Chapter One: Bring Back the Bodies

Chapter One: Bring Back the Bodies

When Serena Williams disappeared from Instagram in early September, 2017, her six million followers thought they knew what had happened. Several months earlier, in March of that year, Williams had accidentally announced her pregnancy to the world via a bathing suit selfie and a caption that was hard to misinterpret: "20 weeks." Now, they assumed, her baby had finally arrived.

But then they waited, and waited some more. Two weeks later, Williams finally re-appeared on Instagram, announcing the birth of her daughter and inviting her followers to watch a video that welcomed Alexis Olympia Ohanian Jr. to the world. A montage of baby bump pics interspersed with clips of a pregnant Williams playing tennis and cute conversations with her husband, Reddit cofounder Alexis Ohanian, segued into the shot that her fans had been waiting for: the first of baby Olympia. Williams was narrating: "So we're leaving the hospital," she explains. "It's been a long time. We had a lot of complications. But look who we got!" The scene fades to white, and ends with a set of stats: Olympia's date of birth, birth weight, and number of grand slam titles: 1. (Williams, as it turned out, was already eight weeks pregnant when she won the Australian Open earlier that year).

Williams's Instagram followers were, for the most part, enchanted. But a fair number of her followers-- many of them Black women like Williams herself-- fixated on the comment she'd made as she was heading home from the hospital with her baby girl. Those "complications" that Williams mentioned-- they'd had them too.

On Williams's Instagram feed, the evidence was anecdotal--women posting about their own experience of childbirth gone horribly wrong. But a few months later, Williams returned to social media--Facebook, this time--armed with data. Citing a 2017 study from the US Centers for Disease Control and Prevention (CDC), Williams wrote that "Black women are over 3 times more likely than white women to die from pregnancy- or childbirth-related causes."

A Facebook post by Serena Williams responding to her Instagram followers who had shared their stories of pregnancy and childbirth-related complications with her. ¶ Credit: Serena Williams ¶ Source: https://www.facebook.com/SerenaWilliams/videos/10156086135726834/

While these disparities were well known to Black women-led reproductive justice groups like Sister Song, the Black Mamas Matter Alliance, and Raising Our Sisters Everywhere, as well as to feminist scholars across a range of disciplines, Williams helped to shine a national spotlight on them. And she wasn't the only one. A few months earlier, Nina Martin of the investigative journalism outfit ProPublica, working with Renee Montagne of NPR, had reported on the same phenomenon. "Nothing Protects Black Women From Dying in Pregnancy and Childbirth," the headline read. In addition to the study also cited by Williams, Martin and Montagne cited a second study from 2016 which showed that neither education nor income level-- the factors usually invoked when attempting to account for healthcare outcomes that diverge along racial lines-- impacted the fates of Black women giving birth. On the contrary, the data showed that Black women with college degrees suffered more severe complications of pregnancy and childbirth than white women without high school diplomas.

But what were these complications, more precisely? And how many women had actually died as a result? ProPublica couldn't find out, and neither could *USA Today*, which took up the issue a year later to see what, after a year of increased attention and advocacy, had changed. What they found was that there was still no national system for tracking complications sustained in pregnancy and childbirth, even as similar systems have long been in place for tracking things like, for instance, teen pregnancy, hip replacements, and heart attacks. They also found that there is also still no reporting mechanism for ensuring that hospitals follow national childbirth safety standards, as is required for both hip surgery and cardiac care. “Our maternal data is embarrassing,” stated Stacie Geller, a professor obstetrics and gynecology at the University of Illinois, when asked for comment. The Chief of the CDC’s Maternal and Infant Health Branch, William Callaghan, makes the significance of this “embarrassing” data more clear: “What we choose to measure is a statement of what we value in health,” he explains. We might edit his statement to add: it’s a measure of *who* we value in health, too.

The lack of data about maternal health outcomes, and its impact on matters of life and death, underscores how it is *people* who end up affected by the choices we make in our practices of data collection, analysis, and communication. More than that, it’s almost always the bodies of those who have been disempowered by forces they cannot control, such as sexism, racism, or classism--or, more likely, the intersection of all three--who experience the most severe consequences of these choices. Serena Williams acknowledged this exact phenomenon when asked by *Glamour* magazine about the statistics she cited in her Facebook post. “If I wasn’t who I am, it could have been me—” she said, referring to the fact that she had to demand that her medical team perform additional tests in order to diagnose her own postnatal complications, and because she was Serena Williams, 23-time grand slam champion, they listened. But, she told *Glamour*, “that’s not fair.”

It is absolutely not fair. But without a significant intervention into our current data practices, this unfairness--and many other inequities with issues of power and privilege at their core-- will continue to get worse. Stopping that downward spiral is the real reason we wrote this book. We wrote this book because we are data scientists and data feminists. We think that data science and the fields that rely upon it stand to learn significantly from feminist writing, thinking, scholarship, and action.

1

Feminism is one key conceptual orientation that can help mitigate inequality and work towards justice, but it is not the only one. We talk about some others in *Now Let's Multiply*.

As we explain in *Why Data Science Needs Feminism*, feminism isn’t only about women. It isn’t even only about issues of gender. Feminism is about power--about who has it, and who doesn’t. In a world in which data is power, and that power is wielded unequally, feminism can help us better understand how it operates and how it can be challenged. As data feminists--a group that includes women, men, nonbinary and genderqueer people, and everyone else--we can take steps, together, towards a more just and equal world.

A good starting point is to understand how power operates on bodies and through them. "But!" you might say. "Data science is premised on things like objectivity and neutrality! And those things have nothing to do with bodies!" But that is precisely the point. Data science, as it is generally understood in the world today, has very little to do with bodies. But that is a fundamental misconception about the field, and about data more generally. Because even though we don't see the bodies that data science is reliant upon, it most certainly relies upon them. It relies upon them as the sources of data, and it relies upon them to make decisions about data. As we discuss more in depth in a couple of pages, it even relies on them to decide what concepts like "objective" and "neutral" really mean. And when not all bodies are represented in those decisions-- as in the case of the federal and state legislatures which might fund data collection on maternal mortality--well, that's when problems enter in.

