

13 Planetary Potemkin AI: The Humans Hidden inside Mechanical Minds

Jathan Sadowski

In 1770, the Hungarian inventor Wolfgang von Kempelen unveiled the Mechanical Turk, a chess-playing contraption that “consisted of a wooden cabinet behind which was seated a life-size figure of a man, made of carved wood, wearing an ermine-trimmed robe, loose trousers and a turban—the traditional costume of an Oriental sorcerer” (Standage 2002, 22). The chess-playing robot was toured around Europe and America, and exhibition matches were staged with such famous opponents as Napoleon Bonaparte. All the while, Kempelen maintained that the automaton operated by its own accord.

To prove there was no trickery, he opened the cabinet before every exhibition and showed spectators the dense tangle of gears, wheels, and levers. The sheer mechanical complexity on display convinced observers that advanced technology—beyond what they could comprehend—was powering the system. But Kempelen had actually created an elaborate illusion, not a robot. Inside was a human chess master who used magnets and levers to operate the Mechanical Turk; he simply hid behind the fake machinery when Kempelen opened the cabinet. In other words, the complex mechanical system that Kempelen showed people was meant to distract their attention from how the automaton really worked: human labor. Kempelen sold the idea of an intelligent machine, but what people witnessed was just human effort disguised by clever engineering.

In the 1730s, a French inventor named Jacques de Vaucanson constructed a copper-plated cyborg called *le Canard Digérateur*, or the Digesting Duck. It was the size of a duck, walked like a duck, and quacked like a duck. But its real trick, which amazed and baffled audiences, was that it could poop like a duck. The automaton “ate food out of the exhibitor’s hand, swallowed it, digested it, and excreted it, all before an audience” (Wood 2002).

Vaucanson claimed that he had built a “chemical laboratory” in the duck’s stomach to decompose the food before expelling it from the mechanical butt. While Vaucanson was an expert engineer—the duck was an intricate piece of machinery—like a good magician, he did not reveal how the duck worked. After his death, the secret was uncovered: there was no innovative chemical technology inside the duck; rather, there

were just two containers, one for the food and one for preloaded excrement. (Strangely, the Digesting Duck and the Mechanical Turk were both destroyed by museum fires around the same time in the mid-nineteenth century.)

Kempelen and Vaucanson would fit very well into Silicon Valley today. They could spend their days making mysterious machines and uttering wondrous proclamations about their supposed abilities, while attracting major venture capital investment. Perhaps Vaucanson—who literally snuck duck excrement into his technological system and called it innovation—would create the next biotech darling built on deception (à la Theranos). Whereas Kempelen’s Mechanical Turk was a forerunner of current systems of artificial intelligence (AI) not because it managed to play a game well, like IBM’s Deep Blue or Google’s AlphaGo, but because many AI systems are, in large part, also technical illusions designed to fool the public. Whether it’s content moderation for social media (Gray and Suri 2019) or image recognition for police surveillance (Mozur 2018), claims abound about the effectiveness of AI-powered analytics, when in reality the cognitive labor may very well come from an office building full of (low-waged) workers in outsourcing destinations such as the Philippines or India (Roberts 2019). This is not to say that the only options are a binary divide of either 100 percent machine power or 100 percent human power. Rather, the reality is a hybrid relationship where workers use technology to create value, managers use technology to discipline labor, and entrepreneurs use technology to erase the existence of the other two groups.

I call this way of building and presenting such systems—whether analog automatons or digital software—“Potemkin AI.” Potemkin refers to a façade designed to hide the reality of a situation. The term functions as an analytical label meant to help us understand the actual, existing practices and promises of the technology sector, as well as the conditions of workers that make these systems function. It also functions as a theoretical concept meant to reveal the expanded networks of labor, capital, production, and information that are hidden inside so many supposedly AI systems. Much as the theory of technopolitics shows that embedded within any given technology is a tangle of values, interests, and priorities (Winner 1978; Sadowski 2020), Potemkin AI opens the black box of digital labor in the planetary market. In other words, Potemkin AI is, at once, a way to demystify the real operations of these technological systems and the real relations of this political economic system.

