Commissioned by PolicyLink

The People vs. the Algorithmic State

How government is aiding Big Tech's extractivist agenda, and what we can do about it

Ulises Ali Mejias,
Professor, SUNY Oswego

Abstract: This paper focuses on forms of discrimination, including racism, that are produced and replicated when the state uses biased algorithms, and when the state fails to regulate corporations that use biased algorithms. The paper specifically asks what happens when the state—whose purpose is to guarantee the rights of citizens—becomes a political and business partner with corporations whose profit model hinges on exploitative data-driven advertising, platform services, and gig work. The framework of "data colonialism" (Couldry and Mejias 2019) is used to examine the historical roots of these new forms of extractivism. Suggestions are reviewed for how the public can help prevent the deployment of discriminatory algorithms, and hold states and corporations accountable when this happens.

Who Writes the Law? Who Writes the Algorithm?

Laurence Diver corrected Lessig's famous dictum that *code is law* (Lessig 2000) by saying that code "is instead both more, and less, than law" (Diver 2021, 11). Code is more than law because it doesn't require interpretation and it can act directly and instantly on individuals, thus shaping their behavior. But it is less than law because it is created in the absence of the normative social mechanisms that usually keep laws in check. I open with this observation about code's imperfect likeness to law in order to point out another dimension of this relationship. Traditionally, laws are made by governments presumably for the public good, and software is made by corporations for profit. Today, however, corporations are superseding the role of the state in terms of making laws through their code, and states are becoming willing participants in this process. Governments are not only failing to regulate the creation and deployment of these new types of codes/laws by corporations, but by going as far as using code from these corporations, governments are becoming partners in joint ventures that are motivated by profit, and that are sometimes racist and discriminatory.

This paper examines the nature of these state-corporate relationships in the context of the design, use and (lack of) regulation of exploitative—and by extension, discriminatory—algorithms, with a focus on the United States (referred to as "the government"). While a lot has been written about the extractivist and biased algorithms created by corporations like Amazon, Facebook (now a subsidiary of Meta), and Google,[1] this paper addresses the role of the state in facilitating and encouraging the design of biased algorithms.

In the context of this paper, algorithms are calculative devices that, together with other elements like databases, are used to build automated decision systems. These algorithms frequently use machine learning as a method to achieve artificial intelligence, resulting in systems that can make decisions without direct human guidance.[2] My main point in this paper is that by using these systems, the government has a responsibility to guarantee that its delegation of power to make decisions involving citizens unfolds in a way that doesn't impinge on the rights afforded to those citizens.

Thus, my use of the term “algorithmic state” covers three types of situations:

- The **use** by the government of algorithmic decision systems to analyze data collected from citizens, make some sort of assessment using machine learning methods, and recommend or implement a decision following that assessment.
- The **failure** of the government to regulate corporations whose business model is based on forms of extractivist and exploitative advertising, platform services, and gig work that are data-driven and supported by artificial intelligence (AI).
- The **partnerships** between the public and private sectors—usually in the form of government contracts—where the state is a willing political and business accomplice in the use of data to monitor and discriminate against vulnerable populations (and which, as a result of these partnerships, makes effective regulation of the private sector more difficult).

As I argue below, while many of these functions represent new challenges resulting from advances in technology, they also represent forms of intentional discrimination and bias, including racism, that have deep historical roots, roots that are very much central to our national character.

My main thesis in the first half of the paper is that there is an inherent conflict of interest when the state, whose purpose is to guarantee the rights of citizens, becomes involved in the use of algorithms in two types of scenarios. In the first one, the state might get involved in the application of algorithms that are *intentionally* designed to discriminate. In the second one, the state might get involved in the use of algorithms that are *unintentionally* biased, and might fail to act once the bias is exposed.

In the second half of the paper, I review the kinds of options that are available for holding states and corporations accountable when they use discriminatory algorithms. I argue that in the absence of stronger and more responsive and imaginative regulation, and lacking a critical discussion in society about AI and automated decision systems, the use of algorithms by the state is already undermining democracy, increasing inequality and racism, and creating less open forms of governance.

Algorithmic Bias in a Colonial Context

Scandals involving discriminatory algorithms are part of our regular media cycles, along with the accompanying disavowals by corporations that it is not they who are racist or biased, but the artificial intelligence systems they have created. We have seen cases of Facebook housing ads that discriminate against people of color, Amazon hiring algorithms that exclude women, AI language models that exhibit an anti-Muslim bias, predicting policing systems that classify people of color as higher risks than whites, image processing algorithms by Google that label pictures of Black people as "gorillas"... and on and on.[3] We are told that since our society is racist, the algorithms that learn from our society may sometimes also be racist, but that corporations and governments can't do much about this unfortunate state of affairs. Do we need a longer historical lens to understand what is going on?

It's true that algorithms do not create inequality as much as they replicate and augment inequality that already exists, repurposing it along new technological vectors over old social domains. Discussing an Amazon algorithm that determined which neighborhoods would get same-day delivery (which ended up selecting white affluent neighborhoods and avoiding poor Black and Brown areas), Barocas, Hardt and Narayanan (2019, 14) write:

> To understand why the racial disparities in Amazon's system might be harmful, we must keep in mind the history of racial prejudice in the United States, its relationship to geographic segregation and disparities, and the perpetuation of those inequalities over time. Amazon argued that its system was justified because it was designed based on efficiency and cost considerations and that race wasn't an explicit factor. Nonetheless, it has the effect of providing different opportunities to consumers at racially disparate rates. The concern is that this might contribute to the perpetuation of long-lasting cycles of inequality.

In order to articulate a theoretical framework to understand how these long-lasting cycles of inequality are perpetuated by digital technologies, Nick Couldry and I have developed the concept of *data colonialism*. In essence, we define data colonialism as "an emerging order for the appropriation of human life so that data can be continuously extracted from it for profit" (Couldry and Mejias 2019, xiii). To put it simply: whereas the old colonialism grabbed land and human labor, the new one grabs us, our lives, through the medium of data—data that can be fed to algorithms to predict and shape behavior, and to make decisions that might impact different categories of people unequally.

Understandably, the word “colonialism” evokes strong emotions for many people who continue to see its brutal influence in shaping our world. For them, our use in this context might seem insensitive or exaggerated. But we use the word “colonialism” very intentionally. For us, this is not about *metaphorically* suggesting that what we are seeing is *like* colonialism. We argue that we are in fact experiencing a new phase of colonialism that is as monumental as the original phase of colonialism that began with the conquest of America. Furthermore, we argue that we cannot understand the evolution of capitalism without re-assessing the relationship between colonialism and capitalism, which digital data can help us to do.

