2
Default Discrimination
Is the Glitch Systemic?

GLITCH

- a minor problem
- a false or spurious electronic signal
- a brief or sudden interruption or irregularity
- may derive from Yiddish, glitsh – to slide, glide, "slippery place."[1]

When Princeton University media specialist Allison Bland was driving through Brooklyn, the Google Maps narrator directed her to "turn right on Malcolm Ten Boulevard," verbally interpreting the X in the street name as a Roman numeral rather than as referring to the Black liberation leader who was assassinated in New York City in 1965 (Figure 2.1).

Social and legal codes, like their byte-size counterparts, are not neutral; nor are all codes created equal. They reflect particular perspectives and forms of social organization that allow some people to assert themselves – their assumptions, interests, and desires – over others. From the seemingly mundane to the extraordinary, technical systems offer a mirror to the wider terrain of struggle over the forces that govern our lives.

Figure 2.1 Malcolm Ten

Source: Twitter @alliebland, November 19, 2013, 9:42 p.m.

Database design, in that way, is “an exercise in worldbuilding,” a normative process in which programmers are in a position to project their world views – a process that all too often reproduces the technology of race.[2] Computer systems are a part of the larger matrix of systemic racism. Just as legal codes are granted an allure of objectivity – “justice is (color)blind” goes the fiction – there is enormous mystique around computer codes, which hides the human biases involved in technical design.

The Google Maps glitch is better understood as a form of displacement or digital gentrification mirroring the widespread dislocation underway in urban areas across the United States. In this case, the cultural norms and practices of programmers – who are drawn from a narrow racial, gender, and classed demographic – are coded into technical systems that, literally, tell people where to go. These seemingly innocent directions, in turn, reflect and reproduce racialized commands that instruct people where they belong in the larger social order.[3]

Ironically, this problem of misrecognition actually reflects a solution to a difficult coding challenge. A computer’s ability to parse Roman numerals, interpreting an “X” as “ten,” was a hard-won design achievement.[4] That is, from a strictly technical standpoint, “Malcolm Ten Boulevard” would garner cheers. This illustrates how innovations reflect the priorities and concerns of those who frame the problems to be solved, and how such solutions may reinforce forms of social dismissal, regardless of the intentions of individual programmers.

While most observers are willing to concede that technology can be faulty, acknowledging the periodic breakdowns and “glitches” that arise, we must be willing to dig deeper.[5] A narrow investment in technical innovation necessarily displaces a broader set of social interests. This is more than a glitch. It is a form of exclusion and subordination built into the ways in which priorities are established and solutions defined in the tech industry. As Andrew Russell and Lee Vinsel contend, “[t]o take the place of progress, ‘innovation,’ a smaller, and morally neutral, concept arose. Innovation provided a way to celebrate the accomplishments of a high-tech age without expecting too much from them in the way of moral and social improvement.”[6] For this reason, it is important to question “innovation” as a straightforward social good and to look again at what is hidden by an idealistic vision of technology. How is technology already raced?

This chapter probes the relationship between glitch and design, which we might be tempted to associate with competing conceptions of racism. If we think of racism as something of the past or requiring a particular visibility to exist, we can miss how the New Jim Code operates and what seeming glitches reveal about the structure of racism. Glitches are generally considered a fleeting interruption of an otherwise benign system, not an enduring and constitutive feature of social life. But what if we understand glitches instead to be a slippery place (with reference to the possible Yiddish origin of the word) between fleeting and durable, micro-interactions and macro-structures, individual hate and institutional indifference? Perhaps in that case glitches are not spurious, but rather a kind of signal of how the system operates. Not an aberration but a form of evidence, illuminating underlying flaws in a corrupted system.

