the data in question clearly exerts a world-shaping effect (adumbrating the shape of digital computing in doing so). No doubt differences of scale, computational power, and much else in the last two centuries have substantially transformed what data is and does; but that does not mean data was not an active "material agent" in the past. As we have seen, Price urged data collection precisely because he anticipated its material power. Notably, Halpern et al. (2022: 199) describe a longer durée in other parts of their essay, as when they suggest (rightly in our view) that the data practices of the present day are "[h]aunted by the histories of colonialism, race, and population in which statistics had its modern origins." At various points, the essay asserts a nuanced comparatism and historicism that could temper the absolutizing claims about data positivism.
26. On algorithmic auditing see, for example, Costanza-Chock, Raji, and Buolamwini (2022) and Broussard (2018: 148), which argues that algorithmic auditing and associated "accountability reporting are two strains of public interest technology that . . . show the most promise for remedying algorithmic harms along the axes of race, gender, and ability."
27. The term *critical realism* is used comparably by Bhaskar and Callinocos (2003), where the emphasis on structures, multivalent processes, and historical contingencies (also) tends to challenge simple notions of intentionality, individuality, and self-determination.
28. The authors argue that LLMs risk environmental harms, documentation debt, harm to marginalized communities, and underinvestment in alternative language technologies. They recommend assessing environmental costs first, using only documented and consented data, incorporating diverse people into the design process, and encouraging investment in multiple language technologies.
29. Readers familiar with Saussure ([1916] 2011) might interpret this distinction in terms of a language model's access to *signifiers* without any comparable access to *signifieds*.
30. *OED Online*, s.v. "parrot," <https://www.oed-com.virtual.anu.edu.au/view/Entry/138145?rskey=IU6vJ9&result=3&isAdvanced=false> (accessed March 15, 2023). The verb form of "parroting" has carried these associations since at least the seventeenth century. SP's metaphor is somewhat contradictory in describing the stochastic parrot's generation of text as both "haphazard" (random) and "probabilistic" (organized accord to statistical modeling).
31. "Feeling lonely," one day O inserts himself in place of B without A's knowledge, thus leaving A to supplement messages presumed to be from B with "guesses about B's state of mind," "goals," and "communicative intent." "It is not that O's utterances make sense," the authors write, "but rather, that A can make sense of them." When A asks O to help her construct a weapon to fend off an "angry bear," she discovers that O "has no idea about what A 'means'" (5188–89).
32. On the ELIZA effect, named after a rudimentary chat program created by Joseph Weizenbaum in 1966, see Weizenbaum (1976) and Hofstadter (1996). As Weizenbaum (1976: 11) wrote with respect to the response of many humans to this seemingly intelligent software, the experience showed that "extremely short exposures to a relatively simple computer program" can "induce powerful delusional thinking in quite normal people."
33. What Smith identifies as first-wave AI (also known as "GOFAI," for "good old-fashioned AI") is "symbolic" in enlisting coded rules, taxonomies, and formal logics in the effort to reproduce knowledge of the world and humanlike reasoning about it. By contrast, second-wave AI is trained on massive datasets to *infer* such knowledge and reasoning through statistical modeling. First-wave AI, according to Smith (2019: 20), built "interconnected network[s] of discrete symbols or data structures, organized in roughly propositional form, operating according to a regimen dictated by an equally symbolic program." Such "symbolic" architectures continue to underlie many databases, systems for digital record-keeping, and programs such as Microsoft Word.
34. The same assumption underwrites the transmission of "implicit meaning" between human individuals in the articles by Bender and colleagues. In contrasting the latter articles to Smith's book, it is important to recognize that Smith's analysis of data-driven ML seldom addresses language models per se and is largely premised on (smaller) machine vision systems that preceded the advent of GPT-3 and that thus precede the move—memorialized in Stanford's Center for Research on Foundation models (e.g., Liang 2022)—to single out these large systems as exceptional "foundations" for the future of "AI" research. Hence, one should not exaggerate the differences between Smith's book and the two articles, which boil down to 1) Smith's relative enthusiasm for the epistemological gains of data-driven ML and 2) his sympathy for the "AI" project writ large. In other respects, Smith's distinctions between reckoning and judgment harmonize with SP, as when he (2019: 126) argues that no "disconnected reckoner" that does not "participate" in the world is likely to meet the standard for judgment. In one of the rare instances of Smith's addressing a system that outputs language, he writes that such ostensible "use" of natural language "can obscure" the lack of humanlike understanding (110). He also makes clear that his enthusiasm for ML depends on the design of the system in question (implying, perhaps, that he would affirm Bender et al.'s call for documentable datasets). "Contrary to current fashion," he writes, "the mere ability to statistically combine data assembled from diverse sources without exercising judgment at each and every step may actually militate against, rather than for, anything we should want to call intelligence" (90). Notably, for Bender et al., questions as to how the recommended turn to

- smaller and documentable datasets will impact something called “AI” are never at issue since “AI” as such is never upheld as the relevant research domain.