When we begin that reassessment, we can see the emergence of industrial capitalism not as an innovative and self-spontaneous achievement of white Europeans, but as something built on the backs of non-white people during colonialism. In short, there wouldn’t be any factories without plantations. The wealth generated by the latter financed the former. And if we start our timeline with colonies and plantations as precursors to factories, which in themselves are precursors to today’s data centers, we can discern the 500-year-old feature that characterizes all of these formations: the doctrine of *cheap.*

Let’s start with cheap nature. To colonize the world, nature had to be framed as something cheap: it was said to be abundant, free . . . just there for the taking. From a legal perspective, it was said to be *without an owner,* at least a civilized one. The legal concept of *terra nullius* (literally, “no man’s land”) was used by colonizers to award themselves the status of rightful occupiers and owners of any territory they deemed empty. By the time Indigenous inhabitants understood this new language of private property (if they hadn’t been enslaved or exterminated in the process), it was too late. While our situation today is incomparable to theirs, by the time we understand what it is that we agreed to when we clicked the “I accept” button on the User License Agreement, it is also too late.

From cheap nature we move to cheap labor. In colonialism, some humans (and this was mostly determined by race) were seen as naturally predisposed to provide the labor required to transform nature into wealth. Starting with colonialism, and through capitalism, exploitation and abuse were thus framed as social progress. The colonizer’s rule was said to be good for all of humanity (it even brought salvation). And in the process, life itself—that of the colonized—became cheap.

We move from cheap nature to cheap labor to arrive at . . . cheap data. In cheap data, we see some of the same extractive rationality we saw with nature and labor: cheap data is said to be abundant, without owner (at least in its aggregate form), and just there for the taking. Yes, individually we might feel like we own our personal data, but by the time it is collected and analyzed, it no longer belongs to us. And that is the point: data might be cheap, but to become something useful it needs processing, like oil needs to be refined. This requires intensive capital and technical know-how, most often in the forms of sophisticated algorithms which only big corporations or states can provide. Our only job is to generate the data, and we are told this division of labor is what progress looks like.

Data colonialism can seem like a controversial concept, particularly because of the question of whether we are entitled to use the word "colonialism" in this context. It is important to keep in mind that we are not establishing a one-to-one correspondence between historical forms of colonialism and data colonialism. In our work, Couldry and I (2019) emphasize the many differences between the two forms of colonialism: the modes, intensities, scales, and contexts are very different. But we do point out that there is one crucial similarity. The *historical function* of both kinds of colonialism is the same, and that function is to *extract,* to *dispossess.*

Racism as an Algorithmic Function

As the data colonialism argument makes clear, old forms of extraction and discrimination are given new lease through digital technologies, to the detriment of populations that are already vulnerable from centuries of colonial legacies. Race is often a primary dimension along which this algorithmic discrimination operates, a phenomenon that has been labeled "The New Jim Code" by Ruha Benjamin (2019).

According to a report by the group Public Citizen, algorithms are already responsible for people of color paying higher auto insurance premiums; getting lower credit scores; being assigned higher mortgage rates or being turned down for loans altogether; being assigned worse grades and getting rejected from schools; denied life-saving care; being handed tougher sentences; and experiencing more surveillance and policing (Chung 2021).

Why does this happen? According to Barocas, Hardt, and Narayanan (2019, 163-172), multiple factors are at play in the operationalization of racism in machine learning systems:
These systems don't operate in a vacuum. They are trained with real world data, and the real world is racist.

- These systems are trained to make predictions based on differences among users. Differences, however, may not just be a function of preference, but of structural dynamics (like racism). In such cases, the system may in fact reinforce stereotypes and institutionalized discrimination.
- Machine learning automates the discovery of correlations (for example, a machine learning system might discover that people applying for jobs in programming seem to like chess). But when those correlations are deployed as selection criteria, the system basically creates its own

discriminatory reality (the system might determine that people applying for programming jobs must list chess as a hobby, for example).

- Algorithmic recommendation systems encourage segregation. In essence, they make it easy for people who think alike to form conclaves, and protect themselves from outside opinions (just like it happens offline, but at a bigger scale).
- These systems also encourage homogeneity in decision making. The tyranny of trending topics means that the algorithms of a handful of global platforms decide what becomes mainstream, whose voice is heard, and whose is not.
- This new reality shifts power to those who can understand and "game" the system. Algorithms require a particular kind of expert to design, understand, and analyze them. Those who can make decisions about these systems acquire new forms of power in society (think of the power the leaders of Big Tech, almost exclusively white men, have acquired in the last decade).

When it comes to race, social problems are compounded by the application of technologies whose impact has not been properly ascertained. Thus, old structural problems with racialized dimensions—such as poverty, lack of access to proper education and health care, exploitation of workers, and police profiling, all of which are discussed below—are aggravated by the explicit targeting of Black and Brown people that algorithms make possible.

Consider poverty. Algorithmic state discrimination is often used against the poor. In her book *Automating Inequality*, Virginia Eubanks (2018) analyzes public assistance systems related to food assistance, Medicaid, homelessness, and at-risk children in the United States. In many of these systems, eligibility is determined automatically by algorithms, whose operation is opaque and seemingly arbitrary. Eubanks describes the punitive and harmful effects inflicted on those who are deemed ineligible—women of color, disproportionally—and brings to light the demeaning and excessive amount of surveillance that is inflicted on the poor.

Or consider the impact that algorithms can have on the education of people of color. Recent research has revealed that algorithms making admission decisions at top-performing high schools across New York City systematically screen out Black and Latinx applicants; that algorithms in charge of assigning grades to students who could not take exams during Covid closures failed native Spanish-speaking students in their Spanish classes because the algorithm determined those students, living in a low-income district in Colorado, would not do well; and that an algorithm used by over 500 universities in the US to assign risk scores for retention systematically assigned higher risk scores to Black students over white students (Chung 2021).

A similar impact can be seen when it comes to health care. One case involves an algorithm used widely by hospitals and insurance companies in the US to determine the level of care patients should receive. A study published in *Science* revealed that the algorithm displayed rampant racism, which skewed the medical care received by millions of Black people: “Black people had to be sicker than white people before being referred for additional help. Only 17.7% of patients that the algorithm assigned to receive extra care were Black. The researchers calculate that the proportion would be 46.5% if the algorithm were unbiased” (Obermeyer et al., 2019).

When it comes to labor issues, we can consider what is happening in the sharing or gig economy. Companies in this sector often promote their platforms as opportunities for people of color to be their own bosses, and achieve financial security. But the truth is that these companies’ business models harm Black and Brown workers disproportionately, since they are the ones who do most of the work (a recent study in San Francisco, for example, found that 56% of gig workers are immigrants, and only 22% are white (Schor 2020). As an American Civil Liberties Union (ACLU) report puts it, the gig economy’s business models “deepen the desperation of workers who have been excluded from stable employment” (Ongweso 2021b). Whether we are talking about Airbnb, Uber, or TaskRabbit, the experience of gig workers is largely determined by their race: non-whites earned less money, got lower ratings, and experienced more discrimination when using these kinds of apps, all because of the algorithms (not to mention the racial gentrification effect that companies like Airbnb are having, pushing residents of color out of neighborhoods) (Schor 2020).