What kind of problems? Structural ones. Structural problems refer to problems that are systemic in nature, rather than due to a specific point (or person) of origin. It might be counterintuitive to think that individual bodies can help expose structural problems, but that's precisely what the past several decades--centuries, even--of feminist activism and critical thought has allowed us to see. Because many of the problems that individual people face are often the result of larger systems of power, but they remain invisible until those people bring them to light. In a contemporary context, we might easily cite the #MeToo movement as an example of how individual experience, taken together, reveals a larger structural problem of sexual harassment and assault. We might also cite the fact that the movement's founder was a Black woman, Tarana Burke, whose contributions have largely been overshadowed by the more famous white women who joined in only after the initial--and therefore most dangerous--work had already taken place.

Burke's marginalization in the #MeToo movement is only one datapoint in a long line of Black women who have stood on the vanguard of feminist advocacy work, only to have their contributions subsumed by white feminists after the fact. This is a structural problem too. It's the result of several intersecting differentials of power--differentials of power that must be made visible and acknowledged before they can be challenged and changed.

To be clear, there are already a significant number of data scientists, designers, policymakers, educators, and journalists, among others, who share our goal of using data to challenge inequality and help change the world. These include the educators who are introducing data science students to real-world problems in health, economic development, the environment, and more, as part of the Data Science for Social Good initiative; the growing number of organizations like DataKind, Tactical Tech, and the Engine Room, that are working to strengthen the capacity of the civil sector to work with data; newsrooms like ProPublica and the Markup that use data to hold Big Tech accountable; and public information startups like MuckRock, which streamlines public records requests into reusable databases. Even a commercial design firm, Periscopic, has chosen the tagline, “Do Good With Data.” We agree that data can do good in the world. But we can do only do good with data if we acknowledge the inequalities that are embedded in the data practices that we ourselves rely upon. And this is where the bodies come back in.

In the rest of this chapter, we explain how it’s people and their bodies who are missing from our current data practices. Bodies are missing from the data we collect; bodies are extracted into corporate databases; and bodies are absent from the field of data science. Even more, it’s the bodies with the most power that are ever present, albeit invisibly, in the products of data science. Each of these is a problem, because without these bodies present in the field of data science, the power differentials currently embedded in the field will continue to spread. It’s by bringing back these bodies--into discussions about data collection, about the goals of our work, and about the decisions we make along the way--that a new approach to data science, one we call *data feminism*, begins to come into view.

Bodies uncounted, undercounted, silenced

One person already attuned to certain things missing from data science, and to the power differentials responsible for those gaps, is artist, designer, and educator Mimi Onuoha. Her project, *Missing Data Sets*, is a list of precisely that: descriptions of data sets that you would expect to already exist in the world, because they describe urgent social issues and unmet social needs, but in reality, do not. These include “People excluded from public housing because of criminal records,” “Mobility for older adults with physical disabilities or cognitive impairments,” and “Measurements for global web users that take into account shared devices and VPNs.” These data sets are missing for a number of reasons, Onuoha explains in her artist statement, many relating to issues of power. By compiling a list of the data that are missing from our “otherwise data-saturated” world, she states, we can “reveal our hidden social biases and indifferences.”

An Incomplete List of Missing Data Sets

This list will always be incomplete, and is designed to be illustrative rather than comprehensive. It also comes primarily from the perspective of the U.S, though the complete list of datasets features far more international examples.

- ~~Civilians killed in encounters with police or law enforcement agencies~~ [update: this is no longer a missing dataset]
- Sales and prices in the art world (and relationships between artists and gallerists)
- People excluded from public housing because of criminal records
- Trans people killed or injured in instances of hate crime (note: existing records are notably unreliable or incomplete)
- Poverty and employment statistics that include people who are behind bars
- Muslim mosques/communities surveilled by the FBI/CIA
- Mobility for older adults with physical disabilities or cognitive impairments
- LGBT older adults discriminated against in housing
- Undocumented immigrants currently incarcerated and/or underpaid
- Undocumented immigrants for whom prosecutorial discretion has been used to justify release or general punishment
- Measurements for global web users that take into account shared devices and VPNs
- Firm statistics on how often police arrest women for making false rape reports
- Master database that details if/which Americans are registered to vote in multiple states
- Total number of local and state police departments using stingray phone trackers (IMSI-catchers)
- How much Spotify pays each of its artists per play of song
-

Onuoha's list of missing datasets includes "People excluded from public housing because of criminal records," "Mobility for older adults with physical disabilities or cognitive impairments," and "Measurements for global web users that take into account shared devices and VPNs." By hosting the project on GitHub, Onuoha allows visitors to the site to suggest additional missing datasets that she might include. ¶ Credit: Mimi Onuoha ¶ Source: https://github.com/MimiOnuoha/missing-datasets ¶

The lack of data about women who die in childbirth makes Onuoha's point plain. In the absence of U.S. government-mandated action or federal funding ProPublica had to resort to crowdsourcing to find out the names of the estimated 700 to 900 U.S. women who died in childbirth in 2016. So far, they've identified only 134. Or, for another example: In 1998, youth of color in Roxbury, Boston, were sick and tired of inhaling polluted air. They led a march demanding clean air and better data collection, which led to the creation of the AirBeat community monitoring project. Just south of the U.S. border, in Mexico, a single anonymous woman is compiling the most comprehensive dataset on femicides – gender-related killings. The woman, who goes by the name "Princesa," has logged 3,920 cases of femicide since 2016. Her work provides the most up-to-date information on the subject for Mexican journalists and legislators--information that, in turn, has inspired those journalists to report on the subject, and has compelled those legislators to act.