35. Bender et al. would agree, we think, with Smith’s related pronouncement, that data is the “primary locus of bias in machine learning” (2019: 67n16), but for inverse reasons: while he thinks human systems and categorizations introduce bias into ML systems, they understand human curation and documentation as potentially key ways of ameliorating harm.
 36. Costanza-Chock’s inspiration, the indigenous Zapatista separatist movement in Mexico in the 1990s, resonates with a presentation in our Data Ontologies series. The event featured Angie Abdilla, Adam Goodes, and Baden Pailthorpe, who were interviewed by Genevieve Bell about their “Tracker Data Project,” which interrelates an enormous cache of biometric data (created during Goodes’s career as one of Australia’s most celebrated and racially vilified Australian Football League players) with Goodes’s North Wind-Ararru kinship system through a 3D scan of a sacred ancient ancestral *wirra*—red river gum tree—and audio recordings of creation stories translated into the sound of the wind using ML. These interventions move far beyond the opposition of human and nonhuman through creative participation in First Nations Australian ontologies that center “Country” and kinship—an ethico-onto-epistemic merger of people, place, and history. As Escobar (2017: 20) points out, the ontological turn did not invent the more-than-human.
 37. See also D’Ignazio and Klein (2020: 8), who ask: “How can we use data to remake the world?” The problematic use of data in a “flawed history,” they argue, “does not mean ceding control of the future to the powers of the past. . . . The power of data can be wielded back” (17). Such foregrounding of alternatives recalls Williams’s appeal to activate “certain projects towards the future.”
 38. Compare to recent reporting on the “vast new underclass” or low-paid and exploited workers now being enlisted to provide human feedback at industrial scale to sustain the illusion of an automatic technology (Wong 2023; Perrigo 2023b).
 39. Zuboff cites Arendt (1951) and Harvey (2005), both of whom expand Marx’s notion of primitive accumulation as a one-time world-historical event to entail recurrent phases, often tied to projects of imperial domination. See also Harney and Moten (2013) for a theorization of the *undercommons* that focalizes the ongoing effects of Atlantic slavery and looks beyond academic discourse.
 40. Andrejevic (2007: 297) has proposed *digital enclosure* to describe the surveillant business model attached to “the misleading image of the internet cloud,” defining such enclosure as a “relationship between a material, spatial process” and “the private expropriation of information.”
 41. Similarly, a stopgap for the “psychologically harmful work” of tagging toxic content should involve better pay and conditions; in the longer run, better tools and processes could “ease the human burden of this work” (Costanza-Chock 2023)—processes, we might add, that could involve the use of documented (smaller) datasets.
 42. As Sadowski elaborates (2020: 59), “datafication can be dehumanizing: it reduces people to abstractions about their attributes and associations, which are atomized, accumulated, analyzed, and administered. . . . It enables smart systems in various sectors of society to profile, sort, stratify, score, rank, reward, punish, include, exclude, and otherwise calculate decisions that determine what sociologists call our ‘life-chances.’”

Works Cited

- Abebe, Rediet, and Maximilian Kasy. 2021. “The Means of Prediction.” *Boston Review*, May 20. https://www.bostonreview.net/forum_response/the-means-of-prediction/.
- Adema, Janneke. 2021. *Living Books: Experiments in the Posthumanities*. Cambridge, MA: MIT Press.
- Alvarado, Rafael, and Paul Humphreys. 2017. “Big Data, Thick Mediation, and Representational Opacity.” *New Literary History* 48, no. 4 (Autumn): 729–49.
- Amoore, Louise. 2020. *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Durham, NC: Duke University Press.
