EGM/STI/EP.8
October 2022
ENGLISH ONLY

UN Women
Expert Group Meeting
'Innovation and technological change, and education in the digital age for achieving gender equality and the empowerment of all women and girls'
10-13 October 2022

A Social Justice Framework for Leveraging Data Science to Advance Gender Equity

Expert paper prepared by:

Lauren Klein
Emory University
&
Brandeis Marshall
DataedX Group[1]

[1] The views expressed in this paper are those of the author and do not necessarily represent those of the United Nations.

A Social Justice Framework for Leveraging Data Science to Advance Gender Equity

Lauren Klein, Ph.D.
Winship Distinguished Research Professor
Departments of English and Quantitative Theory & Methods
Emory University

Brandeis Marshall, Ph.D.
Founder and CEO
DataedX Group

Introduction

In the world today, data is a form of power. Data has been used in scientific research to advance medical discovery, in crisis response to direct resources to communities in need, and in analysis and evaluation settings to provide evidence of programmatic success (or failure). But far more often, data has been used to discriminate, police, and surveil. Perhaps most famously, Amazon was required to scrap its automated first-round resume-screening system when it was discovered to have down-graded the resumes of women applicants (Goodman 2018). In the United States, fears about the restrictions recently placed on abortion access–the result of the rollback of Roe v. Wade–have led to widespread calls for women and girls to delete all menstruation data from their health-tracking apps (Garamvolgyi 2022). And advocacy by the Algorithmic Justice League has called attention to the global risks of facial recognition software when employed by both corporate and state actors, particularly risks to Black and brown women and to nonbinary people, for whom the software performs particularly poorly while simultaneously being subject to disproportionate monitoring and surveillance (2022). The interrelated and intersectional harms brought about by these data-driven systems can be traced to the fact that the power of data is currently wielded unequally. More specifically, it is corporations, governments, and other well-resourced institutions–institutions that represent the values and views of those in positions of power–who have the ability to design and deploy these data systems, while those whose lives and livelihoods are most dependent on the output of these systems remain absent from any conversations about their potential uses or potential harms.

We write as a multiracial pair of cisgender women data scientists, living in the US South, in the Global North. Our formal training spans the digital humanities (Klein) and computer science (Marshall). Both of us employ data science techniques in our research and teaching. We have also each written books: *Data Feminism* (Klein, coauthored with Catherine D'Ignazio, MIT Press, 2020), and *Data Conscience* (Marshall, Wiley, 2022). These books draw from the long history of intersectional feminist organizing and critical thought, and broader social movements, in order to propose a pair of frameworks for rebalancing and redistributed the unequal power relationships that currently structure the field of data science. These frameworks can be applied to the entire lifecycle of data science research, from the communities involved in the initial phases of research ideation; to the categories of data collection and the context that surrounds any particular dataset; to questions of transparency and accountability that surround any particular data analysis; to the governance required to ensure that any output of a data-driven project are appropriately put to use. In short, we believe that it is possible to leverage data science to advance gender equity, but only if that data science and the research and actions that it enables are ethically and intentionally envisioned from the start.

In what follows, we draw from *Data Feminism* and *Data Conscience* in order to provide a short account of the major problems that exist with standard approaches to data science, particularly as they relate to data science involving women and girls; a framework–equal parts conceptual and practical–that can help structure interventions into these problems; and a set of technical and policy recommendations for governments and other agencies seeking to enable these interventions. Throughout, we interweave examples of current data science projects that demonstrate how this framework can be operationalized in the world.