And then there is the role of algorithms in policing and the criminal justice system, areas already marked by long histories of institutionalized racism. Police departments are using algorithms to conduct facial recognition, automatic license plate recognition, voice identification, phone call metadata analysis, security and body camera footage analysis, and social media data harvesting and monitoring, to name a few applications (Williams and Kind 2019). While official data is difficult to obtain, there have already been documented cases of people of color being wrongfully arrested as a result of these systems (General and Sarlin 2021). Meanwhile, courts are using algorithms to conduct “risk assessments” to predict whether a defendant is likely to commit a crime in the future. This information can impact bail or sentencing decisions, and unsurprisingly, race has been shown to be a factor in how the algorithm treats defendants. According to one report, “Black defendants are 45% to 77% more likely to be assigned higher risk scores than white defendants” (Larson et al., 2016).

In cases such as the ones just described, is the state using the algorithm to blame, or the company that sold the algorithm to the state? This question brings us to the complicated relationship between the public and private sectors.

The Algorithmic State and Big Tech

We can assume that all federal agencies and most state and local agencies use algorithms in one way or another. Algorithms used by government agencies can, for instance, determine Covid-19 vaccine allocations, establish eligibility for welfare programs, check student papers for plagiarism at a state college, or monitor potential terrorist threats. In many cases, the application of these algorithms involves close collaboration with corporations who design the algorithms the government uses. In other cases, the government is supposed to regulate these very same corporations as they apply their own algorithms, making the relationships between the state and technology companies complicated and multifaceted. In this section, I will attempt to differentiate between the various ways the state uses algorithms, and how the private and public sector relate to each other during this process.

Use of Algorithms by the State

As defined earlier, the algorithmic state is characterized by the use of automated decision systems to analyze data collected from citizens, make some sort of assessment based on the data, and recommend or implement a decision following that assessment. Of these applications, surveillance algorithms have received a lot of attention in the last decade, probably because of the secrecy in which they happen, and the dramatic way in which they came to light as a result of the Snowden revelations in 2013.

There are many surveillance programs in the US that we know about, and some that we probably don't even know exist. They include: Boundless Informant, a program that has collected billions of data reports generated from global electronic communication; BULLRUN, a program intended to preserve the NSA's ability to eavesdrop on encrypted communications; DCSNet, an FBI interface that allows for instant wiretaps on any telecommunications device in the US; the personal and financial database known as Main Core, containing information on millions of US citizens; and PRISM, a clandestine surveillance program in which allegedly most major Big Tech companies participated.[4]

In the absence of transparency about these kinds of systems, regulation about the manner in which they can be used, and lacking a larger critical discussion in society about their limits, we can expect the use of these algorithms by the state to undermine democracy and create less open forms of governance. They are also likely to increase inequality and discrimination, as many of the algorithms used by the state have not been adequately tested for bias.

Consider the use of facial recognition systems deployed at airports, such as the ones used in passport control. Studies have already shown that similar systems are more prone to errors in instances where the subject is not white (Buolamwini and Gebru 2018). In that case, maybe the repercussions are limited to longer wait times at airport lines for people of color. But suppose, in another example, that the government decides to use an algorithm to detect online hate speech (which, given the list of surveillance programs listed above, we can assume that it does). Again, studies have demonstrated that hate speech detection algorithms are biased against people of color (Sap et al., 2019). Instead of a minor inconvenience, the algorithms in this case can lead to further surveillance and the individual's loss of privacy or even freedom.

Political ideology can also play an important role in shaping the application of corporate algorithms by the state. To mention only one example, the Trump administration awarded Palantir (a company co-founded by Trump supporter Peter Thiel) $150 million in federal contracts to develop Immigration and Customs Enforcement (ICE) tools such as the Integrated Case Management system and the FALCON analytical platform, tools which the agency uses to track down, detain, and charge immigrants and their US families ("Who's behind ICE?" 2018).

At the other end of the political spectrum, some state governments have passed laws to protect citizens from the unregulated application of algorithms. Maine, for instance, recently approved legislation restricting the use of facial recognition (Woodruff 2021), California banned predictive policing (Asher-Schapiro 2020), and the city of Seattle passed the nation's strongest regulations on surveillance technologies ("Seattle Adopts Nation's Strongest Regulations" 2017). The point is that even within the US, we can see a variety of political agendas when it comes to the use and regulation of algorithms by the state.

One area of growing concern is the use of algorithms by governments and political parties to basically advertise to citizens in the hope of influencing their behavior. These new ways of political "nudging" use the same psychological profiling tools used by advertising companies to target customers, but instead of trying to sell them a product, the goal is to influence the public during elections or campaigns. According to a *Guardian* article, the shift to this form of governance "stems from a marriage between the introduction of nudge theory in policy making and an online advertising infrastructure that provides unforeseen opportunities to run behavioral adjustment campaigns" (Hern 2021). In essence, what started as a scandal when Cambridge Analytica did it is now a common practice of the algorithmic state. And ast we saw in the 2016 election, these techniques can also be abused by foreign powers in order to sow disinformation. In one case, Russian agents purchased Facebook ads seen by 126 million Americans encouraging them to click the Like button if they wanted Jesus to win in his fight against Satan a.k.a. Hillary Clinton (Isaac and Wakabayashi 2017). While these are examples of propaganda and disinformation, rather than simply “nudging,” the point is that those ads weren’t targeted to random users, but thanks to the use of psychometric profiling and algorithms, to users who were more likely to be driven to action (from voting to violence) by a particular kind of political message.

And while some politicians and lawmakers might be receptive to the idea of regulating the use of algorithms by corporations, it is less likely that they will be equally open to the idea of being transparent about the state’s own use of algorithms, especially when it comes to surveillance. That is why Edward Snowden is considered a traitor for exposing the work of the National Security Agency, while Facebook whistleblower Frances Haugen is invited to provide testimony before Congress. It is easier for some government representatives to hate on Mark Zuckerberg than to be critical of themselves.

More often than not, however, the government has been content to let Big Tech run amok, to the detriment of citizens. This dangerous kind of neglect is examined next.

Failure to Regulate Algorithms

The Palantir example, mentioned above, highlights the fact that corporations are important players in the way the state uses algorithms. Whether it is as the designers and suppliers of the systems used by the state, or as the entities that the government tries to regulate, corporations—particularly those we’ve come to know as Big Tech—are very much part of this equation.

While the nature of the partnerships between Big Tech and the state will be examined in the next section, here I will briefly mention the failure of the state to regulate an industry whose business model is founded on extractivist and exploitative forms of advertising, platform services, and gig work that are all data-driven and supported by AI algorithms.

The reason commonly given for the state’s failure to legislate the data sector is that the algorithms are too new and too complex for politicians and lawmakers to properly understand and regulate (the clueless performance of some politicians during congressional hearings focusing on Big Tech seems to corroborate this observation). Closer to the truth, probably, is the fact that Big Tech is a big “investor” in our government. In only the second quarter of 2021, Apple spent $1.64 million on lobbying, Amazon $4.86 million, Microsoft $2.47 million, Facebook $4.77 million, and Google’s parent company Alphabet spent $2.09 million (Edgerton and Allison 2021). Palantir, by the way, spent more than $2 million on lobbying in 2020, and received $40 million in Covid-related contracts from Health and Human Services (Schwartz 2021).