This type of technology is the inverse of decision-making systems that seem to have humans-in-the-loop but actually do not, what Brennan-Marquez, Levy, and Susser (2019) term “skeuomorphic humanity.” Digital assistants and robocallers, for example, have been programmed to mimic the qualities and characteristics of humans—even including vocal tics designed to make us feel like we are interacting with a person.

When it comes to Potemkin AI, however, there is a long list of services that purport to be powered by sophisticated software but actually rely on hidden humans acting like robots (Solon 2018; Taylor 2018; Altenried 2020). Autonomous vehicles commonly use remote driving and human drivers disguised as seats to hide their Potemkin AI.¹ App developers for email-based services like personalized ads, price comparisons, and automated travel-itinerary planners use humans to read private emails (MacMillan 2018). A service that converted voicemails into text, SpinVox, was accused of using humans and not machines to transcribe audio (Cellan-Jones 2009). Facebook's much-vaunted personal assistant, M, relied on humans to answer questions—until Facebook shut down the service to focus on other AI projects (Newton 2018). The *Wall Street Journal* has called the pervasive use of human eyes marketed as AI “tech’s ‘dirty secret’” (MacMillan 2018).

The list of Potemkin AI continues to grow with every cycle of venture capital investment. The structural obfuscation of human labor is critical to the functioning of the financial system that props up the technology sector, pushing forward the constant development of so-called innovation. It is a system built on speculation, or the production of expectations that certain things will materialize and value will be realized (Pollock and Williams 2010); it is a system where the accumulation of data capital and the technology needed to generate and process data is treated as the overriding imperative (Sadowski 2019), and where the interests of massive institutions, like SoftBank's Vision Fund 2 (a proposed \$108 billion investment fund for AI companies), dictate the very trajectory of technological development (van Doorn and Badger 2020). If the machinery of real AI is not advanced enough to fulfill the wild promises and infinite desires of finance, then we should expect Potemkin AI to continue propagating as financiers and technologists sell the next best thing. Rather than the old marketing slogan “Accept no imitations,” Potemkin AI is an ideological project to convince the public that they should “accept only imitations.”

This chapter fleshes out the political economy of Potemkin AI. The first section situates this argument in recent work on digital labor, particularly regarding the emerging geographies of what Graham and Anwar (2019) have analyzed as a “planetary labor market” for gig work on digital platforms. These networks of (data) capital and (digital) labor are a crucial component of the “artificial artificial intelligence” that is now behind so much of what is sold as innovation. Following this, we will further unpack the psychopolitics of these black-boxed labor machines. For it is the *illusion* of being powered by AI that grants these systems efficacy. It's not only that their operations are hidden but that their proponents mistake an ideology for a technology, that they believe—or at least act as if—human labor has been replaced by mechanical minds. The chapter then concludes by exploring a simple yet critical question: Why go to the trouble of creating Potemkin AI?

Global Potemkin Village

The term *Potemkin* derives from the name of a Russian minister who in 1787 built temporary, fake villages to impress Empress Catherine II and disguise the true state of her empire. Potemkin technology, then, constructs a facade that not only hides what is actually going on but deceives potential users, investors, and the general public alike. In place of the Wizard of Oz telling us to pay no attention to the man behind the curtain, we have entrepreneurs telling us to pay no attention to the people behind the platform.

Importantly, the predecessors to Potemkin AI were constrained in space and time. The chess master sat inside the Mechanical Turk, manipulating the machinery in response to his opponent. The village facades stood in particular places and only had to be propped up long enough for the empress to pass through. When it comes to Potemkin AI, the homunculi are not inside the system *per se*; rather, their labor is transmitted into it over space and time. “Digital technologies extend and intensify working activity, rendering the boundaries of the workplace emergent” (Richardson 2018, 244). For instance, the cognitive labor of African workers today may power the American autonomous vehicles of the future (Anwar and Graham 2020). Or the “human computation” done by Indian workers in the past may support the valuation of a European start-up now (Irani 2015). In short, Potemkin AI is propped up by planetary networks of (data) capital and (digital) labor.