Princesa has undertaken this important data collection effort because women's deaths are being neglected and going uncounted by the local, regional, and federal governments of Mexico. But it's not better anywhere else. *The Washington Post* and *The Guardian US* currently compile the most comprehensive national count of police killings of citizens in the United States, and not the U.S. federal government. But it's powerful institutions like the federal government that, more often than not, control the terms of data collection--for several reasons that Onuoha's *Missing Data Sets* points us towards. In the present moment, in which the most powerful form of evidence is data--a fact we may find troubling, but is increasingly true--the things that we do not or cannot collect data about are very often perceived to be things that do not exist at all.

Even when the data are collected, however, they still may not be disaggregated or analyzed in terms of the categories that make issues of inequality apparent. This is, in part, what is responsible for the lack of data on maternal mortality in the United States. While there is (as of 2003) a box to check on the official U.S. death certificate that indicates whether the person who died, if female, was pregnant at the time or within a year of death, it would require a researcher who was already interested in racial disparities in healthcare to combine those data with the data collected on race for the "three times more likely" stat that Serena Williams cited in her Facebook post to be revealed.

As feminist geographer Joni Seager states, "If data are not available on a topic, no informed policy will be formulated; if a topic is not evident in standardized databases, then, in a self-fulfilling cycle, it is assumed to be unimportant." Princesa's femicide map is an outlier, a case when a private citizen stood up and took action on behalf of the bodies that were going uncounted. ProPublica solicited stories and trawled Facebook groups and private crowdfunding sites in order to compile their list of the women who would otherwise go uncounted and unnamed. But this work is precarious in that it relies upon the will of individuals or the sustained attention of news organizations in order to take place. In the case of Princesa, this work is even more precarious in that it places herself and her family at risk of physical harm.

Sometimes, however, it's the subjects of data collection who can find themselves in harm's way. When power in the collection environment is not distributed equally, those who fear reprisal have strong reasons not to come forward. Collecting data on the locations of undocumented immigrants in the United States, for example, could on the one hand be used to direct additional resources to them; but on the other hand, it could send ICE officials to their doors. A similar paradox of exposure is evident among transgender people. Journalist Mona Chalabi has written about the challenges of collecting reliable data on the size of the transgender population in the U.S. Among other reasons, this is because transgender people are afraid to come forward for fear of violence or other harms. And so many choose to stay silent, leading to a set of statistics that does not accurately reflect the populations they seek to represent.

There is no universal solution to the problem of uncounted, undercounted, and silenced bodies. But that's precisely why it's so important to listen to, and take our cues from, the communities that we as data scientists, and data feminists, seek to support. Because these communities are disproportionately those of women, people of color, and other marginalized groups, it's also of crucial importance to recognize how data and power, far too often, easily and insidiously align. Bringing the bodies back into our discussions and decisions about what data gets collected, by whom, and why, is one crucial way in which data science can benefit from feminist thought. It's people and their bodies who can tell us what data will help improve lives, and what data will harm them.

2

There is a growing body of work dedicated the difficulties of uncounted and undercounted populations, and related phenomena. The emerging field of Critical Data Studies advocates for using frameworks from cartography and GIS which "have long been concerned with the nature of missing data", including theorizing their origins in power imbalances as well as determining ethical courses of action for mappers in diverse situations. Jonathan Gray, Danny Lämmerhirt, and Liliana Bounegru wrote a report, Changing What Counts, which includes case studies of citizen involvement in collecting data on drones, police killings, water supplies and pollution. Environmental health and justice represents an area where communities are out front collecting data when agencies refuse or neglect to do so. For example, Sara Wylie, co-founder of Public Lab, works with communities impacted by fracking to measure hydrogen sulfide using low-cost DIY sensors. The lack of data on women impacted by police violence in the U.S. led Kimberlé Crenshaw and the African American Policy Forum to develop the Black Women Police Violence database, designed to challenge the narrative that policy violence only affects males of color. Erin McElroy's work on community-collected eviction data in San Francisco, as part of the Anti-Eviction Mapping Project, demonstrates how data that originates in communities can be more complete and grounded than outside data collection efforts. Indigenous cartographers Margaret Pearce and Renee Pualani Louis describe cartographic techniques for recuperating indigenous perspectives and epistemologies (often absent or misrepresented) into GIS maps. And through methods like crowdsourcing or sensor journalism, the data journalism community is not just reporting with existing data, but increasingly undertaking projects that involve compiling their own databases in the absence of official data sources. That said, participatory data collection efforts have their own silences, as Heather Ford and Judy Wajcman show in their study of the 'missing women' of Wikipedia.

Bodies extracted for science, surveillance, and selling

Far too often, the problem is not that bodies go uncounted or undercounted, or that their existence or their interests go unacknowledged, but the reverse: that their information is enthusiastically scooped up for the narrow purposes of our data-collecting institutions. For example, in 2012, *The New York Times* published an explosive article by Charles Duhigg, "How Companies Learn Your Secrets," which soon became the stuff of legend in data and privacy circles. Duhigg describes how Andrew Pole, a data scientist working at Target, synthesized customers' purchasing histories with the timeline of those purchases in order to detect whether a customer might be pregnant. (Evidently, pregnancy is the second major life event, after leaving for college, that determines whether a casual shopper will become a customer for life). Pole's algorithm was so accurate that he could not only identify the pregnant customers, but also predict their due dates.

But then Target turned around and put this algorithm into action by sending discount coupons to pregnant customers. Win-win. Or so they thought, until a Minneapolis teenager's dad saw the coupons for maternity clothes that she was getting in the mail, and marched into his local Target to read the manager the riot act. Why was his daughter getting coupons for pregnant women when she was only a teen?!