Research Foundations

Our research in *Data Feminism* and *Data Conscience* builds on the foundational work of critical data studies scholars such as Safiya Noble, whose *Algorithms of Oppression* (NYU Press, 2018) shone a light on the racism and sexism encoded in data systems as pervasive as Google search results; Meredith Broussard, whose *Artificial Unintelligence* (MIT Press, 2018) punctured the myth that computers could take over for human decision-making, particularly when accounting for underrepresented groups; and Mary Gray and Siddharth Suri's *Ghost Work* (HarperCollins, 2019) which exposed the global underclass of tech workers who keep our data-driven systems afloat. We also place our work in dialogue with more recent contributions to this space, including Paola Ricaurte's "Data Epistemologies, the Coloniality of Power, and Resistance" (2019), which calls attention to the colonial pathways reinscribed by technical infrastructure; Ruha Benjamin's *Race After Technology* (Polity, 2020), which names phenomenon by which data systems accelerate racism and discrimination as "the New Jim Code"; and Wendy Chun's *Discriminating Data* (MIT Press, 2021), which traces the ways by which contemporary data analysis methods amplify racism, misogyny, and dis/misinformation to their historical roots.

We also draw from technical research, including work by Joy Buolamwini and Timnit Gebru on the disparate performance of gender recognition software on pale vs. dark skin tones, as well as their attention to the potential for harm of gender and facial recognition software more generally (2018); on Os Keyes's further exploration of the harms of automated gender recognition for trans and nonbinary people (2018); on Abeba Birhane et al.'s work on the values more broadly encoded in machine learning research and the values that are not (2021) and on the limits of participatory models for data science and AI development (2022); Birhane's additional work on the "algorithmic colonization" of the African continent (2020); Kotaro Hara et al. (2018) and Carlos Toxtli et al.'s (2021) investigations into the extractive labor associated with cultural data work; and Inioluwa Deborah Raji et al. on the limits of benchmarks (2021) and audits (2019, 2020, 2022) for keeping data systems in check. Taken together, these papers affirm the inadequacy of existing approaches to ethical data science, particularly with respect to gender equity.

But in exposing the gaps in these approaches–both in terms of who is included (or not) in envisioning this work and in terms of how the work is itself envisioned–this research points the way to alternative frameworks that might better leverage data science to advance gender equity and gender justice more broadly. In what follows, we elaborate one such framework that derives from our own research.

A Social Justice Framework for Leveraging Data Science to Advance Gender Equity
Data Feminism advances a way of thinking about datasets, data systems, and data science that is informed by the rich history of feminist activism and feminist critical thought. In *Data Conscience*, Marshall calls out the lackadaisical attitude in implementing transparency, accountability, and governance in data systems and imposes a blunt call to action. Together, our books offer conceptual and practical guidelines that, when unified into the set of guidelines elaborated below, can be leveraged to advance gender equity with data science.

1. Recognize how gender equity is about more than women and more than gender

The goal of achieving equality for people of all genders requires a commitment to examining the root causes of the inequalities that women and girls, and other gender minorities, face today. With this in mind, it is important to underscore that gender equality is about more than women; it takes more than one gender to have gender inequality and more than one gender to work toward justice. Furthermore, achieving gender equality involves attending to more than gender; intersectional feminists like Kimberlé Crenshaw, bell hooks, and the Combahee River Collective, have taught us how race, class, sexuality, ability, age, religion, geography, and more are factors that work together to influence each person's experiences and opportunities in the world. Intersectional feminism also teaches us that these experiences and opportunities (or the lack of opportunities, as the case may be) are the result of larger structural forces of power, which must be challenged and changed. And because the power of data is currently wielded unequally, it must be challenged and changed as well.

To this end, we choose to employ the term "equity" over "equality" in this paper. The difference between equity and equality is that equality is measured from a starting point in the present, with resources and/or punishments measured out according to what is happening now. But this approach of *equal* resource allocation means that those who are ahead in the present will go even further, achieve even more, and stay on top, whereas those who start out behind will find it difficult to catch up. Working toward a world in which everyone is treated equally means taking present power differentials into account and distributing (or redistributing) resources accordingly. This is an *equity* approach. Equity is much harder to model through data-scientific methods than equality, as it needs to take time, history, and differential power into account, but it is not impossible.