Default Discrimination

At a recent workshop sponsored by a grassroots organization called Stop LAPD Spying, the facilitator explained that community members with whom she works might not know what algorithms are, but they know what it feels like to be watched. Feelings and stories of being surveilled are a form of "evidence," she insisted, and community testimony is data.[7] As part of producing those data, the organizers interviewed people about their experiences with surveillance and their views on predictive policing. They are asked, for example: "What do you think the predictions are based on?" One person, referring to the neighborhood I grew up in, responded:

> Because they over-patrol certain areas – if you're only looking on Crenshaw and you only pulling Black people over then it's only gonna make it look like, you know, whoever you pulled over or whoever you searched or whoever you criminalized that's gonna be where you found something.[8]

Comments like this remind us that people who are most directly impacted by the New Jim Code have a keen sense of the default discrimination facilitated by these technologies. As a form of social technology, institutional racism, past and present, is the precondition for the carceral technologies that underpin the US penal system. At every stage of the process – from policing, sentencing, and imprisonment to parole – automated risk assessments are employed to determine people's likelihood of committing a crime.[9] They determine the risk profile of neighborhoods in order to concentrate police surveillance, or the risk profile of individuals in order to determine whether or for how long to release people on parole.

In a recent study of the recidivism risk scores assigned to thousands of people arrested in Broward County, Florida, ProPublica investigators found that the score was remarkably unreliable in forecasting violent crime. They also uncovered significant racial disparities:

> In forecasting who would re-offend, the algorithm made mistakes with black and white defendants at roughly the same rate but in very different ways. The formula was particularly likely to falsely flag black defendants as future criminals, wrongly labeling them this way at almost twice the rate as white defendants. White defendants were mislabeled as low risk more often than black defendants.[10]

The algorithm generating the risk score builds upon already existing forms of racial domination and reinforces them precisely because the apparatus ignores how race shapes the "weather." Literary scholar Christina Sharpe describes the weather as "the total climate; and the climate is antiblack."[11] For example, the survey given to prospective parolees to forecast the likelihood that they will recidivate includes questions about their criminal history, education and employment history, financial history, and neighborhood characteristics (among many other factors). As all these variables are structured by racial domination – from job market discrimination to ghettoization – the survey measures the extent to which an individual's life chances have been impacted by racism without ever asking an individual's race.[12]

Likewise, predictive policing software will always be more likely to direct police to neighborhoods like the one I grew up in, because the data that this software is drawing from reflect ongoing surveillance priorities that target predominantly Black neighborhoods.[13] Anti-Blackness is no glitch. The system is accurately rigged, we might say, because, unlike in natural weather forecasts, the weathermen are also the ones who make it rain.[14]

Even those who purportedly seek "fairness" in algorithmic decision-making are not usually willing to assert that the benchmark for whether an automated prediction is "unwarranted" is whether it strays from the proportion of a group in the larger population. That is, if a prediction matches the current crime rate, it is still unjust! Even so, many who are grappling with how to enact ethical practices in this arena still use the crime rate as the default measure of whether an algorithm is predicting fairly, when that very measure is a byproduct of ongoing regimes of selective policing and punishment.[15]

Default Discrimination

$$\frac{\delta A}{\delta t} = B + \frac{\eta D}{4}\nabla^2 A - \omega A + \theta\omega\delta$$

Figure 2.2 Patented PredPol Algorithm

Source: http://www.predpol.com/technology

Interestingly, the most commonly used algorithm in Los Angeles and elsewhere, called PredPol, is drawn directly from a model used to predict earthquake aftershocks (Figure 2.2). As author of *Carceral Capitalism*, Jackie Wang gives us this description: "In police departments that use PredPol, officers are given printouts of jurisdiction maps that are covered with red square boxes that indicate where crime is supposed to occur throughout the day … The box is a kind of *temporary crime zone*." She goes on to ask:

> What is the attitude or mentality of the officers who are patrolling one of the boxes? When they enter one of the boxes, do they expect to stumble upon a crime taking place? How might the expectation of finding crime influence what the officers actually find? Will people who pass through these temporary crime zones while they are being patrolled by officers automatically be perceived as suspicious? Could merely passing through one of the red boxes constitute probable cause?[16]

Let me predict: yes. If we consider that institutional racism in this country is an ongoing unnatural disaster, then crime prediction algorithms should more accurately be called crime *production* algorithms. The danger with New Jim Code predictions is the way in which self-fulfilling prophecies enact what they predict, giving the allure of accuracy. As the man behind PredPol's media strategy put it, "it sounds like fiction, but its more like science fact."[17]

Predicting Glitches

One of the most iconic scenes from *The Matrix* film trilogy deals with the power of predictions and self-fulfilling prophecies. The main protagonist, Neo, goes to visit the Oracle, a software program depicted as a Black woman in her late sixties. Neo is trying to figure out whether he is who others think he is – "the one" who is supposed to lead humanity in the war against the machines. As he tries to get a straight answer from the Oracle and to figure out whether she really has the gift of prophecy, she says, "I'd ask you to sit down, but you're not going to anyway. *And don't worry about the vase*."