It turned out that the young woman was, indeed, pregnant. Pole's algorithm informed Target before the teenager informed her father. Evidently, there are approximately twenty-five common products, including unscented lotion and large bags of cotton balls, that, when analyzed together, can predict whether or not a customer is pregnant, and if so, when they are due to give birth. But in the case of the Minneapolis teen, the win-win quickly became a lose-lose, as Target lost a potential customer and the pregnant teenager lost far worse: her privacy over information related to her own body and her health. In this way, Target's pregnancy prediction model helps to illustrate another reason why bodies must be brought back to the data science table: without the ability of individuals and communities to shape the terms of their own data collection, their bodies can be mined and their data can be extracted far too easily--and done so by powerful institutions who rarely have their best interests at heart.

At root, this is another question of power, along with a question of priorities and resources-- financial ones. Data collection and analysis can be prohibitively expensive. At Facebook's newest data center in New Mexico, the electrical cost alone is estimated at $31 million annually. Only corporations like Target, along with well-resourced governments and elite research universities, have the resources to collect, store, maintain, and analyze data at the highest levels. It's the flip side of the lack of data on maternal health outcomes. Put crudely, there is no profit to be made collecting data on the women who are dying, but there is significant profit in knowing whether women are pregnant.

Data has been called "the new oil" for, among other things, its untapped potential for profit and its value once it's processed and refined. But just as the original oil barons were able to use that profit to wield outsized power in the world--think of John D. Rockefeller, J. Paul Getty, or, more recently, the Koch brothers-- so too do the Targets of the world use their data capital to consolidate control over their customers. But it's not petroleum that's extracted in this case; it's data that's extracted from people and communities with minimal consent. This basic fact creates a profound asymmetry between who is collecting, storing, analyzing and visualizing data, and whose information is collected, stored, analyzed, and visualized. The values that drive this extraction of data represent the interests and priorities of the universities, governments, and corporations that are dominated by elite, white men. We name these values the three S's: science (universities), surveillance (governments) and selling (corporations).

3

In their widely cited paper Critical Questions for Big Data, danah boyd and Kate Crawford outlined the challenges of unequal access to big data, noting that the current configuration (in which corporations own and control massive stores of data about people) creates an imbalance of power in which there are "Big Data rich" and "Big Data poor." Media scholar Seeta Peña Gangadharan has detailed how contemporary data profiling disproportionately impacts poor, communities of color, migrants and indigenous groups. Social scientist Zeynep Tufecki warns that corporations have emerged as "power brokers" with outsized potential to influence politics and publics precisely because of their exclusive data ownership. Building on this, Mark Andrejevic has outlined a "big data divide" in which only elite institutions have abilities to capture, mine and utilize data whereas individuals do not, privileging "a form of knowledge available only to those with access to costly resources and technologies." Jeff Warren describes how this gives "data shepherds" (technologists) disproportionate power over knowledge production and discourse, circumscribing the kinds of questions that can be asked in a democracy. And in advancing the idea of "Black data" to refer to the intersection of informatics and Black queer life, Shaka McGlotten states, "How can citizens challenge state and corporate power when those powers demand we accede to total surveillance, while also criminalizing dissent?"

In the case of Target and the pregnant teen, the originating charge from the marketing department to Andrew Pole was: "If we wanted to figure out if a customer is pregnant, even if she didn't want us to know, can you do that?" But did the teenager have access to her purchasing data? No. Did she or her parents have a hand in formulating any of the questions that Target might wish to ask of its millions of records of consumer purchases? No. Did they even know that their family's purchasing data was being analyzed and recorded? No no no. They were not invited to the design table, even though it was one on which their personal data was put out on (corporate) display. Instead, it was Target--a company currently valued at $32 billion dollars--that determined what data to collect, and what questions to ask of it.

The harms inflicted by this asymmetry don't only have to do with personal exposure and embarrassment, but also with the systematic monitoring, control, and punishment of the people and groups who hold less power in society. For example, Paola Villareal's data analysis for the ACLU reveals clear racial disparities in the City of Boston's approach to policing marijuana-related offenses. (Additional analyses have found this phenomenon to be true in cities across the United States). In *Automating Inequality*, Virginia Eubanks provides another example of how the asymmetrical relationship between data-collecting institutions and the people about which they collect data plays out. The Allegheny County Office of Children, Youth, and Families, in Pennsylvania, employs an algorithmic model to predict the risk of child abuse. Additional methods of detecting child abuse would seem to be a good thing. But the problem with this particular model, as with most predictive algorithms in use in the world today, is that it has been designed unreflexively. In this case, the problem is rooted in the fact that it takes into account every single data source that it can get. For wealthier parents, who can more easily access private health care and mental health services, there is simply not that much data. But for poor parents, who primarily access public resources, the model scoops up records from child welfare services, drug and alcohol treatment programs, mental health services, jail records, Medicaid histories, and so on. Because there is far more data about poor parents, they are oversampled in the model, and disproportionately targeted for intervention. The model "confuse[s] parenting while poor with poor parenting," Eubanks explains-- with the most profound of results.

Ensuring that bodies are not simply viewed as a resource, like oil, that can be "extracted" and "refined," is another way that data feminism can intervene in our current data practices. Like the process of data collection, this process of extracting bodies is one that disproportionately impacts women, people of color, low-income people, and others who are more often subject to power rather than in possession of it. And it's another place where bringing the bodies back into discussions about data collection, and its consequences, can begin to challenge and transform the unequal systems that we presently face.

Bodies absent from data work

One place where these conversations need to be happening is in the field of data science itself. It's no surprise to observe that women and people of color are underrepresented in data science, just as they are in STEM fields as a whole. The surprising thing is that the problem is getting worse. According to a research report published by the American Association of University Women in 2015, women comprised 35% of computing and mathematical occupations in 1990, but this percentage dropped to 26% in 2013.

4

For comparison, this is the same percentage of female information science graduates in 1974. And in subfields like machine learning, the proportion of women is even less. As per the points made in this chapter, even knowing the exact extent of the disparity is challenging. According to a 2014 Mother Jones report about diversity in Silicon Valley, tech firms convinced the U.S. Labor Department to treat their demographics as a trade secret, and didn't divulge any data until after they were sued by Mike Swift of the *San Jose Mercury News*. There are analyses that have obtained the data in other ways. For example, a gender analysis by data scientists at

LinkedIn has shown that tech teams at tech companies have far *less* gender parity than tech teams in other industries including healthcare, education, and government.