Lobbying might also explain in part why the government has not gotten around to revising the controversial Section 230[5] of the Communications Decency Act, which protects companies like Twitter and Facebook from being sued for libel if users publish defamatory content on their platforms. Even President Trump, with newly found antitrust vigor, tried to get rid of this law when social media companies dared to fact-check and slap warnings on some of his tweets. But while it is important to defend the principles of free speech behind Section 230, there are other ways of regulating social media companies to hold them accountable for the spread of false information. Unfortunately, the government seems unwilling to come up with a better version of Section 230, even after the deadly disinformation campaigns following Covid-19. In allowing Big Tech to try to fix itself, the government is showing its own complicated conflicts of interest with this sector.

Partnerships and Conflicts of Interest

From the perspective of the data colonialism thesis, data is an instrument of extraction. Extracting data related to our social lives—which happens essentially every time we use our phones, computers, or smart devices—allows companies to better target advertisements to us, to create platform services that are addictive and that can nudge our behavior, and to create gig platforms that exploit workers. By failing to adequately regulate these companies, and by buying services from them (over 6,500 government agencies are reportedly using Amazon's cloud services[6]), the algorithmic state becomes a willing participant in the process of data colonialism, a process where all of us pay a price, but where the heaviest price is paid by historically vulnerable populations: people of color, women, and the poor.

It is important to point out that the exercise of colonialism and racism has frequently been—and continues to be—a partnership between the state and the private sector. To colonize the world, corporations (regardless of whether they are named East India Company or Alphabet) need the legal frameworks and sources of moral legitimacy provided by the state (and previously, the Church). And the state directly benefits financially from the privatized colonialist enterprise by collecting taxes, resources, and benefits. In an age when data is becoming "the new oil" (*The Economist* 2017), the colonialist partnership between the state and Big Tech becomes something we should be concerned about.

The global machinery conducting this new form of extractivism is something Couldry and I call the *social quantification sector,* or SQS (2019, 46): the collection of public and private organizations devoted to converting more and more aspects of our lives into data that can be used to generate profit and influence behavior. The SQS includes—of course—Big Tech (in the US and Europe, as well as China), but also all sorts of small and large players including manufacturers of hardware and software, media companies, and data analytics companies. The SQS also features many strategic partnerships between public and private players. Out of necessity or choice, most states have, in the last few decades, relinquished many of the functions and responsibilities they used to have in terms of collecting and analyzing data about their citizens, and increasingly rely on private corporations to fulfill their data and information technology needs, a process that has resulted in an economic boon for companies in the SQS (particularly post Covid-19), and which is financed by the taxes of the same citizens subjected to this new form of extractivism.

From 2007 to 2019 (a period corresponding to the War on Terror), Big Tech saw a 50-fold increase in government contracts: Amazon and Microsoft enjoyed a 400% and 800% increase, respectively; Google netted $16 million in contracts with the Pentagon, $2 million with the Department of Homeland Security, and $4 million with the Department of Justice; Facebook booked $167,000 in contracts with the Pentagon and $363,600 with the DHS (Ongweso 2021a). These numbers, however, do not tell the whole story, since it has been discovered that thousands of contracts between the Department of Defense and federal law enforcement agencies—including Immigration and Customs Enforcement, the FBI, the Drug Enforcement Agency, and the Federal Bureau of Prisons—recruit subcontractors as a way to hide the connection between the state and Big Tech companies (Glaser 2020).

In other cases, corporations are giving their products to governments at no cost, as a way to minimize detection by the public or the media. Palantir, for instance, gifted their analytics platform to the New Orleans Police Department, making the acquisition practically invisible to the city council (Winston 2018). Similar examples of these Trojan gifts from corporations to states have been documented in Latin America and throughout the world (Pisnau and Arroyo 2021).

This exemplifies a continuing trend where the government, instead of developing its own technologies for conducting delicate state operations like surveillance and policing, outsources the development of these technologies to private corporations. For instance, the FBI used to have a program called Carnivore to monitor email and electronic communications. These days, states can purchase more sophisticated global spyware tools such as NarusInsight and Pegasus to do the same thing (Pegasus was developed by an Israeli cyberarms firm to spy on mobile phones, and is used by many countries around the world) (Feldstein 2021).

What further proof of the close relationship between public and private SQS is needed than the fact that it is a common practice for Big Tech to hire former state employees to help them get access to lucrative government contracts (Lippman and Birnbaum 2021)? The conflict of interest is obvious: the state cannot effectively regulate the companies it is in bed with.

Possible Solutions

How do we begin to decolonize our data and create technological infrastructures not based on extractivism? Two obvious solutions seem to be to design better (non-discriminatory) technologies and algorithms, and to regulate the current use of algorithms.[7]

Non-Discriminatory Design

As far as the "better design" solution, plenty has been written about the importance of incorporating more diverse people and ideas into the design of technologies.[8] Proposals like these resonate with work being done under the rubric of terms like AI for Social Good, Indigenous Data Sovereignty, and others. These initiatives are crucial and valuable, but their premises should be critically questioned. As Couldry and I (2019) suggest, there is something fundamentally extractive about AI and data that makes it impossible to simply "fix" them and apply them for social justice without addressing the macro dimensions of data colonialism. In other words, focusing on a liberal narrative—that an individual design can make a difference—ignores the larger structural issues.

And of course, if the premise is that we should let Big Tech figure out how to design more just algorithms, then the problem is that we are asking the institutions responsible for creating inequality to design tools that correct that inequality. While this is not necessarily impossible, it does contribute to the sense that *it's easier to fix bias in technology than in people*, which again ignores the larger structural issues of why we have created unjust systems in the first place.

Legislating the Legislators

The other option is to expect that the algorithmic state will regulate the algorithms, theirs as well as Big Tech's. As we saw in the discussion above, this is complicated. Consider China, an extreme example of what happens when the algorithmic state doesn't have to answer to anyone.

In China, the Communist Party has essentially built a physical and digital barricade around the region of Xinjiang, for the purpose of what some are calling a genocide. This complete surveillance apparatus, aided by algorithms, is meant to limit the freedom of ethnic Muslims. Apart from the physical infrastructure of checkpoints and detention centers, the system relies on cameras and recognition software trained to identify Uyghur faces, quick response (QR) tags to identify Uyghur homes, GPS to follow Uyghur movements, a vast tracking system to monitor Uyghur phones and computers (harvesting everything from photos to messages to location data) and an extensive database of Uyghur biometrics, including blood types, DNA, fingerprints, voice prints, and iris patterns for essentially every resident of Xianjiang (Khatchadourian 2021). The system runs on what is supposedly the world’s fastest supercomputer which uses algorithms to correlate data, predict “threats,” and issue warnings. The infrastructure for this massive system is built by Chinese companies, which also export it to other countries (many times at low or no initial cost, as an incentive for adoption). It was recently discovered that even some cities in the US were buying and using the same technologies that are facilitating this repression in China (Whittaker 2021). The selling point seems to be that governments can train the system on any local group they wish to monitor, including social justice activists.