> NEO: What vase? [Neo knocks a vase to the floor]
>
> THE ORACLE: That vase.
>
> NEO: I'm sorry.
>
> THE ORACLE: I said don't worry about it. I'll get one of
>
> my kids to fix it.
>
> NEO: How did you know?
>
> THE ORACLE: What's really going to bake your noodle later on is, *would you still have broken it if I hadn't said anything*.[18]

This scene invites a question about real-life policing: Would cops still have warrants to knock down the doors in majority Black neighborhoods if predictive algorithms hadn't said anything?

The Matrix offers a potent allegory for thinking about power, technology, and society. It is set in a dystopian future in which machines overrun the world, using the energy generated by human brains as a vital source of computing power. Most of humanity is held captive in battery-like pods, their minds experiencing an elaborate life-like simulation of the real world in order to pacify humans and maximize the amount of energy brains produce. The film follows a small band of freedom fighters who must convince Neo that the simulated life he was living is in fact a digital construction.

Early on in his initiation to this new reality, Neo experiences a fleeting moment of déjà vu when a black cat crosses his path – twice. Trinity, his protector and eventual love interest, grows alarmed and explains that this "glitch in the matrix" is not at all trivial but a sign that something about the program has been changed by the agents of the Matrix. The sensation of déjà vu is a warning sign that a confrontation is imminent and that they should prepare to fight.

The film's use of déjà vu is helpful for considering the relationship between seemingly trivial technical glitches and meaningful design decisions. The glitch in this context is a not an insignificant "mistake" to be patched over, but rather serves as a signal of something foundational about the structure of the world meant to pacify humans. It draws attention to the construction and reconstruction of the program and functions as an indication that those seeking freedom should be ready to spring into action.

A decade before the *Matrix* first hit the big screen, Black feminist theorist Patricia Hill Collins conceptualized systemic forms of inequality in terms of a "matrix of domination" in which race, class, gender, and other axes of power operated together, "as sites of domination and as potential sites of resistance."[19] This interlocking matrix operates at individual, group, and institutional levels, so that empowerment "involves rejecting the dimensions of knowledge, whether personal, cultural, or institutional, that perpetuate objectification and dehumanization."[20] Relating this dynamic to the question of how race "gets inside" technology, the Roman numeral glitch of Google Maps and others like it urge us to look again at the way our sociotechnical systems are constructed – by whom and to what ends.

Racist glitches – such as celebrity chef Paula Dean's admission that "yes, of course" she has used the N-word alongside her desire to host a "really southern plantation wedding" with all-Black servers;[21] or a tape-recorded phone call in which former Los Angeles Clippers owner and real estate mogul Donald Sterling told a friend "[i]t bothers me a lot that you want to broadcast that you're associating with black people"[22] – come and go, as provocative sound bites muffling a deeper social reckoning. In my second example, the scandal associated with Sterling's racist remarks stands in stark contrast with the hush and acceptance of a documented pattern of housing discrimination exercised over many years, wherein he refused to rent his properties to Black and Latinx tenants in Beverly Hills and to non-Korean tenants in LA's Koreatown.[23] In the midst of the suit brought by the Department of Justice, the Los Angeles chapter of the National Association for the Advancement of Colored People nevertheless honored Sterling with a lifetime achievement award in 2009. Only once his tape-recorded remarks went public in 2014 did the organization back out of plans to award him this highest honor for a second time, forcing the chapter president to resign amid criticism.

Dragging individuals as objects of the public condemnation of racist speech has become a media ritual and pastime. Some may consider it a distraction from the more insidious, institutionalized forms of racism typified by Sterling's real estate practices. The déjà vu regularity of all those low-hanging N-words would suggest that stigmatizing individuals is not much of a deterrent and rarely addresses all that gives them license and durability.