They are being pushed out as "data analysts" have become rebranded as "data scientists," in order to make room for more highly valued and more highly compensated men.

5

This phenomenon, while new to data science, is unfortunately as old as time. Scholars such as Marie Hicks and Nathan Ensmenger have shown how the push to professionalize computer science resulted in the pushing out of the women who had previously performed those same roles. Historians of medicine often point to the history of obstetrics, in which female midwives were replaced by male obstetricians after the advent of formal medical schools. The same phenomenon can be found in the kitchen, with women performing most home cooking, unpaid altogether, while men attend culinary school to become celebrity chefs.

We identify this later in the book as what we call a "privilege hazard," one in which discrimination becomes hard-coded into so-called "intelligent systems," because the people doing the coding are the most privileged-- and therefore the least well-equipped-- to acknowledge and account for inequity.

6

Social scientist Kate Crawford has advanced the idea that the biggest threat from artificial intelligence systems is not that they will become smarter than humans, but rather that they will hardcode sexism, racism and discrimination into the digital infrastructure of our societies. This is evident not only in data products and systems themselves but also in the divisions of labor in the data economy. The book Ghost Work by anthropologist Mary Gray and computer scientist Siddharth Suri details the existence of a "global underclass" performing work like content moderation, transcription, and captioning. While Silicon Valley tech workers remain steadily young, white and male, these "ghost workers" are often older, often female and minority, and always precarious.

This privilege hazard is a risk that can rear its head in harmful ways. For example, in 2016, MIT Media Lab graduate student Joy Buolamwini, founder of the Algorithmic Justice League, was experimenting with software libraries for the Aspire Mirror project. This project used computer vision software to overlay inspirational images (like a favored animal or an admired celebrity) onto a reflection of the user's face. She would open up her computer and run some code that she'd written, built on a free JavaScript library that used her computer's built-in camera to detect the contours of her face. Buolamwini's code was bug-free, but she couldn't get the software to work for a more basic reason: it had a really hard time detecting her face in front of the camera. Buolamwini has dark skin. While her computer's camera picked up her lighter-skinned colleague's face immediately, it took much longer for the camera to pick up Buolamwini's face, when it did at all. Even then, sometimes, her nose was identified as her mouth. What was going on?

Joy Buolamwini had to resort to "white face" to get a computer vision algorithm to detect her face. Many facial detection algorithms have only been trained on pale and male faces. ¶ Credit: Joy Buolamwini ¶ Source: https://medium.com/mit-media-lab/the-algorithmic-justice-league-3cc4131c5148 ¶ Permissions: Pending

What was going on was this: facial analysis technology, which uses machine learning approaches, learns how to detect faces based on existing collections of data that are used to train, validate, and test models that are then deployed. These datasets are constructed in advance, in order to present any particular learning algorithm with a representative sample of the kinds of things it might encounter in the real world. But problems arise very quickly when the biases that already exist in the world are replicated in these datasets. Upon digging into the benchmarking data for facial analysis algorithms, Buolamwini learned that they consisted of 78% male faces and 84% pale faces, sharply at odds with a global population that is majority female and majority non-pale.

7

Specifically, the breakdown for the Labeled Faces in the Wild (LFW) dataset was 77.5% male faces and 83.5% white faces. And Buolamwini and Timnit Gebru showed that the breakdown for the IARPA Janus Benchmark A (IJB-A) dataset published by the US government was 75% male and 80% pale faces (as determined by the Fitzpatrick skin type). But Buolamwini makes the additional point that population parity in the test data is not always the answer, because small populations like Native Americans might not have enough test cases to determine whether the model was working.

How could such an oversight have happened? Easily, when most engineering teams have 1) few women or people of color; and 2) no training to think about #1 as a problem.

Oversights like this happen more often than you might think, and with a wide range of consequences. Consider a craze that (briefly) swept the internet in Spring 2018. In order to promote awareness of its growing number of digitized museum collections, Google released a new feature for its Arts and Culture app. You could take a selfie, upload the image, and the app would find the face from among its millions of digitized artworks that looked the most like you. All over Facebook, Twitter, and Instagram, people were posting side-by-side shots of themselves and--for instance, the *Mona Lisa*, *American Gothic,* or a Vermeer.

Well, white people were. Because most of the museums with collections that Google had helped to digitize came from the U.S. and Europe, most featured artworks from the Western canon. And because most artworks from the Western canon feature white people, the white users of the Arts and Culture app found really good matches for their faces. But some Asian users of the app, for example, found themselves matched with one of only the handful of portraits of Asian people included in those collections.

On Twitter, the response to this inadequacy was tellingly resigned. One user, @pitchaya, whose Tweet was quoted in a *digg.com* article on the subject, tweeted sarcastically: “If you do that whole Google Arts & Culture app portrait comparison as an Asian male, it gives you one of 5-6 portraits that hardly resembles you but, hey, looks Asian enough.” Another user, @rgan0, also quoted in the piece, called out Google directly: “The Google Arts and Culture app thinks I look like a “Beautiful [Japanese] Woman”! :p get more Asian faces in your art database, Google.”

And if the disparities of representation in Western art museums weren’t enough of a problem, some Art and Culture App users worried about something more insidious taking place. For app users to upload their images for analysis, they had to agree to allow Google to access those images. Were their images also being stored for future internal research? Was Google secretly using crowdsourcing to improve its training data for its own facial recognition software, or for the NSA? A short-lived internet uproar ensued, ending only when Google updated the user agreement to say: “Google won’t use data from your photo for any other purpose and will only store your photo for the time it takes to search for matches.”