While the Chinese Communist Party (CCP) is carrying out the repression and re-education of Uyghurs, it is also legislating digital technology in a way that might seem interventionist to some, but downright sensible to others. The CCP has recently announced limits on how many hours per week children can spend playing video games, abolished certain forms of influencer culture, and banned cryptocurrency transactions (which are disastrous for the environment and promote criminal activity). The government will increase the frequency and intensity of inspections for online companies to ensure time limits and anti-addiction systems are put in place (Goh 2021).

In the United States, some attempts towards the regulation of algorithms are centered on the Algorithmic Justice and Online Platform Transparency Act of 2021, HR3611 (itself an enhanced version of the 2019 Algorithmic Accountability Act). The Act has as its main goal to “prohibit the discriminatory use of personal information by online platforms in any algorithmic process” and to “require transparency in the use of algorithmic processes and content moderation” (Matsui 2021).

The tone of the Act reveals how mainstream a critique of Big Tech is becoming. It alludes to “manipulative dark patterns” of data collection by platforms, and “opaque algorithmic processes” that result in discriminatory experiences. But it also clearly mentions the “disproportionate harm to populations that already experience marginalization,” (Matsui 2021, Section 2 Findings) which brings the analysis in line with approaches such as data colonialism.

HR3611 is largely concerned with algorithms used in content moderation and advertising; in other words, algorithms used in “determining the order or manner that a set of information is provided, recommended to, or withheld from a user” (Matsui 2021, Section 3 Definitions). Its approach to a solution is mostly based in creating reporting systems. For instance, the Act would require each online platform to publish a “Notice and Review of Algorithmic Process” that lists the categories of personal information the platform collects, the manner in which it does so, how the online platform uses such information, and the method by which the type of algorithmic process prioritizes, assigns weight to, or ranks different categories of information.

Most of this information is already identified in many platforms' Terms of Use or End User License Agreements, so this requirement is not necessarily groundbreaking. The Act would require all of these Notices to be publicly available for analysis and consultation. And the Act does add the provision that these Notices need to include a statement of whether the platform utilizes any algorithmic process that relates to opportunities for housing, education, employment, insurance, and credit, and if so, a description of how the algorithm was tested for accuracy and fairness, and against bias and discrimination. Online platforms are also required to maintain a public library of advertisements, and the bill criminalizes "discriminatory advertising," which includes "any harm caused by predatory or manipulative advertising practices, including practices targeting vulnerable populations" (Matsui 2021, Section 6, Prohibited Conduct).

While it would be good to have all of these requirements made into law, the limits of HR3611 are evident. Its narrow view of the kind of work algorithms perform (content moderation and advertising) covers a lot of ground, but also leaves out many uses, such as prediction, recommendation, eligibility checks, etc. The bill also confines itself to the commercial sphere, while saying nothing of algorithms used by public agencies or by private-public partnerships. Fortunately, important work has already been done to provide more comprehensive ways to assess algorithms, including their use by governments.

Algorithmic Impact Assessments

Impact Assessments are commonly used to measure the economic, environmental, and social ramification of policies implemented by state or private actors. They are evidence-based studies meant to provide transparency and involve the public in a discussion of the impact of big-scale projects, with the view towards eliminating any one stakeholder from having a disproportionate power in determining the course of action.

Both the Data & Society and AI Now institutes have recently published white papers on how Impact Assessments can be applied to algorithms. In its report, *Assembling Accountability: Algorithmic Impact Assessment for the Public Interest* (Moss et al., 2021), Data & Society explains that Impact Assessments in general—and by extension Algorithmic Impact Assessments (AIA), in particular—address three important questions: "what a system does; who can do something about what that system does; and who ought to make decisions about what the system is permitted to do" (Moss et al., 2021, 5). Contrary to HR3611, which basically leaves corporations to monitor themselves, Data & Society suggests that corporate accountability practices should not remain internal, and that accountability, in this case, cannot be reduced merely to "voluntary commitments to auditing and transparency" (Moss et al., 2021, 9).

According to Data & Society, there are various interrelated components of an AIA. The *source of legitimacy* is an important component because it establishes the entire exercise as grounded in recognized and publicly valued norms (such as laws and regulations) that have the power to hold actors accountable. Another important component are the *actors* themselves, as well as the *forum* in which responsibility for potential consequences is established. AIAs are usually triggered by a *catalyzing event* (such as a discovery of algorithmic discrimination) and are followed by a mandated *time frame* within which an assessment must be performed. During the assessment, *public access* guarantees that the process and documentation remain open and transparent, and *public consultation* creates the conditions for feedback to be collected from a broad set of stakeholders. The AIA itself is carried out by independent and impartial *assessors* following established *methods*, which result in abstract and evaluative *impacts*. This information is reported to the forum, and a list of *harms* and procedures to achieve *redress* in light of those harms are identified to secure justice (Moss et al., 2021).

The process for conducting AIA, as outlined by Data & Society, is considerably more difficult to implement than the simple reporting requirements of a bill like HR3611, but, by the same token, it is likely to have a more equitable outcome. In other words, it does more to answer Ulbricht and Yeung's call "to study algorithms in context, in ways that attend to the embeddedness of these computational systems within larger socio-technical and organizational infrastructures, legal frameworks, political cultures, and user communities as well as the technical architectures, protocols, and configurations through which they are designed to function" (Ulbricht and Yeung 2021, 16).

The Data & Society report recognizes that the primary challenge of AIA will be to build a social ecosystem with balanced relationships between regulatory agencies, private corporations, and public advocates—an ecosystem where these stakeholders agree on the scope, applicability, and consequences of the AIA.

While the Data & Society report is somewhat agnostic on whether AIA should be used in private or public sectors, the AI Now report, *Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability* (Reisman et al., 2018), directly addresses the issue of algorithms used by state agencies. The report, in essence, calls "for an end to the use of unaudited 'black box' systems in core public agencies," including those involved in criminal justice, predictive policing, energy use allocation, labor rights and access to education, since it is here that "disparate impacts" are more consequential, and we need ways to hold these agencies accountable when they occur (Reisman et al., 2018, 9).

The key policy goals of Algorithmic Impact Assessments, according to AI Now, are to empower the public with transparent information about the automated decision systems that impact their lives; to empower public agencies with expertise to review the systems they acquire; to ensure greater accountability by providing meaningful and comprehensive opportunities for review by experts; and to guarantee that the public also has opportunities to respond to and dispute the results of the assessment. The process for conducting AIA, they suggest, includes a pre-acquisition review, an initial disclosure of requirements by the agency, a comment period, a due process challenge period, and a process for renewing AIA. The procedure is designed so that no automated decision system—regardless of how it is procured—is deployed without this form of assessment (regardless of whether corporations or states claim the operations of their algorithms are a matter of trade secrecy).

Like the Data & Society report, the AI Now paper recognizes that a balanced multi-stakeholder environment where each party has the power to hold the other accountable is necessary for AIA to work. Without the appropriate levels of transparency and shared governance, AIA "could become a checkbox that agencies mark off and forget, potentially sidelining community concerns" (Reisman et al., 2018, 28).[9]

Limits of the Proposed Solutions

As a whole, the solutions discussed above present valuable proposals to move forward. But they are also characterized by two important limitations.