This difficulty also underscores the point that a broad formulation of "ethics" (of data collection, data science, or in people) is not a strong enough concept in which to anchor ideas about gender equity. We must instead adopt an justice-oriented approach, one which looks to understand and design systems that intervene in the root cause of gender inequality: unequal power and the structural forces that produce and maintain it. For broader evidence of this shift, we might look to the expansion of DEI (diversity, equity and inclusion) efforts to JEDI (justice, equity, diversity and inclusion). For example, the J.E.D.I. Collaborative works to dismantle barriers within the natural products industry through a justice-oriented approach. They are also developing resources and toolkits to help mitigate microaggressions, bias, and marginalization tactics. This is not to say that a shift in name is always matched by a shift in action, as Sasha Costanza-Chock has observed (2020). Rather, we include this discussion to underscore the importance of continually evaluating the concepts that structure our work so that its impact is as substantive and transformative as it can be–and is in fact required if we are to ever hope to address the full extent of the injustices that we face today.

2. Identify how unequal power impacts research data and research questions

We see the unequal balance of power with respect to data science consistently play out in decisions about what research data to collect, and what research to undertake on the basis of that data. The Nigerian-American artist Mimi Ọnụọha has a project, *The Library of Missing Datasets*, which calls attention to how the interests of the powerful–corporations, governments, and the like–are what determine which issues are addressed via data science and which issues or not. “That which we ignore reveals more than what we give our attention to,” Ọnụọha explains. “It’s in these things that we find cultural and colloquial hints of what is deemed important. Spots that we've left blank reveal our hidden social biases and indifferences” (2018).

But the issue of missing data is not only the subject of artistic inquiry; it can also be very real. The example of maternal mortality in the United States underscores this point. On the one hand, corporations like Target have famously directed internal data-scientific research efforts towards predicting whether a customer is pregnant or not, with a goal of profiting off of their increased pregnancy and baby purchases (Hill 2012). While on the other hand, hospitals have chosen not to maintain a national tracking system for data about whether those same pregnant people are actually surviving childbirth–particularly Black and brown women, for whom maternal mortality remains an outsized risk (SisterSong et al. 2014). This stands in contrast to issues like heart attacks and hip replacements, for which such systems have long been in place (Young 2021). The reason for this lack of data points back to the structural imbalance of power with respect to data collection. In the United States, women remain underrepresented in the field of data science, as they do at the highest levels of corporate and medical leadership–and women of color even more so (Mangurian et al. 2018). While women, and particularly Black and brown women, have consistently reported their own experiences of poor maternal healthcare; myriad families have experienced the tragic loss of loved ones; and grassroots advocacy groups have long organized around the issue of maternal mortality, those who directly control the mechanisms of data collection have not identified the issue of maternal mortality as an issue worth of prioritizing. “Our maternal data is embarrassing,” stated Stacie Geller, a professor of 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 (Field and Sexton 2017). We might edit his statement to add that it’s a measure of *who* we value in health, too.

In the opening chapter of *Data Feminism*, D'Ignazio and Klein argue for the importance of asking three "who questions" (Muller 2011) about power in data science: *Data science for whom? Data science by whom?* and *Data science with whose interests and goals in mind?* When asked about the datasets, models, and other data projects related to questions of gender justice (or the lack thereof), these questions help to expose how the terms of data collection–and, in turn, the research questions that those data might help to answer–are not aligned with those who stand to benefit the most from that research.

3. Ensure meaningful and inclusive categories of data collection

The power imbalances that determine which data are collected and which data are not carry over into the categories of data collection as well. In the United States, women were wholly excluded from medical trials until 1993–the result of gender bias as well as concerns over fertility and reproduction as well as that women's fluctuating hormonal levels would confound any particular trial's data. This has resulted in generations of medical research that presumes that there are no meaningful sex differences in terms of prevalence of illness, response to treatment, severity of outcomes, etc. (Liu and Major 2017). Even today, most trials still only consider cisgender women, tracking gender data in binary form (Hodshire 2022). But gender, as we know, includes more than two genders, and for true gender equality to be achieved, we must ensure categories of data collection that account for gender beyond the binary.