As Graham and Anwar (2019) explain in their geographical study of online outsourcing work mediated by platforms, by embedding digital work in “stretched-out networks of production” and information, the planetary labor market “facilitates a confluence that can transcend the spatial boundaries that constrained the convergence of employers and workers but remained shaped and characterized by multi-scalar and asymmetrical technological, political, social, cultural, and institutional factors.” It’s no coincidence that these geographies of (data) capital and (digital) labor tend to be organized along familiar colonialist political-economic relations wherein the “margins” are made to serve the “core,” providing it with sources of value that can be exploited to power the machine of growth and progress (Couldry and Mejias 2019; Sadowski 2019). The dominant “imaginaries of the digital economy” marshal the language of development and inclusion to justify deeply asymmetrical power-geometries (Wahome and Graham 2020). The fact that many citizens in wealthy nations also earn precarious wages through piecemeal digital work—producing datasets, training AI systems, serving platforms—says more about how the periphery is also nestled unevenly inside the imperial core of innovation (Gregory and Maldonado 2020).

Peeking behind the Curtain

When the inner workings of a technology are obscured, it is often labeled a “black box”—a term derived from engineering diagrams where you can see a system’s inputs and outputs but not what happens in between. An algorithm, for example, might effectively be black-boxed because the technical details are described using dense jargon decipherable by only a small group of experts. Or, with more advanced machine learning systems, even these experts might be epistemically shut out from understanding their automated decision-making processes (Burrell 2016). Accusations of willful obscurantism are often reserved for postmodernism, but as a paper on “troubling trends in machine learning scholarship” points out, research and applications in this field are rife with ambiguous details, shaky claims, and deceptive obfuscation (Lipton and Steinhart 2018). Being baffled by abstruse critical theory is one thing, but not being able to discern, say, how an AI application makes medical diagnoses or assesses insurance risk is much more consequential (Prince and Schwarcz 2020).

Algorithms might also be black-boxed through the force of law by the tech companies who claim them as trade secrets. In *The Black Box Society*, Frank Pasquale (2015) details how many of the algorithms that govern information and finance—the circulation of data and dollars—are shrouded in opacity. Algorithms are often described as a type of recipe. Just as Coca Cola keeps their formula a tightly guarded secret, so too do tech companies fiercely protect their “secret sauce.” In this case, analysts and regulators are institutionally shut out from inspecting the ingredients and probing the processes of these technologies. Again, it’s one thing to enjoy a beverage we cannot reverse engineer but quite another to take on faith proprietary software that makes sentencing decisions in criminal cases (Angwin et al. 2016).

Potemkin AI is related to black boxing, but it pushes obfuscation into deception. The Mechanical Turk, like many of the much-discussed AI systems today, was not just a black box that hid its inner workings from prying eyes. After all, Kempelen opened his automaton’s cabinet and explained the workings of what looked to be a complex machine. Except that he was lying. Similarly, marketing about AI systems deploys technical buzzwords that work like a magician’s incantations: Smart! Intelligent! Automated! Cognitive computing! Deep learning! Abracadabra! Alakazam!

Weaving the right spell can endow an AI system with seeming powers of objectivity, neutrality, authority, efficiency, and other desirable attributes and outcomes. As with any good trick, it matters less if the system actually works that way than if people believe it does and act accordingly. This power operates somewhat differently from a purely panoptical gaze wherein you are always potentially being observed by an unseen

watcher. Disciplinary power is certainly part of Potemkin AI's purpose, though, as the example of policing in the next section shows. Byung-Chul Han's (2017) work on psychopolitics helps tell a different story of Potemkin AI as a manifestation of what he calls smart power: "Power that is smart and friendly does not operate frontally—i.e., against the will of those who are subject to it. Instead, it guides their will to its own benefit. . . . It leads astray instead of erecting obstacles. . . . Smart power cosies up to the psyche, rather than disciplining through coercion" (Han 2017, 14). In other words, this is the power of coaxing and cajoling, of implanting beliefs and inducing action.

There's an undeniably seductive quality to AI. For consumers, it offers new heights of convenience, responsiveness, and personalization. We're not supposed to say out loud that we want a clever companion whose only purpose is to serve us, but most people would likely gladly accept their very own robot—from *robota*, the old Slavic word for "servitude" or "forced labor"—or at least get used to owning one pretty quickly. For capital, the prospect of possessing and controlling workers who "lack nothing but a soul," as the Czech playwright who coined the word *robot* described them (Intagliata 2011), has always been the dream. But in most applications AI is still an unsatisfying reality, if not a total fantasy. Potemkin AI is role-play. It's people masquerading as soulless systems. There's nothing wrong with this game per se, so long as everybody is aware and honest about who is actually serving whom. Yet as I've explained, Potemkin AI is seduction based on deception.