First, each proposed solution seems to limit its scope too narrowly. HR3611 proposes to legislate the private sector, using a limited view of what algorithms can do in that sector, while saying almost nothing about the use of algorithms by the state. At the other extreme, the AI Now report focuses on government agencies, while saying little about how to regulate the corporations who build the algorithms used by the state. Once a discriminatory algorithm enters the market, the damage is done; even if it is rejected by some states, it might be used by others. In other words, when it comes to regulating partnerships like the one discussed earlier between Palantir and ICE, it seems there are no AIA or bills like HR3611 being proposed that would hold members of these partnerships accountable. The Data & Society report does ask the question, in passing, of whether the same AIA mechanism should be applied to private and public sector uses of algorithms (Moss et al., 2021, 50). But the real issue is the complicated nature of assessing the impact of algorithms in a Social Quantification Sector increasingly characterized by symbiotic partnerships between corporations and the state. Who builds the algorithm, with what data, and for what purpose are questions that today seem to cross the boundaries of the public and private sectors, and we need better methods for holding these partnerships accountable.

Which brings us to the second limitation. These solutions, which focus on design and legislation, are based on an incomplete ontological view of the computer system (cf. Ulbricht and Yeung, 2021). If this ontology is based on a simple input-processor-output model, the solutions discussed above seem to focus on the processor-output side of the equation only. In other words, questions about the ways in which the input is generated are mostly ignored in arguments about designing and regulating algorithms. That is why frameworks such as data colonialism, which attempt to provide a model for understanding where the data comes from, and at what cost, must be integrated into the models for assessing the impact of algorithms.

Policy Recommendations

In conclusion, I offer specific recommendations for confronting the problems created by the algorithmic state.

1. Promote better, faster, and more extensive *ex post* legislation.

Until the culture around the creation and deployment of algorithms and AI changes significantly, we will need to continue to rely on legislation that attempts to place controls on algorithms that are already operating in society, such as HR3611. The problem with this approach is that the public, presumably acting through the government, will always be playing catch-up, and the process will be slow, narrow, prone to corporate influence and ineffective. Still, it needs to be done, and hopefully in a faster and better way, while pursuing the other measures described below. President Biden's newly created Artificial Intelligence Research Resource Task Force[10] could be a step in the right direction.

2. Develop innovative models for *ex ante* legislation.

Instead of waiting to regulate algorithms *after* they have been created and deployed, we need to find ways of intervening in the process *before* the algorithm is even designed. When an automobile manufacturer sets out to design a new engine, it knows that the engine must meet certain emission standards. We need the same concept for algorithms. This is what Laurence Diver (2021) calls *digisprudence*: whereas legisprudence is concerned with the design of legal rules, digisprudence is concerned with the design of software-based rules. Standards need to be embedded directly into the tools that are used for the development of algorithms.

3. **Promote the widespread use of Algorithmic Impact Assessments.**
   AIA, which has been described above, are comprehensive ways of ex post legislation, and should be adopted widely. It is important to note here that AIA should not serve merely and simplistically as a call for "the public" (i.e., citizens exposed to the effects of automated decision systems) to be included in the design and oversight of algorithms. For one thing, calling for human oversight "creates shallow protection that companies and governments can easily avoid in superficial ways," which in fact may end up legitimizing opaque systems (Green and Kak 2021). And for another, the complexity of these computer systems is becoming so abstract that sometimes it is difficult even for their creators to explain why algorithms do what they do, so the public needs to be adequately prepared to enter into this analysis (see my previous recommendation).

4. **Get ahead of the wave of mistrusting government by promoting transparency and shared governance.**
   Looking at the future impact of these legislation efforts, we should also be cognizant of what our efforts to regulate algorithms might unleash. For that, we need to listen to the echoes of the past. We seem to be entering a period similar to that of the mid-1960's to early 1970's in which a group of liberal reformers—led by Ralph Nader—revealed the abuses and lack of safety standards in the automobile, meat, and coal-mining industries, among others. Their efforts resulted in the National Traffic and Motor Vehicle Safety Act (1966), the Natural Gas Pipeline Safety Act (1968), the Federal Coal Mine Health and Safety Act (1969), the Clean Air Act (1970), the Occupational Safety and Health Act (1970), and the Clean Water Act (1972). Is HR3611 the first of a similar flurry of Acts?

   While it is incontrovertible that those Acts improved the quality of life in the US, a recent book by historian Paul Sabin (2021) argues that they also achieved a less desirable outcome in the long term: they diminished our faith in government. Sabin argues that Nader and his team went not only after greedy corporations, but also after the corrupt and dysfunctional government agencies that protected them. Organizations like the Audubon Society and the Sierra Club acquired the legal status that allowed them to take the government and corporations to court on behalf of the public. The net result was a message that even conservatives could get behind: the government cannot be trusted. And while there are many aspects of the government that indeed should not be trusted, we have to consider whether tomorrow's well-meaning equivalent of the Sierra Club, in suing the government agencies that are unable to legislate increasingly complex algorithms, will be creating a crisis of confidence further down the road. In order to preempt that, we need to bolster legislation that promotes transparency and shared governance. Only through transparency and shared governance can we regain trust in the government to carry out or delegate—to machines and the people who make those machines—the power to make decisions that have a bearing on the rights of citizens.

5. **Develop stringent regulation on local and federal governments' ability to purchase automated decision systems from corporations.**
   Speaking of delegation, government contracts that involve the purchase of algorithms should be held to more stringent standards than contracts that involve the purchase of, say, pens or paper cups. Government has an obligation to demonstrate that the algorithmic tools it is buying will not harm citizens. The relationships between the government and social quantification sector suppliers need to be completely transparent and open to public scrutiny, not hidden in subcontracts or secret agreements.

   Litigation can be an important tool to keep the government in check in this regard, as long as we realize that the goal is "structural change within government agencies and their programs, rather than through isolated or narrow changes to specific [Algorithmic Decision Systems]" (Richardson et al., 2019, 4).

6. **Restrict the international sale of spyware.**
   The United States should lead the way in placing international restrictions on the traffic of spyware (such as the Israeli-manufactured Pegasus software) and other algorithmic weapons, which has led to rampant abuse through the targeting of minorities and dissidents. At home, government officials should continue the work of repealing the Patriot Act (Coleman 2020) and dismantling its unconstitutional surveillance infrastructure.

7. **Spend lots more money on public education.**
   We need to support legislation geared towards the improvement of public education in general, and the education of programmers and legislators in particular. We need to invest in short-term algorithmic literacy campaigns (such as my proposal for an Algorithm Observatory[11]), medium-term efforts to decolonize computer science and media curriculums, and more importantly, long-term efforts to re-introduce the humanities and arts at all levels of instruction, and in a meaningful way. Why are the humanities and arts important? Because the big questions about our roles and responsibilities in designing and using complex technologies will never go away, and we need to form individuals capable of confronting those big questions; individuals who are as well versed in the arts, humanities, and civics as they are in the sciences; individuals who understand this nation's racist history and how it continues to shape us and the tools we create today. Only a well-funded public education can help us do that. Unfortunately, that is not what we are currently offering young generations. In the absence of that, we can expect to continue to see the worst elements of the algorithmic state enforced over society for the benefit of the few.