But what if they had been? The art selfie conspiracy theorists weren’t actually too far from reality, given that earlier that year, Amazon had briefly been contracted by the Orlando Police Department to use its own proprietary facial recognition software, trained on its own proprietary data, to help the police automatically identify suspects in real time. How representative was Amazon's training, benchmarking, or validation data? Was it more or less representative than the data that Buolamwini explored in her research? There was no way to know. And while a best match of 44% between Asian Art and Culture App users and Terashima Shimei's *Beautiful Woman* (which is the painting @rgon0 matched with) might earn RTs of solidarity on Twitter, a best match of 44% between a suspected criminal and a random person identified through traffic camera footage--the image source for the Amazon project--could send an innocent person to jail.

Who any particular system is designed for, and who that system is designed by, are both issues that matter deeply. They matter because the biases they encode, and often unintentionally amplify, remain unseen and unaddressed--that is, until someone like Buolamwini literally has to face them. What's more, without women and people of color more involved in the coding and design process, the new research questions that might yield groundbreaking results don't even get asked--because they're not around to ask them. As the example of facial analysis technology, or the Google Arts and Culture app help to show, there is a much higher likelihood that biases will be designed into data systems if the bodies of the system's designers themselves only represent the dominant group.

Bodies invisible: The view from nowhere is always a view from somewhere

So far, we've shown how bringing the bodies back into data science can help expose the inequities in the scope and contents of our data sets, as in the example of the hundreds of unnamed U.S. women who die in childbirth each year. We've also shown how bringing back the bodies can help avoid their data being mined without their consent, as in the example of the Minneapolis teenager who Target identified as pregnant. And we've also shown how bringing bodies that are more representative of the population into the field of data science can help avert the increasing number of racist, sexist data products that are inadvertently released into the world, as in the example of the Google Arts and Culture app, or of the facial recognition software that is the focus of Joy Buolamwini's research. (We'll have more to say about some of the worst applications of computer vision, like state surveillance, in the chapters to come).

But there are other bodies that need to be brought back into the field of data science not because they're not yet represented, but because of the exact opposite reason: they are overrepresented in the field. They are so overrepresented that their identities and their actions are simply assumed to be the default. An example that Yanni Loukissas includes in his book, *All Data are Local*, makes this point crystal clear: Marya McQuirter, a former historian at the Smithsonian Institution's National Museum of African American History and Culture, recalls searching the Smithsonian's internal catalog for the terms "black" and "white." Searching the millions of catalog entries for “black” yielded a rich array of objects related to Black people, Black culture, and Black history in the US : the civil rights movement, the jazz era, the history of enslavement, and so on. But searching for “white” yielded only white-colored visual art. Almost nothing showed up relating to the history of white people in the United States.

McQuirter, who is Black, knew the reason why: in the United States, it’s white people and their bodies who occupy the “default” position. Their existence seems so normal that they go unremarked upon. They need not be categorized, because-- it is, again, assumed-- most people are like them. This is how the perspective of only one group of bodies--the most dominant and powerful group --becomes invisibly embedded in a larger system, whether it’s a system of classification, as in the case of McQuirter’s catalog search; a system of surveillance, as in the case of Amazon and the Orlando police; or a system of knowledge, as reflected in a data visualization, as we’ll now explain--

Whose perspective are we seeing when we see a visualization like this one of global shipping routes?

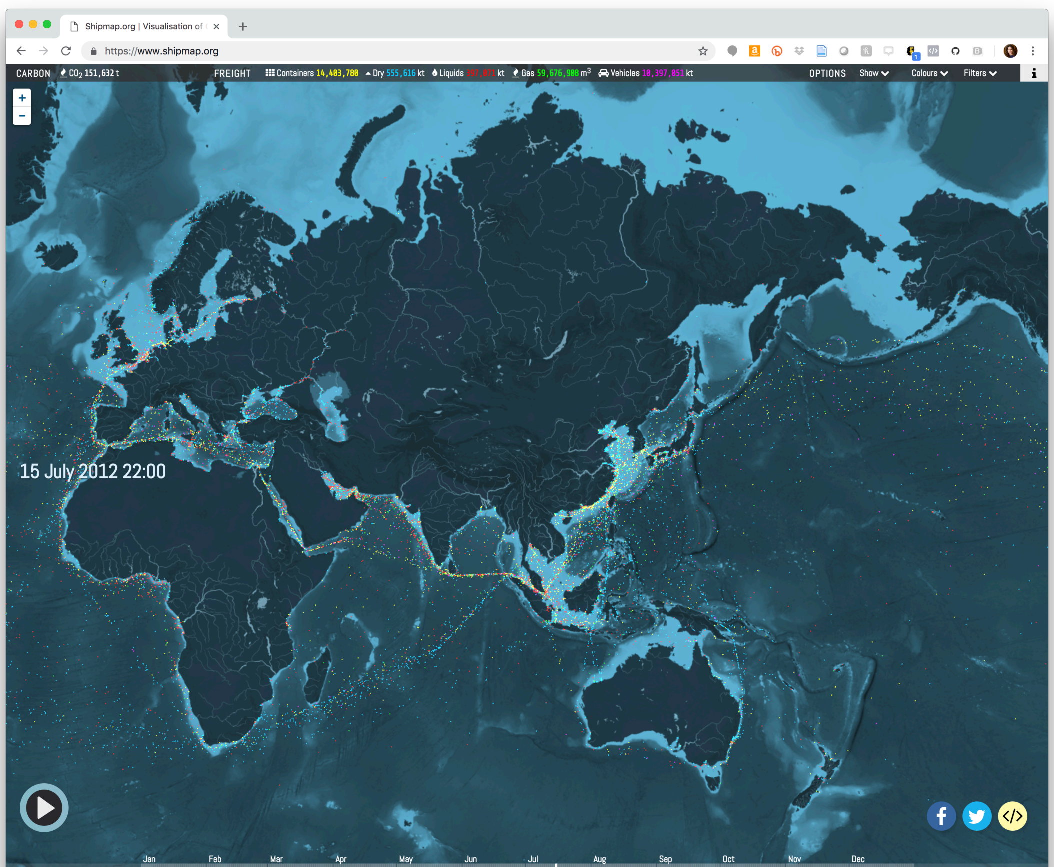

Time-based visualization of global shipping routes designed by Kiln based on data from the UCL Energy Institute. ¶ Credit: Website created by Duncan Clark & Robin