- Anderson, Chris. 2008. “The End of Theory: The Data Deluge Makes the Scientific Method Obsolete.” *Wired*, June 23. <https://www.wired.com/2008/06/pb-theory/>.
- Andrejevic, Mark. 2007. “Surveillance in the Digital Enclosure.” *Communication Review* 10, no. 4: 295–317.
- Arendt, Hannah. 1951. *The Origins of Totalitarianism*. New York: Harcourt.
- Barad, Karen. 2007. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Durham, NC: Duke University Press.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜” In *FAccT ’21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–23. New York: Association for Computing Machinery.
- Bender, Emily M. and Alexander Koller. 2020. “Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data.” In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–98. <https://aclanthology.org/2020.acl-main.463>.
- Benjamin, Ruha. 2019. *Race after Technology: Abolitionist Tools for the New Jim Code*. Cambridge, MA: Polity.
- Bennett, Jane. 2010. *Vibrant Matter: A Political Ecology of Things*. Durham, NC: Duke University Press.

- Bhaskar, Roy, and Alex Callinicos. 2003. "Marxism and Critical Realism: A Debate." *Journal of Critical Realism* 1, no. 2: 89–114.
- Booten, Kyle. "Build Word-Gyms, Not Word-Factories." Paper presented at the conference "AI Futures: An Interdisciplinary Conversation on Large Language Models and the Future of Human Writing," Rutgers University, New Brunswick, NJ, February 16.
- Bourdieu, Pierre, and Loïc Wacquant. 2013. "Symbolic Capital and Social Class." *Journal of Classical Sociology* 13, no. 2: 292–302.
- boyd, danah, and Kate Crawford. 2012. "CRITICAL QUESTIONS FOR BIG DATA: Provocations for a Cultural, Technological, and Scholarly Phenomenon." *Information, Communication, and Society* 15, no. 5 (June): 662–79.
- Broussard, Meredith. 2018. *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, MA: MIT Press.
- Burgess, Jean, Kath Albury, Anthony McCosker, and Rowen Wilken. 2022. *Everyday Data Cultures*. Cambridge, MA: Polity.
- Charette, Robert N. 2018. "Michigan's MiDAS Unemployment System: Algorithmic Alchemy Created Lead, Not Gold." *IEEE Spectrum*, January 24. <https://spectrum.ieee.org/michigans-midas-unemployment-system-algorithm-alchemy-that-created-lead-not-gold#toggle-gdpr>.
- Chen, Angela. 2018. "How Computers Misunderstand the World." *The Verge*, May 23. <https://www.theverge.com/2018/5/23/17384324/meredith-broussard-artificial-unintelligence-technology-criticism-technochauvinism>.
- Chun, Wendy Hui Kyong. 2021. *Discriminating Data: Correlation, Neighborhoods, and the New Politics of Recognition*. Cambridge, MA: MIT Press.
- Ciston, Sarah. 2023. *A Critical Field Guide for Working with Machine Learning Datasets*. Edited by Mike Ananny and Kate Crawford. Knowing Machines. <https://knowingmachines.org/critical-field-guide> (accessed June 17).
- Clark, Andy. 2016. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford: Oxford University Press.
- Costanza-Chock, Sasha (@schock). 2023. "Generative AI systems are trained upon vast datasets of centuries of human creative and intellectual work [. . .]." Twitter thread, March 26, starting at 12:15 p.m. <https://mobile.twitter.com/schock/status/1640024767704227840>.
- Costanza-Chock, Sasha, Inioluwa Deborah Raji, and Joy Buolamwini. 2022. "Who Audits the Auditors? Recommendations from a Field Scan of the Algorithmic Auditing Ecosystem." *Proceedings of the 2022 ACM Conference*, 1571–83. Seoul: ACM.
- Coundouriotis, Eleni, and Lauren M. E. Goodlad. 2020. "What Is and Isn't Changing?" *Modern Language Quarterly* 81, no. 4: 399–418.
- Crawford, Kate. 2021. *Atlas of AI*. New Haven, CT: Yale University Press.
- Croak, Marian, and Jen Gennai. 2023. "Responsible AI: Looking Back at 2022, and to the Future." *Google: The Keyword*, January 11. <https://blog.google/technology/ai/responsible-ai-looking-back-at-2022-and-to-the-future/>.