References

Asher-Schapiro, Avi. 2020. "California City Bans Predictive Policing in U.S. First." Reuters, June 24, 2020, sec. everythingNews. https://www.reuters.com/article/us-usa-police-tech-trfn-idUSKBN23V2XC.

Barocas, Solon, Moritz Hardt, and Arvind Narayanan. 2019. Fairness and Machine Learning. Limitations and Opportunities (2019). fairmlbook.org.

Barr, Alistair. 2015. "Google Mistakenly Tags Black People as 'Gorillas,' Showing Limits of Algorithms." Wall Street Journal, July 1, 2015, sec. Digits. https://www.wsj.com/articles/BL-DGB-42522.

Benjamin, Ruha. 2019. Race After Technology: Abolitionist Tools for the New Jim Code. 1st edition. Medford, MA: Polity.

Bucher, Taina. 2018. If...Then: Algorithmic Power and Politics. New York: Oxford University Press.

Buolamwini, Joy, and Timnit Gebru. 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." In Proceedings of the 1st Conference on Fairness, Accountability and Transparency, 77–91. PMLR. https://proceedings.mlr.press/v81/buolamwini18a.html.

Chung, Jane. 2021. "Racism In, Racism Out." Public Citizen. https://www.citizen.org/article/algorithmic-racism/.

Coleman, Justine. 2020. "Massie, Gabbard Team up on Bill to Repeal the Patriot Act." Text. TheHill. December 16, 2020. https://thehill.com/homenews/house/530498-massie-gabbard-team-up-on-bill-to-repeal-the-patriot-act.

Costanza-Chock, Sasha. 2020. Design Justice: Community-Led Practices to Build the Worlds We Need. The MIT Press.

Couldry, Nick, and Ulises A. Mejias. 2019. The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism. 1st edition. Stanford, California: Stanford University Press.

Crawford, Kate. 2021. Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. New Haven: Yale University Press.

Dastin, Jeffrey. 2018. "Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women." Reuters, October 10, 2018, sec. Retail. https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G.

Diver, Laurence. 2021. "Digisprudence: The Design of Legitimate Code." Law, Innovation and Technology 0 (0): 1–30. https://doi.org/10.1080/17579961.2021.1977217.

Dunbar-Hester, Christina. 2019. Hacking Diversity: The Politics of Inclusion in Open Technology Cultures. Princeton, New Jersey: Princeton University Press.

Edgerton, Anna, and Bill Allison. 2021. "Big Tech Spent Millions on Lobbying Amid Antitrust Scrutiny." Bloomberg.Com, July 21, 2021. https://www.bloomberg.com/news/articles/2021-07-21/big-tech-spent-millions-on-lobbying-amid-antitrust-scrutiny.

Eubanks, Virginia. 2018. Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor. New York, NY: St. Martin's Press.

Feldstein, Steven. 2021. "Governments Are Using Spyware on Citizens. Can They Be Stopped?" Carnegie Endowment for International Peace. July 21, 2021. https://carnegieendowment.org/2021/07/21/governments-are-using-spyware-on-citizens.-can-they-be-stopped-pub-85019.

Furl, Nicholas, P. Jonathon Phillips, and Alice J O'Toole. 2002. "Face Recognition Algorithms and the Other-Race Effect: Computational Mechanisms for a Developmental Contact Hypothesis." Cognitive Science 26 (6): 797–815. https://doi.org/10.1016/S0364-0213(02)00084-8.

General, John, and Jon Sarlin. 2021. "A False Facial Recognition Match Sent This Innocent Black Man to Jail." CNN. April 29, 2021. https://www.cnn.com/2021/04/29/tech/nijeer-parks-facial-recognition-police-arrest/index.html.

Glaser, April. 2020. "Google, Amazon, and Microsoft Hold Thousands of Contracts with the U.S. Military and Federal Law Enforcement, New Research Shows." NBC News. July 8, 2020. https://www.nbcnews.com/tech/tech-news/thousands-contracts-highlight-quiet-ties-between-big-tech-u-s-n1233171.

Goh, Brenda. 2021. “Three Hours a Week: Play Time’s over for China’s Young Video Gamers.” Reuters, August 31, 2021, sec. China. https://www.reuters.com/world/china/china-rolls-out-new-rules-minors-online-gaming-xinhua-2021-08-30/.

Green, Ben, and Amba Kak. 2021. “The False Comfort of Human Oversight as an Antidote to A.I. Harm.” Slate, June 15, 2021. https://slate.com/technology/2021/06/human-oversight-artificial-intelligence-laws.html.

Hern, Alex. 2021. “Study Finds Growing Government Use of Sensitive Data to ‘Nudge’ Behaviour.” The Guardian, September 8, 2021, sec. Technology. https://www.theguardian.com/technology/2021/sep/08/study-finds-growing-government-use-of-sensitive-data-to-nudge-behaviour.

Isaac, Mike, and Daisuke Wakabayashi. 2017. “Russian Influence Reached 126 Million Through Facebook Alone.” The New York Times, October 30, 2017, sec. Technology. https://www.nytimes.com/2017/10/30/technology/facebook-google-russia.html.

Khatchadourian, Raffi. 2021. “Surviving the Crackdown in Xinjiang.” The New Yorker, April 5, 2021. https://www.newyorker.com/magazine/2021/04/12/surviving-the-crackdown-in-xinjiang.

Lessig, Lawrence. 2000. “Code Is Law.” Harvard Magazine. January 1, 2000. https://www.harvardmagazine.com/2000/01/code-is-law-html.

Lippman, Daniel, and Emily Birnbaum. 2021. “The Secret behind Amazon’s Domination in Cloud Computing.” Politico. June 4, 2021. https://www.politico.com/news/2021/06/04/amazon-hiring-former-government-officials-491878.

Matsui, Doris O. 2021. “H.R.3611 - 117th Congress (2021-2022): Algorithmic Justice and Online Platform Transparency Act.” Legislation. 2021/2022. May 28, 2021. https://www.congress.gov/bill/117th-congress/house-bill/3611/.

Mattu, Jeff, Julia Larson, Lauren Angwin, and Surya Kirchner. 2016. “How We Analyzed the COMPAS Recidivism Algorithm.” ProPublica. May 23, 2016. https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm?token=PrSqc58cn7gW-eRA_isLN03ygBJN8E84.

Moss, Emanuel, Elizabeth Ann Watkins, Ranjit Singh, Madeleine Clare Elish, and Jacob Metcalf. 2021. “Assembling Accountability: Algorithmic Impact Assessment for the Public Interest.” Data & Society. June 29, 2021. https://datasociety.net/library/assembling-accountability-algorithmic-impact-assessment-for-the-public-interest/.