But, as with Trinity's response to Neo in the *Matrix* regarding his path being crossed twice by a black cat, perhaps if we situated racist "glitches" in the larger complex of social meanings and structures, we too could approach them as a signal rather than as a distraction. Sterling's infamous phone call, in this case, would alert us to a deeper pattern of housing discrimination, with far-reaching consequences.

Systemic Racism Reloaded

Scholars of race have long challenged the focus on individual "bad apples," often to be witnessed when someone's racist speech is exposed in the media – which is typically followed by business as usual.[24] These individuals are treated as glitches in an otherwise benign system. By contrast, sociologists have worked to delineate how seemingly neutral

policies and norms can poison the entire "orchard" or structure of society, systematically benefiting some while subjugating others.[25]

Whereas racist glitches are often understood as transient, as signals they can draw our attention to discriminatory design as a durable feature of the social landscape since this nation's founding. As sociologists Joe Feagin and Sean Elias write, "[i]n the case of US society, systemic racism is foundational to and engineered into its major institutions and organizations."[26] This reorientation is also exemplified by Eduardo Bonilla-Silva's *Racism without Racists*, in which he defines "racialized social systems, or white supremacy for short … as the totality of the social relations and practices that reinforce white privilege. Accordingly, the task of analysts interested in studying racial structures is to uncover the particular social, economic, political, social control, and ideological mechanisms responsible for the reproduction of racial privilege in a society."[27]

Taken together, this work builds upon the foundational insights of Charles V. Hamilton and Kwame Ture (née Stokely Carmichael), who developed the term "institutional racism" in 1967. While the authors discuss the linkage between institutional racism and what they describe as individual racism, they also state:

> This is not to say that every single white American consciously oppresses black people. He does not need to. Institutional racism has been maintained deliberately by the power structure and through indifference, inertia, and lack of courage on the part of the white masses as well as petty officials … The line between purposeful suppression and indifference blurs.[28]

But taking issue with the overwhelming focus on top-down forces that characterize work on systemic racism, including Feagin and Elias' "theory of oppression," Michael Omi and Howard Winant highlight the agency and resistance of those subordinated by such systems. They say:

> To theorize racial politics and the racial state, then, is to enter the complex territory where structural racism encounters self-reflective action, the radical practice of people of color (and their white allies) in the United States. It is to confront the instability of the US system of racial hegemony, in which despotism and democracy coexist in seemingly permanent conflict.[29]

Strikingly, throughout this early work on institutional racism and structural inequality, there was very little focus on the role of technologies, beyond mass media, in advancing or undermining racial ideologies and structures. As Jessie Daniels notes in "Race and Racism in Internet Studies":

> The role of race in the development of Internet infrastructure and design has largely been obscured (Taborn, 2008). As Sinclair observes, "The history of race in America has been written as if technologies scarcely existed, and the history of technology as if it were utterly innocent of racial significance."[30]

Daniels' (2009) *Cyber Racism* illuminates how "white supremacy has entered the digital era"

while acknowledging how those "excluded by the white-dominated mainstream media" also use the Internet for grassroots organizing and antiracist discourse.[31] In so doing, she challenges both those who say that technology is only a "source of danger" when it comes to the active presence of White supremacists online and those who assume that technology is "inherently democratizing."[32] Daniels echoes Nakamura's (2002, 2008) frustration with how race remains undertheorized in Internet studies and urges more attention to the technology of structural racism. In line with the focus on glitches, researchers tend to concentrate on how the Internet perpetuates or mediates racial prejudice at the individual level rather than analyze how racism shapes infrastructure and design. And, while Daniels does not address this problem directly, an investigation of how algorithms perpetuate or disrupt racism should be considered in any study of discriminatory design.

Architecture and Algorithms

On a recent visit that I made to University of California at San Diego, my hosts explained that the design of the campus made it almost impossible to hold large outdoor gatherings. The "defensive" architecture designed to prevent skateboarding and cycling in the interest of pedestrians also deliberately prevented student protests at a number of campuses following the Berkeley free speech protests in the mid-1960s. This is not so much a trend in urban planning as an ongoing feature of stratified societies. For some years now, as I have been writing and thinking about discriminatory design of all sorts, I keep coming back to the topic of public benches: benches I tried to lie down on but was prevented because of intermittent arm rests, then benches with spikes that retreat after you feed the meter, and many more besides.