The lessons associated with accounting for gender beyond the binary, and with including gender among the categories of data collection, extend far beyond medical research. There is a gender component to most social issues, and yet gender data is very often not considered as core to the issue being explored. The author Caroline Criado-Perez has memorably documented how, in Sweden, choices in the routes taken by snow-plows result in disparate outcomes for women (2019). The explanation is as follows: when roads are cleared before sidewalks, benefits accrue to people who drive vs. people who walk; and because the majority of people who drive are men communing to work, while the majority of those who walk are women–often traveling with small children–and because a significant number of accidents among pedestrians occur in the snow, the decision to prioritize clearing roads vs. sidewalks and bike paths resulted in women experiencing 69% of all pedestrian injuries, even as they represent only half of the population. Examples like this demonstrate how structural inequalities with respect to gender are manifested in unexpected ways, and collecting gender data even in settings where it might not seem immediately relevant may turn out to have significant explanatory or statistical power in the end.

At the same time, it is important to remain aware of what D'Ignazio and Klein name the *paradox of exposure*: the double bind that places those who stand to significantly gain from having data about their lives officially collected in the most danger from that same collection (or classifying) act. Women and gender minorities, when outliers in the dataset, are often made vulnerable due to their small sample size. In these cases, the best approach may be to collect data in multiple categories, but aggregate the data when reporting sensitive results.

4. Include impacted communities in the design phase of any data science project

Another crucial component of ensuring that the appropriate data is collected in the appropriate categories–and that the appropriate questions are being asked–is to include impacted communities as full research partners in any data science project. No one would argue with the fact that community members themselves possess what Anita Gurumurthy, executive director of IT for Change, has called "the empiricism of lived experience" (2018). This knowledge about how things truly *are* is essential to ensuring that any particular data science project is intervening in an appropriate and invited way.

In *Data Feminism*, D'Ignazio and Klein describe this approach as embracing pluralism, with the underlying premise being that we can gain better, more detailed, more accurate, and ultimately more truthful knowledge if we pool together a wide range of perspectives, especially the perspectives of those who are most directly impacted by the issues at hand. We can see the benefit of this approach in the work of the advocacy group Pollicy, which embraced pluralism when seeking to understand how sub-Saharan African feminist movements employed data in their work (2021). Pollicy took a mixed-methods approach that included qualitative as well as quantitative research, and that centered the experiences and expertise of the feminist organizers themselves. The research team began by interviewing key stakeholders across twenty countries, and used the knowledge gained through those interviews to structure a series of focus groups in which larger groups helped to imagine enhancements to their data collection, data management, and data privacy needs.

Participatory design processes such as these can help to ensure that any data-scientific research that is undertaken is directed towards the issues and opportunities that are desired by communities themselves. While participatory processes may involve more time at the outset to organize and establish trust, the end result is more effective, as it reduces the possibility of unnecessary research questions that do not address the real issues being experienced on the ground, as well as the potential harm brought about by unwanted interventions.

5. Attend to the context of any dataset or data project

A complement to incorporating participatory processes into the data science research design process is to attend to the broader context–social, political, and historical–of the data being collected and/or analyzed. The world has seen the importance of context play out with respect to the data on Covid-19. Almost every country has been releasing data on the number of cases, fatality rates, the percentage of the population affected, and so on. Yet each country's data is subject to the particular conditions of its collection. How many tests are conducted in each country, which sub-populations are being sampled, and how truthful are the countries being in reporting their numbers, are only some of the questions being asked in order to understand the uncertainty surrounding the numbers.