- de Castro, Eduardo Viveiro. 2015. "Who's Afraid of the Ontological Wolf: Some Comments on an Ongoing Anthropological Debate." *Cambridge Journal of Anthropology* 33, no. 1: 2–17.
- Dejean, Mathieu, and Jean-Marc Lalanne. 2021. "Jacques Rancière: 'The Issue Is to Manage to Maintain Dissensus.'" Verso Book Club, August 10.
- Denton, Emily, Alex Hanna, Razvan Amironesei, Andrew Smart, and Hilary Nicole. 2021. "On the Genealogy of Machine Learning Datasets: A Critical History of ImageNet." *Big Data and Society* 8, no. 2 (July): 1–14.
- de Vynck, Gerrit, and Will Oremus. 2023. "As AI Booms, Tech Firms Are Laying Off Their Ethicists." *Washington Post*, March 30.
- D'Ignazio, Catherine, and Lauren F. Klein. 2020. *Data Feminism*. Cambridge, MA: MIT Press.
- Doctorow, Cory. 2020. *How to Destroy Surveillance Capitalism*. New York: Medium Editions.
- Escobar, Arturo. 2017. *Designs for the Pluriverse: Radical Interdependence, Autonomy, and the Making of Worlds*. Durham, NC: Duke University Press.
- Ensmenger, Nathan. 2021. "The Cloud Is a Factory." In *Your Computer Is on Fire*, edited by Thomas S. Mullaney, Benjamin Peters, Mar Hicks, and Kavita Philip, 29–49. Cambridge, MA: MIT Press.
- Eubanks, Virginia. 2017. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's.
- Foster, John Bellamy, and Robert W. McChesney. 2014. "Surveillance Capitalism: Monopoly-Finance Capital, the Military-Industrial Complex, and the Digital Age." *Monthly Review* 66, no. 3 (July/August): 1–31.
- Foucault, Michel. 1995. *Discipline and Punish: The Birth of the Prison*. Translated by Alan Sheridan. New York: Vintage Books.
- Fourcade, Marion, and Kieran Healy. 2017. "Seeing Like a Market." *Socio-Economic Review* 15, no. 1 (January): 9–29.
- Frenkel, Sheera, and Cecilia Kang. 2021. *An Ugly Truth: Inside Facebook's Battle for Domination*. New York: Harper Collins.
- Gitelman, Lisa, ed. 2013. "Raw Data" Is an Oxymoron. Cambridge, MA: MIT Press.
- Goodlad, Lauren M. E. 2020. "A Study in Distant Reading: Genre and the *Longue Durée* in the Age of AI." *Modern Language Quarterly* 81, no. 4: 491–525.
- Goodlad, Lauren M. E. 2022. "Victorian 'Artificial Intelligence': Or How George Eliot's Fiction Helps Us to Understand Statistical Modelling." Sally Ledger Memorial Lecture, London, April 7.
- Gray, Mary L., and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston: Houghton Mifflin Harcourt.

- Halevy, Alon, Peter Norvig, and Fernando Pereira. 2009. "The Unreasonable Effectiveness of Data." *IEEE Intelligent Systems* 24, no. 2: 8–12.
- Halpern, Orit, Patrick Jagoda, Jeffrey West Kirkwood, and Leif Weatherby. 2022. "Surplus Data: An Introduction." *Critical Inquiry* 48, no. 2 (Winter): 197–210.
- Hao, Karen. 2020. "We Read the Paper That Forced Timnit Gebru Out of Google. Here's What It Says." *MIT Technology Review*, December 4. <https://www.technologyreview.com/2020/12/04/1013294/google-ai-ethics-research-paper-forced-out-timnit-gebru/>.
- Hardesty, Larry. 2017. "Explained: Neural Networks." *MIT News*, April 14. <https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>.
- Harney, Stefano, and Fred Moten. 2013. *The Undercommons: Fugitive Planning and Black Study*. Wivenhoe: Minor Compositions.
- Harvey, David. 2005. *A Brief History of Neoliberalism*. Oxford: Oxford University Press.
- Haslanger, Sally. 1995. "Ontology and Social Construction." *Philosophical Topics* 23, no. 2 (Fall): 95–125.
- Hicks, Marie. 2017. *Programmed Inequality: How Britain Discarded Women Technologists and Lost Its Edge in Computing*. Cambridge, MA: MIT Press.
- Hofstadter, Douglas R. 1996. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York: Basic Books.