Houston from Kiln. Data compiled by Julia Schaumeier & Tristan Smith from the UCL EI. The website also includes a soundtrack: Bach's Goldberg Variations played by Kimiko Ishizaka. ¶ Source: https://www.shipmap.org/

We are not seeing any particular person's perspective when we look at this map (unless you are an astronaut on the space station and you have weird blue glasses on that make all the continents blue). In terms of visualization design, this is for good reason - it is precisely this impossible, totalizing view which makes any particular visualization so dazzling and seductive, so rhetorically powerful, and so persuasive.

8

Sociologist Helen Kennedy and her colleagues have shown how visual conventions such as two-dimensional layouts, and geometric shapes, contribute to the pervasive view of data visualization as a neutral and scientific method of display.

This image appears to show us the "big picture" of the entire world. Because we do not see the designers of this image, nor can we detect any visual indicators of human involvement, the image appears truthful, accurate, and free of bias.

This is what feminist philosopher Donna Haraway describes as "the god trick." By the "god" part, Haraway refers to how data is often presented as though it inhabits an omniscient, godlike perspective. But the "trick" is that the bodies who helped to create the visualization – whether through providing the underlying data, collecting it, processing it, or designing the image that you see–have themselves been rendered invisible. There are no bodies in the image anymore.

Haraway terms this "the view from nowhere." But the view from nowhere is always a view from somewhere: the view from the default. Sometimes this view comes into focus when considering what isn't revealed, as in the case of McQuirter's search query. But when we do not remind ourselves to ask what we are *not* seeing, and about *who* we are not seeing--well, that is the most serious body issue of all. It's serious because all images and interactions, the data they are based on, and the knowledge they produce, comes from bodies. As a result, this knowledge is necessarily incomplete. It's also necessarily culturally, politically, and historically circumscribed. Pretending otherwise entails a belief in what sociologist Ruha Benjamin, in *Race After Technology: Abolitionist Tools for the New Jim Code*, describes as the "imagined objectivity of data and technology," because it's not objectivity at all.

To be clear: this does *not* mean that there is no value in data or technology. What this means for data science is this: if we truly care about objectivity in our work, we must pay close attention to whose perspective is assumed to be the default. Almost always, this perspective is the one of elite white men, since they occupy the most privileged position in the field, as they do in our society overall. Because they occupy this position, they rarely find their dominance challenged, their neutrality called into question, or their perspectives open to debate. Their privilege renders their bodies invisible– in datasets, in algorithms, and in visualizations, as in their everyday lives.

Ever heard of the phrase, "History is written by the victors"? It's the same sort of idea. Both in the writing of history and in our work with data, we can learn so much more-- and we can get closer to some sort of truth-- if we bring together as many bodies and perspectives as we can. And when it comes to bringing these bodies back into data science, feminism becomes increasingly instructive, as the rest of the chapters in this book explain.

In *On Rational, Scientific, Objective Viewpoints from Mythical, Imaginary, Impossible Standpoints*, we build on Haraway's notion of the god trick, exploring some reasons why emotion has been kept out of data science as a field, and what we think emotion can, in fact, contribute. We talk about emotional data, among data of many other forms, in *What Gets Counted Counts*--a chapter that emphasizes the importance of thinking through each and every one of the choices we make when collecting and classifying data. The next chapter, *Unicorns, Janitors, Ninjas, Wizards, and Rock Stars*, challenges the assumption that data scientists are lone rangers who wrangle meaning from mess. Instead, we show how working with communities and embracing multiple perspectives can lead to a more detailed picture of the problem at hand. This argument is continued in *The Numbers Don't Speak for Themselves*, in which we show how much of today's work involving "Big Data" prioritizes size over context. In contrast, feminist projects connect data back to their sources, pointing out the biases and power differentials in their collection environments that may be obscuring their meaning. We turn to the contexts and communities that ensure that the work of data science can take place in *Show Your Work*, a chapter that centers on issues of labor. In *The Power Chapter*, it's, well, power, privilege, and structural inequality that we take up and explore. *Teach Data Like an Intersectional Feminist* provides a series of examples of how to implement the lessons of the previous chapters in classrooms, workshops, and offices, so that we can train the next generations of data feminists. And in *Now Let's Multiply*, we speculate about other approaches that might enrich a conversation about data science, its uses, and its limits.

There is growing discussion about the uses and limits of data science, especially when it comes to questions of ethics and values. But so far, feminist thinking hasn't directed the conversation as it might. As a starting point, let's take the language that is increasingly employed to discuss questions of ethics in data and the algorithms that they support, such as the computer vision and predictive policing algorithms we've described just above. The emerging best practices in the field of data ethics involve orienting algorithmic work around concepts like "bias," and values like "fairness, accountability, and transparency." This is a promising development, especially as conversations about data and ethics enter the mainstream, and funding mechanisms for research on the topic proliferate. But there is an additional opportunity to reframe the discussion before it gathers too much speed, so that its orienting concepts do not inadvertently perpetuate an unjust status quo.

Consider this chart, which uses Benjamin's prompt to reconsider the "imagined objectivity of data and technology" in order to develop an alternative set of orienting concepts for the field. These concepts have legacies in intersectional feminist activism, collective organizing, and critical thought, and they are unabashedly explicit in how they work towards justice:

Concepts Which Uphold "Imagined Objectivity"	**Intersectional Feminist Concepts Which Strengthen Real Objectivity**
Because they locate the source of the problem in individuals or technical systems	*Because they acknowledge structural power differentials and work towards dismantling them*
Ethics	Justice
Bias	Oppression
Fairness	Equity
Accountability	Co-liberation
Transparency	Reflexivity
Understanding algorithms	Understanding history, culture, and context

The concept of "bias," for example, locates the source of inequity in the behavior of individuals (i.e. a prejudiced person) or in the outcomes of a technical system (i.e. a system that favors white people or men). Under this conceptual model, a technical goal might be to create an "unbiased" system. First we would design a system, use data to tune its parameters and then we would test for any biases that result. We could even define what might be more "fair," and then we could optimize for that.