Noble, Safiya Umoja. 2018. Algorithms of Oppression: How Search Engines Reinforce Racism. Illustrated edition. NYU Press.

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. “Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations.” Science 366 (6464): 447–53. https://doi.org/10.1126/science.aax2342.

Ongweso Jr, Edward. 2021a. “Big Tech Has Made Billions Off the 20-Year War on Terror.” Vice.Com. 09 2021. https://www.vice.com/en/article/4aveeq/big-tech-has-made-billions-off-the-20-year-war-on-terror.

———. 2021b. “The Gig Economy’s Business Model Is a Racial Justice Issue.” Vice.Com. March 25, 2021. https://www.vice.com/en/article/88nqd4/the-gig-economys-business-model-is-a-racial-justice-issue.

Paul, Katie, and Akanksha Rana. 2019. “U.S. Charges Facebook with Racial Discrimination in Targeted Housing Ads.” Reuters, March 28, 2019, sec. U.S. Legal News. https://www.reuters.com/article/us-facebook-advertisers-idUSKCN1R91E8.

Pisnau, Gaspar, and Veronica Arroyo. 2021. “Surveillance Tech in Latin America.” Access Now. https://www.accessnow.org/surveillance-tech-in-latin-america-made-abroad-deployed-at-home/.

Reisman, Dillon, Jason Schultz, Kate Crawford, and Meredith Whittaker. 2018. “Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability.” AI Now. https://ainowinstitute.org/aiareport2018.html.

Richardson, Rashida, Jason Schultz, and Vincent Southerland. 2019. “Litigating Algorithms 2019 US Report: New Challenges to Government Use of Algorithmic Decision Systems.” AI Now. https://ainowinstitute.org/litigatingalgorithms-2019-us.html.

Sabin, Paul. 2021. Public Citizens: The Attack on Big Government and the Remaking of American Liberalism. W. W. Norton & Company.

Samuel, Sigal. 2021. "AI's Islamophobia Problem." Vox. September 18, 2021. https://www.vox.com/future-perfect/22672414/ai-artificial-intelligence-gpt-3-bias-muslim.

Sap, Maarten, Dallas Card, Saadia Gabriel, Choi Yejin, and Noah Smith. 2019. "The Risk of Racial Bias in Hate Speech Detection." In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 1668–78. Florence, Italy. https://doi.org/10.18653/v1/P19-1163.

Schor, Juliet. 2020. "How the Gig Economy Promotes Inequality." Milken Institute Review. September 18, 2020. https://www.milkenreview.org/articles/how-the-gig-economy-promotes-inequality.

Schwartz, Brian. 2021. "Companies That Received Covid Government Contracts under Trump Administration Ramped up Lobbying." CNBC. April 14, 2021. https://www.cnbc.com/2021/04/14/companies-that-received-covid-government-contracts-under-trump-ramped-up-lobbying.html.

"Seattle Adopts Nation's Strongest Regulations for Surveillance Technology." 2017. ACLU of Washington. August 8, 2017. https://www.aclu-wa.org/news/seattle-adopts-nation%E2%80%99s-strongest-regulations-surveillance-technology.

Srnicek, Nick. 2016. Platform Capitalism. 1st edition. Cambridge, UK ; Malden, MA: Polity.

The Economist. 2017. "The World's Most Valuable Resource Is No Longer Oil, but Data," May 6, 2017. https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data.

Ulbricht, Lena, and Karen Yeung. 2021. "Algorithmic Regulation: A Maturing Concept for Investigating Regulation of and through Algorithms." Regulation & Governance, August. https://doi.org/10.1111/rego.12437.

Whittaker, Zack. 2021. "US Towns Are Buying Chinese Surveillance Tech Tied to Uighur Abuses." TechCrunch (blog). May 24, 2021. https://social.techcrunch.com/2021/05/24/united-states-towns-hikvision-dahua-surveillance/.

"Who's behind ICE? The Tech and Data Companies Fueling Deportations." 2018. Mijente, Immigrant Defense Project, and the National Immigration Project of the National Lawyers Guild. https://mijente.net/wp-content/uploads/2018/10/WHO%E2%80%99S-BEHIND-ICE_-The-Tech-and-Data-Companies-Fueling-Deportations-_v1.pdf.

Williams, Patrick, and Eric Kind. 2019. "Data-Driven Policing: The Hardwiring of Discriminatory Policing Practices across Europe." European Network Against Racism. https://www.citizensforeurope.eu/learn/data-driven-policing-the-hardwiring-of-discriminatory-policing-practices-across-europe.

Winston, Ali. 2018. "Palantir Has Secretly Been Using New Orleans to Test Its Predictive Policing Technology." The Verge. February 27, 2018. https://www.theverge.com/2018/2/27/17054740/palantir-predictive-policing-tool-new-orleans-nopd.

Woodruff, Grace. 2021. "Maine Now Has the Toughest Facial Recognition Restrictions in the U.S." Slate, July 2, 2021. https://slate.com/technology/2021/07/maine-facial-recognition-government-use-law.html.

Zuboff, Professor Shoshana. 2019. The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. S.l.: Profile Books.

Notes

1 Cf. Benjamin 2019, Bucher 2018, Crawford 2021, Noble 2018, Srnicek 2016, Zuboff 2019.

2 The terms "machine learning", "artificial intelligence", and "algorithm" will be used somewhat interchangeably in this paper, although they refer to different things. While there are different kinds of algorithms, this paper is concerned mostly with the kinds of algorithms used to build automated decision systems, that is, algorithms that can get computers to learn and make decisions without explicitly being programmed to make those decisions, thus behaving in an artificially "intelligent" way (machine learning is just one way of achieving artificial intelligence, although there are others, such as Natural Language Processing and automation).

3 Cf. Paul and Rana 2019, Dastin 2018, Samuel 2021, Larson et al., 2016, Barr 2015.

4 "Mass Surveillance in the United States," Wikipedia, https://en.wikipedia.org/wiki/List_of_government_mass_surveillance_projects#United_States

5 Legal Information Institute, 47 U.S. code § 230, https://www.law.cornell.edu/uscode/text/47/230

6 The Trusted Cloud for Government," Amazon Web Services, https://aws.amazon.com/government-education/government/

7 Two other seemingly simple solutions are commonly mentioned: to promote the use of open source technologies, and to pay people for the data they produce. The former ignores that open source software can be used just as well for extractivist purposes, as Google's Android operating system demonstrates. The latter ignores that paying people for their data would leave the extractivist system intact, and would, in a way, legitimize it.

8 For example: Costanza-Chock 2020; Dunbar-Hester 2019.

9 The Canadian Government, for instance, has released a tool that allows agencies to conduct an AIA by simply filling out an online survey. See https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html

10 See https://www.ai.gov/nairrtf/

11 Part media literacy project and part citizen science experiment, Algorithm Observatory is a proposal for a virtual laboratory that will allow the public to engage in the analysis of social computing algorithms. Any algorithm can be studied because the experiments do not require direct access to the algorithm's source code (which is usually proprietary). Instead, the analysis is conducted through empirical observation of the interaction between the algorithm and different kinds of volunteer participants.