- Holbraad, Martin, and Morten Axel Pedersen. 2017. *The Ontological Turn: An Anthropological Exposition*. Cambridge, MA: Cambridge University Press.
- Hu, Lily. Forthcoming. "What Is 'Race' in Algorithmic Discrimination on the Basis of Race?" *Journal of Moral Philosophy*.
- Jameson, Fredric. 2009. *Valences of the Dialectic*. London: Verso.
- Kak, Amba, and Sarah Myers West. 2023. "AI Now 2023 Landscape: Confronting Tech Power." Presented at the AI Now Institute, April 11. <https://ainowinstitute.org/2023-landscape>.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2012. "ImageNet Classification with Deep Convolutional Neural Networks." In *Proceedings of the 25th International Conference on Neural Information Processing Systems*, 1097–1105. Lake Tahoe, NV: ACM.
- Lesjak, Carolyn. 2021. *The Afterlife of Enclosure: British Realism, Character, and the Commons*. Stanford, CA: Stanford University Press.
- Leurs, Koen, and Tamara Shepherd. 2017. "Datafication and Discrimination." In *The Datafied Society: Studying Culture through Data*, edited by Mirko Tobias Schäfer and Karin van Es, 211–32. Amsterdam: Amsterdam University Press.
- Li, Fei-Fei, and John Etchemendy. n.d. "Letter from the Denning Co-directors." HAI: Stanford University Human-Centered Artificial Intelligence. <https://hai.stanford.edu/navigate/welcome> (accessed March 13).
- Liang, Percy. 2022. "Foundation Models Have Forever Changed AI Research. In the Future, They Need to be Released Responsibly." *Protocol*, July 3. <https://www.protocol.com/enterprise/foundation-models-ai-standards-stanford>.
- Manovich, Lev. 2012. "Trending: The Promises and the Challenges of Big Social Data." In *Debates in the Digital Humanities*, edited by Matthew K. Gold, 460–75. Minneapolis: University of Minnesota Press.
- Manyika, James, Michael Chui, Brad Brown, Jacques Bughin, Richard Dobbs, Charles Roxburgh, and Angela Hung Byers. 2011. "Big Data: The Next Frontier for Innovation, Competition, and Productivity." *McKinsey Digital*, May 1. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/big-data-the-next-frontier-for-innovation>.
- Marx, Karl. (1867) 1990. *Capital: Critique of Political Economy*. Vol. 1. London: Penguin.
- Mayer-Schönberger, Viktor, and Kenneth Cukier. 2013. *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. Boston: Houghton Mifflin Harcourt.
- McAfee, Andrew, and Erik Brynjolfsson. 2012. "Big Data: The Management Revolution." *Harvard Business Review*, October. <https://hbr.org/2012/10/big-data-the-management-revolution>.
- McCulloch, Warren S., and Walter Pitts. 1943. "A Logical Calculus of the Ideas Immanent in Nervous Activity." *Bulletin of Mathematical Biophysics* 5: 115–33.
- McKelvey, Fenwick. 2021. "The Other Cambridge Analytics: Early 'Artificial Intelligence' in American Political Science." In *The Cultural Life of Machine Learning: An Incursion into Critical AI Studies*, edited by Jonathan Roberge and Michael Castelle, 117–42. New York: Palgrave Macmillan.
- Mhlambi, Sabelo. 2020. "From Rationality to Relationality: Ubuntu as an Ethical and Human Rights Framework for Artificial Intelligence Governance." Carr Center Discussion Paper Series. Harvard University, Cambridge, MA, July.
- Mol, Annemarie. 2003. *The Body Multiple: Ontology in Medical Practice*. Durham, NC: Duke University Press.
- Moore, Jason W. 2015. *Capitalism in the Web of Life: Ecology and the Accumulation of Capital*. London: Verso.
- Moretti, Franco. 2017. "Franco Moretti: A Response." *PMLA* 132, no. 3 (May): 686–89.
- Murgia, Madhumita. 2023. "Risk of 'Industrial Capture' Looms over AI Revolution." *Financial Times*, March 23.
- Newfield, Christopher, Anna Alexandrova, and Stephen John, eds. 2022. *Limits of the Numerical: The Abuses and Uses of Quantification*. Chicago: University of Chicago Press.