Like the discriminatory designs we are exploring in digital worlds, hostile architecture can range from the more obvious to the more insidious – like the oddly shaped and artistic-looking bench that makes it uncomfortable but not impossible to sit for very long. Whatever the form, hostile architecture reminds us that public space is a permanent battleground for those who wish to reinforce or challenge hierarchies. So, as we explore the New Jim Code, we can observe connections in the building of physical and digital worlds, even starting with the use of "architecture" as a common metaphor for describing what algorithms – those series of instructions written and maintained by programmers that adjust on the basis of human behavior – build. But, first, let's take a quick detour …

The era commonly called "Jim Crow" is best known for the system of laws that mandated racial segregation and upheld White supremacy in the United States between 1876 and 1965. Legal codes, social codes, and building codes intersected to keep people separate and unequal. The academic truism that race is "constructed" rarely brings to mind these concrete brick and mortar structures, much less the digital structures operating today. Yet if we consider race as itself a technology, as a means to sort, organize, and design a social structure as well as to understand the durability of race, its consistency and adaptability, we can understand more clearly the literal architecture of power.

Take the work of famed "master builder" Robert Moses, who in the mid-twentieth century built hundreds of structures, highways, bridges, stadiums, and more, prioritizing suburbanization and upper-middle-class mobility over public transit and accessibility to poor and working-class New Yorkers. In a now iconic (yet still disputed) account of Moses' approach to public works, science and technology studies scholar Langdon Winner describes the low-hanging overpasses that line the Long Island parkway system. In Winner's telling, the design prevented buses from using the roads, which enabled predominantly White, affluent car owners to move freely, while working-class and non-White people who relied on buses were prevented from accessing the suburbs and the beaches. And while the veracity of Winner's account continues to be debated, the parable has taken on a life of its own, becoming a narrative tool for illustrating how artifacts "have politics."[33]

For our purpose, Moses' bridges symbolize the broader architecture of Jim Crow. But, whereas Jim Crow laws explicitly restricted Black people from numerous "White only" spaces and services, the physical construction of cities and suburbs is central to the exercise of racial power, including in our postcivil rights era. And, while some scholars dispute whether Moses intended to exclude Black people from New York suburbs and beaches, one point remains clear: the way we engineer the material world reflects and reinforces (but could also be used to subvert) social hierarchies.

Yet plans to engineer inequity are not foolproof. In April 2018 a group of high school students and their chaperones returning from a spring break trip to Europe arrived at Kennedy Airport and boarded a charter bus that was headed to a Long Island shopping center where parents waited to pick up their kids. As they drove to the mall, the bus driver's navigation system failed to warn him about the low-hanging bridges that line the Long Island parkway and the bus slammed violently into the overpass, crushing the roof, seriously wounding six, and leaving dozens more injured. As news reports pointed out, this was only the latest of hundreds of similar accidents that happened over the years, despite numerous warning signs and sensor devices intended to alert oncoming traffic of the unusually low height of overpasses. Collateral damage, we might say, is part and parcel of discriminatory design.

From what we know about the people whom city planners have tended to prioritize in their designs, families such as the ones who could send their children to Europe for the spring break loom large among them. But a charter bus with the roof shaved off reminds us that tools of social exclusion are not guaranteed to impact only those who are explicitly targeted to be disadvantaged through discriminatory design. The best-laid plans don't necessarily "stay in their lane," as the saying goes. Knowing this, might it be possible to rally more people against social and material structures that immobilize some to the benefit of others? If race and other axes of inequity are constructed, then perhaps we can construct them differently?