Psychopolitics reveals Potemkin AI to be more of an ideology than a technology. But that's not to say there is no material substrate to the ideology or that it's all just ephemeral relations. This is an ideology supported by planetary systems of control and capital (Bratton 2016; Sadowski 2020). After all, there is still labor to manage, data to administer, value to extract. How like capitalism to construct complex tangles of transnational networks for circulating money, information, and commodities, all to perpetuate a fetish for dehumanization! Exercising power is not just about effectively achieving particular outcomes or doing what works; it's also about deciding the parameters for how those ideas and goals will be defined. It's about preserving certain interests over others and reasserting the value of certain people over others. The desire for AI in some places supplants the rights of humans in other places.

Fake It till You Make It

Why go to the trouble of creating Potemkin AI? What's at stake for those propping up the façade? Broadly, we can point to two ancient reasons: profit and power. If an AI application relies heavily on human labor rather than machine learning, that doesn't make for a good sales pitch to venture capitalists and customers, nor does it convince the

public of the technology's capabilities. There are, of course, other motivations like fame and recognition, but I think we can safely label them as secondary to profit and power.

I illustrate these main motivations through the following two examples of Potemkin AI. The first is Amazon's Mechanical Turk platform (or MTurk), which enables mass exploitation of "microwork" distributed across global networks. The second is the advanced surveillance systems being deployed by the Chinese state. In the former, a dehumanized labor platform provides cheap "intelligence," while buying time for innovation to finally arrive. In the latter, a dehumanized monitoring apparatus creates the illusion of inescapable control.

MTurk allows employers to post discrete, often routine tasks like completing surveys or tagging pictures. Workers who complete these microjobs are then paid microwages: one study calculated the median wage at around \$2 an hour (Hara et al. 2018). As the *Financial Times* notes, MTurk is sometimes described by its creators in terms of "humans-as-a-service," or the "human cloud," or even "artificial artificial intelligence" (Hook 2016). These labels capture MTurk's approach of organizing a legion of human workers—hundreds of thousands of people—scattered across the world and hiding them behind a digital platform. Many companies rely on this pool of cheap labor that is ready to click and submit, allowing them to quickly scale up by completing tasks that they hope will one day be accomplished by AI software. Through this microwork, these workers are crucial to the development of AI, whether by producing and verifying datasets used to train AI systems or by simply impersonating the AI (Tubaro, Casilli, and Coville 2020). Yet their contributions to these technologies are belittled—that is, if their very existence is even acknowledged. Undervaluing the work of these people, while profiting off their labor, is all too common in a tech industry heavily biased with gendered, racialized, and classist notions of what entrepreneurs look like, who does innovation, and where venture capital is best invested.

Given that the name "Mechanical Turk" explicitly references the eighteenth-century hoax, it appears that there is no intention to deceive users about the flesh-and-blood foundations of the system. MTurk is indeed up-front about how work is outsourced to real live humans. However, whereas the original Mechanical Turk's inventor overtly claimed that his machine was autonomous, MTurk uses clever design to induce that impression in an audience eager to believe in the platform's Potemkin trick. As Lilly Irani describes, MTurk masks the MTurkers by making them appear to be just another soulless system. "By rendering the requisition of labor technical and infrastructural," she writes, MTurk "limits the visibility of workers, rendering them as a tool to be employed by the intentional and expressive hand of the programmer" (Irani 2015, 730). The platform and its interfaces allow employers to command people as though operating a mindless machine. In this case, Potemkin AI provides a convenient way to rationalize

exploitation—often of precarious workers in places with a large reserve army of disempowered labor—while calling it progress. “The result is that workers can lose a sense of any collective organization and feel replaceable, while clients exploit this lack of associational power of workers to exert their demands on workers” (Graham and Anwar 2019).