But this entire approach is flawed, like the imagined objectivity that shaped it. Just as Benjamin cautions against imagining that data and technology are objective, we must caution ourselves against locating the problems associated with "biased" data and algorithms in technical systems alone. This is a danger that computer scientists have noted in relation to high-stakes domains like criminal justice, where hundreds of years of history, politics, and economics, not to mention the complexities of contemporary culture, are distilled into black-boxed algorithms that determine the course of people's lives. In this context, computer scientist Ben Green warns about the narrowness of computationally conceived fairness, writing that "computer scientists who support criminal justice reform ought to proceed thoughtfully, ensuring that their efforts are driven by clear alignment with the goals of justice rather than a zeitgeist of technological solutionism." And in keynoting the Data Justice Conference in 2018, design theorist Sasha Costanza-Chock challenged the audience to expand their concept of ethics to justice, in particular *restorative justice* which recognizes and accounts for the harms of the past. We do not all arrive in the present moment with equal power and privilege. When "fairness" is a value that does not acknowledge context or history, it fails to acknowledge the systematic nature of the "unfairness" perpetrated by certain groups on other groups for centuries.

Does this make fairness political? Emphatically yes, because all systems are political. In fact, the appeal to avoid politics is a very familiar move for those in power to continue to uphold the status quo. The ability to do so is also a privilege, one held only by those whose existence does not challenge that same status quo. Rather than designing algorithms that are "color blind," Costanza-Chock says, we should be designing algorithms that are *just*. This means shifting from ahistorical notions of fairness to a model of *equity*. This model would take time, history, and differential power into account. Researcher Seeta Peña Gangadharan, co-lead of the Our Data Bodies project, states, "The question is not 'How do we make automated systems fairer?' but rather to think about how we got here. How might we recover that ability to collectively self determine?"

This is why *bias* (in individuals, in data sets, or in algorithms) is not a strong enough concept in which to anchor ideas about equity and justice. In writing about the creation of New York's Welfare Management System in the early 1970s, for example, Virginia Eubanks describes: "These early big data systems were built on a specific understanding of what constitutes discrimination: personal bias." The solution at the time was to remove the humans from the loop, and it remains so today: without potentially bad--in this case, racist-- apples, there would be less discrimination. But this line of thinking illustrates what Robin DiAngelo would call the "'new' racism": the belief that racism is due to individual bad actors, rather than structures or systems. In relation to welfare management, this often means replacing the women of color social workers, who have empathy and flexibility and listening skills, with an automated system that applies a set of rigid criteria, no matter what the circumstances.

Bias is not a problem that can be fixed after the fact. Instead, we must look to understand and design systems that address *oppression* at the structural level. Oppression, as defined by the comic artist Robot Hugs, is what happens "when prejudice and discrimination is supported and encouraged by the world around you. It is when you are harmed or not helped by government, community or society at large because of your identity," they explain. And while the research and energy emerging around algorithmic *accountability* is promising, why should we settle for retroactive audits of potentially flawed systems if we could design for *co-liberation* from the start? Here co-liberation doesn't mean "free the data," but rather "free the people." And the people in question are not only those with less power, but also those with relative privilege (like data scientists, designers, researchers, educators;

like ourselves) who play a role in upholding oppressive systems. Poet and community organizer Tawana Petty defines what co-liberation means in relation to anti-racism in the U.S.: "We need whites to firmly believe that their liberation, their humanity is also dependent upon the destruction of racism and the dismantling of white supremacy." The same goes for gender – men are often not even thought to have a gender, let alone prompted to think about how unequal gender relations seep into our institutions and artifacts and harm all of us. In these situations, it is not enough to do audits after-the-fact. We should be able to dream of data-driven systems that position co-liberation as their primary design goal.

Designing data sets and data systems that dismantle oppression and work towards justice, equity, and co-liberation requires new tools in our collective toolbox. We have some good starting points – building more *understandable algorithms* is a laudable, worthy research goal. And yet, what we need to explain and account for are not only the inner workings of machine learning, but also the *history, culture, and context* that lead to discriminatory outputs in the first place. Did you know, for example, that the concept of homophily which provides the rationale for most contemporary network clustering algorithms in fact derives from 1950s-era models of housing segregation? (If not, we recommend you read Wendy Chun). Or, for another example, did you know that the “Lena” image used to test most image processing algorithms is the centerfold from the November 1972 issue of *Playboy*, cropped demurely at the shoulders? (If not, Jacob Gaboury is the one to consult on the subject). These are not merely bits of trivia to be pulled out to impress dinner party guests. On the contrary, they have very real implications for the design of algorithms, and for their use.

How might we design a network clustering algorithm that does not perpetuate segregation, but actively strives to bring communities together? (This is a question that Chun is pursuing in her current research). How might we ensure that the selection of test data isn’t ever relegated to happenstance? (This is how the “Lena” image, which encoded sexism into the field of image processing, is explained away). The first step requires *transparency* in our methods as well as the *reflexivity* to understand how our own identities, our communities, and our domains of expertise are part of the problem. But they can also be part of the solution.

When we start to ask questions like: "Whose bodies are benefiting from data science?" "Whose bodies are harmed?" "How can we use data science to design for a more just and equitable future?" and "By whose values will we re-make the world?" we are drawing from *data feminism*. It’s data feminism that we describe in the rest of this book. It’s what can help us understand how power and privilege operate in the present moment, and how they might be rebalanced in the future.