When it comes to search engines such as Google, it turns out that online tools, like racist robots, reproduce the biases that persist in the social world. They are, after all, programmed using algorithms that are constantly updated on the basis of human behavior and are learning

and replicating the technology of race, expressed in the many different associations that the users make. This issue came to light in 2016, when some users searched the phrase "three Black teenagers" and were presented with criminal mug shots. Then when they changed the phrase to "three White teenagers," users were presented with photos of smiling, go-lucky youths; and a search for "three Asian teenagers" presented images of scantily clad girls and women. Taken together, these images reflect and reinforce popular stereotypes of Black criminality, White innocence, or Asian women's sexualization that underpin much more lethal and systemic forms of punishment, privilege, and fetishism respectively.[34] The original viral video that sparked the controversy raised the question "Is Google being racist?," followed by a number of analysts who sought to explain how these results were produced:

> The idea here is that computers, unlike people, can't be racist but we're increasingly learning that they do in fact take after their makers … Some experts believe that this problem might stem from the hidden biases in the massive piles of data that algorithms process as they learn to recognize patterns … reproducing our worst values.[35]

According to the company, Google itself uses "over 200 unique signals or 'clues' that make it possible to guess what you might be looking for."[36] Or, as one observer put it, "[t]he short answer to why Google's algorithm returns racist results is that society is racist."[37] However, this does not mean that we have to wait for a social utopia to float down from the clouds before expecting companies to take action. They are already able to optimize online content in ways that mitigate bias. Today, if you look up the keywords in Noble's iconic example, the phrase "Black girls" yields images of Black Girls Code founder Kimberly Bryant and #MeToo founder Tarana Burke, along with images of organizations like Black Girls Rock! (an awards show) and Black Girls Run (a wellness movement). The technical capacity was always there, but social awareness and incentives to ensure fair representation online were lacking. As Noble reports, the pornography industry has billions of dollars to throw at companies in order to optimize content, so advertising cannnot continue to be the primary driver of online content. Perhaps Donald Knuth's proverbial warning is true: "premature optimization is the root of all evil."[38] And so the struggle to democratize information gateways continues.[39]

A number of other examples illustrate algorithmic discrimination as an ongoing problem. When a graduate student searched for "unprofessional hairstyles for work," she was shown photos of Black women; when she changed the search to "professional hairstyles for work," she was presented with photos of White women.[40] Men are shown ads for high-income jobs much more frequently than are women, and tutoring for what is known in the United States as the Scholastic Aptitude Test (SAT) is priced more highly for customers in neighborhoods with a higher density of Asian residents: "From retail to real estate, from employment to criminal justice, the use of data mining, scoring and predictive software … is proliferating … [And] when software makes decisions based on data, like a person's zip code, it can reflect, or even amplify, the results of historical or institutional discrimination."[41]

A team of Princeton researchers studying associations made with Black-sounding names and

White-sounding names confirmed findings from employment audit studies[42] to the effect that respondents make negative associations with Black names and positive associations with White ones. Caliskan and colleagues show that widely used language-processing algorithms trained on human writing from the Internet reproduce human biases along racist and sexist lines.[43] They call into question the assumption that computation is pure and unbiased, warning that, "if we build an intelligent system that learns enough about the properties of language to be able to understand and produce it, in the process it will also acquire historic cultural associations, some of which can be objectionable. Already, popular online translation systems incorporate some of the biases we study … Further concerns may arise as AI is given agency in our society."[44] And, as we shall see in the following chapters, the practice of codifying existing social prejudices into a technical system is even harder to detect when the stated purpose of a particular technology is to override human prejudice.

Notes

1. Merriam-Webster Online, n.d.

2. Personal interview conducted by the author with Princeton digital humanities scholar Jean Bauer, October 11, 2016.

3. See references to "digital gentrification" in "White Flight and Digital Gentrification," posted on February 28 at https://untsocialmedias13.wordpress.com/2013/02/28/white-flight-and-digital-gentrification by jalexander716.

4. Sampson 2009.

5. As Noble (2018, p. 10) writes, "[a]lgorithmic oppression is not just a glitch in the system but, rather, is fundamental to the operating system of the web."

6. Russell and Vinsel 2016.

7. See the conference "Dismantling Predictive Policing in Los Angeles," May 8, 2018, at https://stoplapdspying.org/wp-content/uploads/2018/05/Before-the-Bullet-Hits-the-Body-May-8-2018.pdf.

8. "Dismantling predictive policing in Los Angeles," pp. 38–9.