Irani (2015, 723) explains how, in addition to outsourcing tasks treated as “menial” (another way to devalue this type of work and the people doing it), Potemkin AI—such as MTurk—has helped compensate, both technically and ideologically, for the shortcomings of actual AI in completing cognitive tasks “by simulating AI’s promise of computational intelligence with actual people.” Even clickwork that seems brainless and dull is often still too advanced for “smart” machines. This simple fact does not bode well for funding of AI research and development, especially when investors eventually expect real results and profitable products. Contrary to their cheery marketing copy, investors and corporations don’t funnel their money into AI simply because they are interested in innovation for its own sake. AI promises to solve the problems of capital by unlocking exponential growth, eliminating labor costs, optimizing efficiency, and manifesting a slew of other expected outcomes. But the AI solution will only come about if the systems eventually actually work as promised.

There is a looming fear among those in the tech industry that once reality catches up with the hype, another “AI winter” will arrive, freezing all funding and interest in AI. The first cycle of hype for AI began building in the 1950s and grew until the mid-1970s, when enthusiasm was replaced by disillusionment. The ensuing AI winter lasted until the latter years of the first decade of the twenty-first century, when the combination of big data, processing power, and digital platforms opened up new advances in machine learning research.

As the hype began warming up and AI attracted more attention again, it became a label that start-ups could use as a shorthand for calling their service innovative, disruptive, and all-around superior to that of their dumb competitors. The inflated claims of what AI can achieve feed an expectation economy sustained by a circular logic: investment leads to promises, which leads to branding and more investment . . . and so on until yet another tech bubble bursts. Arguably, right now the most fully actualized product of AI systems is the sociotechnical imaginaries of a smarter planet (Sadowski and Bendor 2019).

Some of the most hyped, most cutting-edge applications of AI are supported by this sort of propaganda. A prime case is the extensive surveillance system being deployed in China. With millions of cameras deployed throughout Chinese cities, the state is looking to upgrade analysis of these feeds via AI and facial recognition that can automatically identify people and even punish criminals (Anderson 2020). For example, a camera at a busy intersection can now witness jaywalkers in action, shame them by

displaying their personal information on a screen, and send them a text message with a fine (Mozur 2018). It is questionable, however, just how accurate and automatic this name-and-shame system actually is right now. Buried at the bottom of a *New York Times* article about China's totalitarian tech is a revealing tidbit that highlights how the AI involved in this system is currently more hype than real:

The system remains more of a digital patchwork than an all-seeing technological network. Many files still aren't digitized, and others are on mismatched spreadsheets that can't be easily reconciled. Systems that police hope will someday be powered by AI are currently run by teams of people sorting through photos and data the old-fashioned way. . . . Still, Chinese authorities who are generally mum about security have embarked on a campaign to persuade the country's people that the high-tech security state is already in place. (Mozur 2018)

Potemkin AI is an effective way of constructing a panopticon. The disciplining power is much greater if people believe that an inhuman force is tirelessly processing feeds from the ubiquitous cameras, rather than groups of human analysts who take time, get fatigued, and make mistakes. Persuading people that the police are using AI is a way to normalize the idea that AI should be and, perhaps more importantly, already is ceaselessly monitoring society. Again, for the purposes of power and discipline, it matters less if the AI is real or fake—what matters is if people believe in the Potemkin deceit and behave accordingly.

It's easy to say that of course the Chinese government would employ propaganda to deceive the public about its power. But it's simply using a tried-and-true tactic of Silicon Valley: fake it till you make it. There is a long history of hiding the dead ends and delays in the process of technological development. This makes the process appear to be linear (no divergences), deterministic (no stopping), and progressive (no worries), while at the same time it suppresses any skepticism and convinces the public that resistance is futile because the tech is so effective and so much better than any alternatives. You can't argue with an algorithm, and the AI in the sky is always watching.

To varying degrees, many applications of AI are more like simulations of AI. This isn't to say that all research and development on artificial intelligence is an elaborate plot to erect a facade of efficacy. Plenty of researchers out there are working to advance the science bit by bit and devise practical applications without falling into the traps of obfuscation and overpromising. Yet at the same time, Potemkin AI is not just limited to a few bad actors, or a few bad apps, in an otherwise healthy industry. It's a core pillar of the political economic structure propping up this sector.