- Nissenbaum, Helen. 2019. "Contextual Integrity up and down the Data Food Chain." *Theoretical Inquiries* 20, no. 1: 221–56.
- Nixon, Rob. 2011. *Slow Violence and the Environmentalism of the Poor*. Cambridge, MA: Harvard University Press.
- Noble, Safiya Umoja. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.

- Noble, Safiya Umoja, and Sarah T. Roberts. 2019. "Technological Elites, the Meritocracy, and Postracial Myths in Silicon Valley." In *Racism Postrace*, edited by Roopali Mukherjee, Sarah Banet-Weiser, and Herman Gray, 113–30. Durham, NC: Duke University Press.
- O'Flaherty, Kate. 2022. "Apple's Privacy Features Will Cost Facebook \$12 Billion." *Forbes*, April 23. <https://www.forbes.com/sites/kateoflahertyuk/2022/04/23/apple-just-issued-stunning-12-billion-blow-to-facebook/?sh=1a210fd51907>.
- O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.
- Pearl, Judea, and Dana Mackenzie. 2018. *The Book of Why: The New Science of Cause and Effect*. New York: Basic Books.
- Perrigo, Billy. 2023a. "DeepMind's CEO Helped Take AI Mainstream. Now He's Urging Caution." *Time*, January 12. <https://time.com/6246119/demis-hassabis-deepmind-interview/>.
- Perrigo, Billy. 2023b. "OpenAI Used Kenyan Workers on Less than \$2 Per Hour to Make ChatGPT Less Toxic." *Time*, January 18. <https://time.com/6247678/openai-chatgpt-kenya-workers/>.
- Pichai, Sundar. 2023. "An Important Next Step on Our AI Journey." *Google: The Keyword*, February 6. <https://blog.google/intl/en-africa/products/explore-get-answers/an-important-next-step-on-our-ai-journey/>.
- Price, Richard. (1771) 1773. *Observations on Reversionary Payments*. London: Cadell.
- Raji, Deborah Inioluwa, Elizabeth I. Kumar, Aaron Horowitz, and Andrew Selbst. 2022. "The Fallacy of AI Functionality." *Proceedings of the 2022 ACM Conference*, 959–72. Seoul: ACM.
- Raley, Rita, and Jennifer Rhee, eds. 2023. "Critical AI: A Field in Formation." Special issue, *American Literature* 92, no. 2 (June).
- Resnick, Brian. 2018. "Cambridge Analytica's 'Psychographic Microtargeting': What's Bullshit and What's Legit." *Vox*, March 26. <https://www.vox.com/science-and-health/2018/3/23/17152564/cambridge-analytica-psychographic-microtargeting-what>.
- Roose, Kevin. 2023. "How ChatGPT Kicked Off an A.I. Arms Race." *New York Times*, February 3.
- Rosiek, Jerry Lee, Jimmy Snyder, and Scott L. Pratt. 2020. "The New Materialisms and Indigenous Theories of Non-Human Agency: Making the Case for Respectful Anti-Colonial Engagement." *Qualitative Inquiry* 26, no. 3–4: 331–46.
- Sadowski, Jathan. 2019. "When Data Is Capital: Datafication, Accumulation, and Extraction." *Big Data and Society* 6, no. 1. <https://doi.org/10.1177/2053951718820549>.
- Sadowski, Jathan. 2020. *Too Smart: How Digital Capitalism Is Extracting Data, Controlling Our Lives, and Taking Over the World*. Cambridge, MA: MIT Press.
- Sadowski, Jathan, and Thao Phan. 2022. "'Open Secrets': An Interview with Meredith Whittaker." In *Economies of Virtue: The Circulation of 'Ethics' in AI*, edited by Thao Phan, Jake Goldenfein, Declan Kuch, and Monique Mann, 147–59. Amsterdam: Institute of Network Cultures.
- Saussure, Ferdinand Mongin. (1916) 2011. *Course in General Linguistics*. New York: Columbia University Press.
- Schäfer, Mirko Tobias, and Karin van Es. 2017. *The Datafied Society: Studying Culture through Data*. Amsterdam: Amsterdam University Press.
- Scheiber, Noam, and Kate Conger. 2020. "The Great Google Revolt." *New York Times Magazine*, February 18.
- Seagate Technology. n.d. "The World of Data as We Know It Keeps Growing." Seagate Blog. <https://blog.seagate.com/intelligent/the-world-of-data-as-we-know-it-keeps-growing/> (accessed July 5, 2023).