9. Ferguson 2017.

10. Angwin et al. 2016.

11. According to Sharpe (2016, p. 106), "the weather necessitates changeability and improvisation," which are key features of innovative systems that adapt, in this case, to postracial norms where racism persists through the absence of race.

12. Meredith Broussard, data journalist and author of *Artificial Unintelligence*, explains:

"The fact that nobody at Northpointe thought that the questionnaire or its results might be biased has to do with technochauvinists' unique worldview. The people who believe that math and computation are 'more objective' or 'fairer' tend to be the kind of people who think that inequality and structural racism can be erased with a keystroke. They imagine that the digital world is different and better than the real world and that by reducing decisions to calculations, we can make the world more rational. When development teams are small, like-minded, and not diverse, this kind of thinking can come to seem normal. However, it doesn't move us toward a more just and equitable world" (Broussard 2018, p. 156).

13. Brayne 2014.

14. As Wang (2018, p. 236) puts it, "the rebranding of policing in a way that foregrounds statistical impersonality and symbolically removes the agency of individual officers is a clever way to cast police activity as neutral, unbiased, and rational. This glosses over the fact that using crime data gathered by the police to determine where officers should go simply sends police to patrol the poor neighborhoods they have historically patrolled when they were guided by their intuitions and biases. This 'new paradigm' is not merely a reworking of the models and practices used by law enforcement, but a revision of the police's public image through the deployment of science's claims to objectivity."

15. I am indebted to Naomi Murakawa for highlighting for me the strained way in which scholars and criminologists tend to discuss "unwarranted disproportion," as if the line between justified and unjustified is self-evident rather than an artifact of racist policing, with or without the aid of crime prediction software. See Murakawa 2014.

16. Wang 2018, p. 241.

17. Wang 2018, p. 237.

18. From sci fi quotes.net, http://scifiquotes.net/quotes/123_Dont-Worry-About-the-Vase; emphasis added.

19. Collins 1990, p. 227.

20. Collins 1990, p. 230.

21. Goodyear 2013.

22. Goyette 2014.

23. Associated Press 2006.

24. Daniels 2013, p. 709.

25. Golash-Boza 2016.

26. Feagin and Elias, 2013, p. 936.

27. Bonilla-Silva 2006, p. 9.

28. Hamilton and Ture 1967, p. 38. Scholar of African American studies Keeanga-Yamahtta Taylor describes the term “institutional racism” as prescient, noting that “it is the outcome that matters, not the intentions of the individuals involved” (Taylor 2016, p. 8).

29. Omi and Winant 1994, pp. 137–8.

30. Sinclair 2004, p. 1; cf. Daniels 2013, p. 696.

31. Daniels 2009, p. 2.

32. Daniels 2009, p. 4.

33. Winner 1980.

34. Helm 2016.

35. Pearson 2016a.

36. See “How search algorithms work,” https://www.google.co.uk/insidesearch/howsearchworks/algorithms.html.

37. See Chiel 2016; in its own defense, the company explained thus: “‘Our image search results are a reflection of content from across the web, including the frequency with which types of images appear and the way they’re described online,’ a spokesperson told the Mirror. This means that sometimes unpleasant portrayals of sensitive subject matter online can affect what image search results appear for a given query. These results don’t reflect Google’s own opinions or beliefs – as a company, we strongly value a diversity of perspectives, ideas and cultures.”

38. Roberts 2018.

39. Sociologist Zeynep Tufekci (2019) puts it thus: “These companies – which love to hold themselves up as monuments of free expression – have attained a scale unlike anything the world has ever seen; they’ve come to dominate media distribution, and they increasingly stand in for the public sphere itself. But at their core, their business is mundane: They’re ad brokers. To virtually anyone who wants to pay them, they sell the capacity to precisely target our eyeballs. They use massive surveillance of our behavior, online and off, to generate increasingly accurate, automated predictions of what advertisements we are most susceptible to and what content will keep us clicking, tapping, and scrolling down a bottomless feed.”

40. Chiel 2016.

41. Kirchner 2015a.

42. Bertrand and Mullainathan 2003.

43. Pearson 2016a.

44. Caliskan et al. 2017, p. 186.