The problem with the industry isn't necessarily with AI not yet working, but rather with the cultural hype, ideological goals, and financial speculation that drive the development of AI. These technologies are not merely corrupted by global capitalism, they are created out of it. They are presented as standalone self-sufficient systems, yet they

are deeply dependent on a planetary network of production and information—not that we would ever know it without critically analyzing the real operations and real relations that their creators work so hard to hide. Too much of the attention and funding for AI is garnered by those who are looking to maximize their profits and/or secure their power. We see so many attempts to use AI as a tool for replacing human decisions, exploiting human labor, and administering human life that it becomes easy to believe these are simply the best, even natural, applications. But AI is an alibi, a way to rationalize these applications. Potemkin AI is a placeholder, a way to normalize this attitude in advance, that only works if we don't look beyond the facade.

Acknowledgments

This chapter is an expanded version of a 2018 essay originally published in *Real Life Magazine* as “Potemkin AI” (reallifemag.com/potemkin-ai/). I thank Mark Graham and Fabian Ferrari for their invitation and feedback on this chapter. And I especially thank Rob Horning for invaluable editorial guidance on the original essay, which really helped refine my thinking on the subject.

Note

1. Companies and researchers doing field tests with autonomous vehicles have used “ghost drivers”—human operators wearing a full body hood meant to make the car look driverless—to test how other drivers and pedestrians react to autonomous vehicles on the street (Lefkowitz 2020).

References

- Altenried, Moritz. 2020. “The Platform as Factory: Crowdwork and the Hidden Labour behind Artificial Intelligence.” *Capital & Class* 44 (2): 145–158.
- Anderson, Ross. 2020. “The Panopticon Is Already Here.” *The Atlantic*, September. <https://www.theatlantic.com/magazine/archive/2020/09/china-ai-surveillance/614197/>.
- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. “Machine Bias.” *ProPublica*, May 23. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- Anwar, Mohammad Amir, and Mark Graham. 2020. “Digital Labour at Economic Margins: African Workers and the Global Information Economy.” *Review of African Political Economy* 47 (163): 95–105.
- Bratton, Benjamin H. 2016. *The Stack: On Software and Sovereignty*. Cambridge, MA: MIT Press.
- Brennan-Marquez, Kiel, Karen Levy, and Daniel Susser. 2019. “Strange Loops: Apparent versus Actual Human Involvement in Automated Decision Making.” *Berkeley Technology Law Journal* 34 (3): 745–772.

- Burrell, Jenna. 2016. "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms." *Big Data & Society* 3 (1): 1–12. <https://doi.org/10.1177/2053951715622512>.
- Cellan-Jones, Rory. 2009. "Voice Technology Firm under Fire." *BBC*, July 23. <http://news.bbc.co.uk/2/hi/technology/8163511.stm>.
- Couldry, Nick, and Ulises A. Mejias. 2019. "Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject." *Television & New Media* 20 (4): 336–349.
- Graham, Mark, and Mohammad Amir Anwar. 2019. "The Global Gig Economy: Towards a Planetary Labour Market?" *First Monday* 24 (4), April 1. <https://doi.org/10.5210/fm.v24i4.9913>.
- Gray, Mary L., and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston: Houghton Mifflin Harcourt.
- Gregory, Karen, and Miguel Paredes Maldonado. 2020. "Delivering Edinburgh: Uncovering the Digital Geography of Platform Labour in the City." *Information, Communication & Society* 23 (8): 1187–1202.
- Han, Byung-Chul. 2017. *Psychopolitics: Neoliberalism and New Technologies of Power*. Translated by Erik Butler. London: Verso.
- Hara, Kotaro, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P. Bigham. 2018. "A Data-Driven Analysis of Workers' Earnings on Amazon Mechanical Turk." In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, Paper 449, 1–14. New York: ACM Press.
- Hook, Leslie. 2016. "The Humans behind Mechanical Turk's Artificial Intelligence." *Financial Times*, October 26.
- Intagliata, Christopher. 2011. "The Origin of the Word 'Robot.'" *Science Friday*, April 22. <https://www.sciencefriday.com/segments/the-origin-of-the-word-robot/>.
- Irani, Lilly. 2015. "The Cultural Work of Microwork." *New Media & Society* 17 (5): 720–739.
- Lefkowitz, Melanie. 2020. "'Ghostdrivers' Test Cultural Reactions to Autonomous Cars." *Cornell Chronicle*, April 22. news.cornell.edu/stories/2020/04/ghostdrivers-test-cultural-reactions-autonomous-cars.
- Lipton, Zachary C., and Jacob Steinhardt. 2018. "Troubling Trends in Machine Learning Scholarship." *Queue* 17 (1): 45–77. <https://doi.org/10.1145/3317287.3328534>.
- MacMillan, Douglas. 2018. "Tech's 'Dirty Secret': The App Developers Sifting through Your Gmail." *Wall Street Journal*, July 2. <https://www.wsj.com/articles/techs-dirty-secret-the-app-developers-sifting-through-your-gmail-1530544442>.
- Mozur, Paul. 2018. "Inside China's Dystopian Dreams: A.I., Shame and Lots of Cameras." *New York Times*, July 8. <https://www.nytimes.com/2018/07/08/business/china-surveillance-technology.html>.
- Newton, Casey. 2018. "Facebook Is Shutting Down M, Its Personal Assistant Service That Combined Humans and AI." *The Verge*, January 8. <https://www.theverge.com/2018/1/8/16856654/facebook-m-shutdown-bots-ai>.

- Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press.
- Pollock, Neil, and Robin Williams. 2010. "The Business of Expectations: How Promissory Organizations Shape Technology and Innovation." *Social Studies of Science* 40 (4): 525–548.
- Prince, Anya E. R., and Daniel Schwarcz. 2020. "Proxy Discrimination in the Age of Artificial Intelligence and Big Data." *Iowa Law Review* 105: 1257–1318. <https://ilr.law.uiowa.edu/assets/Uploads/ILR-105-3-Prince-Schwarcz-6.pdf>
- Richardson, Lizzie. 2018. "Feminist Geographies of Digital Work." *Progress in Human Geography* 42 (2): 244–263.
- Roberts, Sarah T. 2019. *Behind the Screen: Content Moderation in the Shadows of Social Media*. New Haven, CT: Yale University Press.
- Sadowski, Jathan. 2019. "When Data Is Capital: Datafication, Accumulation, and Extraction." *Big Data & Society* 6 (1): 1–12. <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 Roy Bendor. 2019. "Selling Smartness: Corporate Narratives and the Smart City as a Sociotechnical Imaginary." *Science, Technology, & Human Values* 44 (3): 540–563.
- Solon, Olivia. 2018. "The Rise of 'Pseudo-AI': How Tech Firms Quietly Use Humans to Do Bots' Work." *Guardian*, July 6. <https://www.theguardian.com/technology/2018/jul/06/artificial-intelligence-ai-humans-bots-tech-companies>.
- Standage, Tom. 2002. *The Turk: The Life and Times of the Famous Eighteenth-Century Chess-Playing Machine*. New York: Walker.
- Taylor, Astra. 2018. "The Automation Charade." *Logic*, August 1. <https://logicmag.io/failure/the-automation-charade/>.
- Tubaro, Paola, Antonio A. Casilli, and Marion Coville. 2020. "The Trainer, the Verifier, the Imitator: Three Ways in Which Human Platform Workers Support Artificial Intelligence." *Big Data & Society* 7 (1): 1–12. <https://doi.org/10.1177/2053951720919776>.
- van Doorn, Niels, and Adam Badger. 2020. "Platform Capitalism's Hidden Abode: Producing Data Assets in the Gig Economy." *Antipode* 52 (5): 1475–1495. <https://doi.org/10.1111/anti.12641>.
- Wahome, Michel, and Mark Graham. 2020. "Spatially Shaped Imaginaries of the Digital Economy." *Information, Communication & Society* 23 (8): 1123–1138.
- Wood, Gaby. 2002. "Living Dolls: A Magical History of the Quest for Mechanical Life by Gaby Wood." *Guardian*, February 16. <https://www.theguardian.com/books/2002/feb/16/extract.gabywood>.
- Winner, Langdon. 1977. *Autonomous Technology: Technics-Out-of-Control as a Theme in Political Thought*. Cambridge, MA: MIT Press.