- Searle, John R. 1980. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3, no. 3: 417–57.
- Sejnowski, Terrence J. 2018. *The Deep Learning Revolution*. Cambridge, MA: MIT Press.
- Smith, Brian Cantwell. 2019. *The Promise of Artificial Intelligence: Reckoning and Judgment*. Cambridge, MA: MIT Press, 2019.
- Srnicek, Nick. 2016. *Platform Capitalism*. Cambridge: Polity.
- Starmans, Richard. 2016. "The Advent of Data Science: Some Considerations on the Unreasonable Effectiveness of Data." In *Handbook of Big Data*, edited by Peter Bühlmann, Petros Drineas, Michael Kane, and Mark van der Laan, 3–20. Boca Raton: CRC.
- Tapo, Allahsera Auguste, Bakary Coulibaly, Sébastien Diarra, Christopher Homan, Julia Kreutzer, Sarah Luger, Arthur Nagashima, Marcos Zampieri, and Michael Leventhal. 2020. "Neural Machine Translation for Extremely Low-Resource African Languages: A Case Study on Bambara." *Proceedings of the 3rd Workshop on Technologies for MT of Low Resource Languages*, 23–32. Suzhou: ACL.
- Teal, Kelly. 2023. "Activist Investor Out for Blood at Google, Wants More Layoffs." *Channel Futures*, January 24. <https://www.channelfutures.com/business-models/activist-investor-out-for-blood-at-google-wants-138k-more-layoffs>.
- Thatcher, Jim, David O'Sullivan, and Dillon Mahmoudi. 2016. "Data Colonialism through Accumulation by Dispossession: New Metaphors for Daily Data." *Environment and Planning D: Society and Space* 34, no. 6 (December): 990–1006.
- Turing, A. M. 1950. "Computing Machinery and Intelligence." *Mind* 49, no. 236: 433–60.
- van Dijk, José, Thomas Poell, and Martijn de Waal. 2018. *The Platform Society: Public Values in a Connective World*. Oxford: Oxford University Press.
- Vinsel, Lee. 2021. "You're Doing It Wrong: Notes on Criticism and Technology Hype." *Medium*, February 1. <https://sts-news.medium.com/youre-doing-it-wrong-notes-on-criticism-and-technology-hype-18b08b4307e5>.
- Wang, Angelina, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2022. "Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms that Optimize Predictive Accuracy." *Social Science Research Network*: 1–45. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4238015.

- Weizenbaum, Joseph. 1976. *Computer Power and Human Reason: From Judgment to Calculation*. New York: Freeman.
- West, Sarah Myers. 2020. "AI and the Far Right: A History We Can't Ignore." AI Now Institute, May 4. <https://ainowinstitute.org/publication/ai-and-the-far-right-a-history-we-cant-ignore-2>.
- West, Sarah Myers, Meredith Whittaker, and Kate Crawford. 2019. "Discriminating Systems: Gender, Race, and Power in AI." AI Now Institute, April. <https://ainowinstitute.org/publication/discriminating-systems-gender-race-and-power-in-ai-2>.
- Whittaker, Meredith. 2021. "The Steep Cost of Capture." *Interactions* 28, no. 6: 50–55.
- Whittaker, Meredith. 2023. "Origin Stories: Plantations, Computers, and Industrial Control." *Logic(s)*, no. 19. <https://logicmag.io/supa-dupa-skies/origin-stories-plantations-computers-and-industrial-control/>.
- Wigner, Eugene. 1960. "The Unreasonable Effectiveness of Mathematics in the Natural Sciences." *Communications in Pure and Applied Mathematics* 13, no. 1 (February): 1–14.
- Williams, Raymond. 1989. "The Future of Cultural Studies." In *The Politics of Modernism: Against the New Conformists*, 151–62. London: Verso.
- Wong, Matteo. 2023. "America Already Has an AI Underclass." *Atlantic Monthly*, July 26. <https://www.theatlantic.com/technology/archive/2023/07/ai-chatbot-human-evaluator-feedback/674805/>.
- Zuboff, Shoshana. 2015. "Big Other: Surveillance Capitalism and the Prospects of an Information Civilization." *Journal of Information Technology* 30, no. 1: 75–89.
- Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.

PROOF
Duke University Press/Journals