Another example underscores the importance of considering the context surrounding the data as it relates to women and girls: the issue of sexual assault. As has long been documented, these data are consistently plagued by both underreporting and misrepresentation– underreporting the result of social stigma and concerns over retaliatory action, and misrepresentation the result of institutions with much to lose, in terms of both profit and reputation, should the data reflect reality. As information studies scholars Lisa Gitelman and Virginia Jackson have memorably explained, the idea of "'raw data' is an oxymoron" (2013). Data always enter into research projects "fully cooked"—the result of a complex set of social, political, and historical circumstances. Attending to the context of any particular dataset leads not only to more accurate and more truthful data analysis, but also helps to ensure the efficacy and appropriateness of any intervention developed in response to that analysis. The approaches documented in "Datasheets for Datasets" (Gebru et al. 2021) and "Model Cards for Model Reporting" (Mitchell et al. 2018) offer two possible approaches to ensuring that the context of any particular dataset remains attached to that data as it travels through the analysis pipeline and into the world.

6. Codify transparency for sustained accountability

Shifting the focus from the context surrounding the data associated with a data science project to the impact of the analysis that is enabled by that data points to the additional need for transparency. The goal of transparency is to reveal the outcomes and impact of the data, code, algorithms and systems by companies, organizations and groups. Increasing numbers of data and tech professionals are questioning the nature of their work and calling for more strategies to prevent digital harms and audit current platforms.

One way is to broaden how we understand the idea of "bias," a term which has become a shorthand for a multiplicity of algorithmic harms. In *Data Conscience*, Marshall proposes that we expand our focus from siloed instances of bias to the "bias wheel," a model in which we account for the cascading effects of human bias, business bias, data bias and algorithmic bias on data systems and across communities. For example, the CoNLL-2003 dataset has served as a benchmark for natural language processing (NLP) performance and other text analysis techniques ever since. The dataset consists of Reuters news stories between August 1996 and August 1997. There are 1,393 English-language articles and 909 German-language ones. These articles were written by mostly white men on topics championed by other white men, and has perpetuated racial and gender oppression (Field 2020). In other words, human bias manifested itself into data and algorithmic bias with cascading impacts to how business practices have marginalized or excluded non-male voices in their products. Given our knowledge of these multiple dimensions of bias, it's time for the dataset to be retired.

Actions like these can be further supported by emphasizing a concept related to accountability, *reflexivity*: the ability to reflect on and take responsibility for one's own position within a world constituted by unequal power. Valuable work in this space comes from academic researchers (see "Research Foundations" above). But even within Big Tech, there is evidence of an increasing sense of reflexivity among employees for their role in creating harmful data systems. Employees have pushed back against Google's work with the Department of Defense (DoD) on Project Maven, which uses AI to improve drone strike accuracy; Microsoft's decision to take $480 million from the Department of Defense to develop military applications of its augmented reality headset HoloLens; and Amazon's contract with US Immigration and Customs Enforcement (ICE) to develop its Rekognition platform for use in targeting individuals for detention and deportation at US borders. This pushback led to the canceling of the Google and

Microsoft projects, as well as political consciousness raising across the tech sector, as evinced by the recent creation of the Alphabet Workers Union.

7. Hold institutions accountable for the harms and failures of data-driven systems

Accountability–or, the act of holding institutions responsible for the consequences when their AI system fails–would seem to be an easy principle to endorse and adopt. But the conundrum lies in the limitations of data inside a digital structure (see #5, on context, above). Working with data necessarily involves stripping the data of its context in some way, whether the data consists of text, audio, image or media:

- When we have TEXT, we lose the context of the messenger's tone
- When we have AUDIO, we lose the context of speaker's body language
- When we have a STILL VISUAL, we lose the description or story behind the creation of that visual
- When we have MEDIA, we can't include our other senses to heighten the experience (like smelling a dish made on a cooking show) (Marshall, p.145)

Organizations like the Algorithmic Justice League, the Distributed Artificial Intelligence Research Institute (DAIR), and Hugging Face are attempting to build data-driven systems with accountability at their core. But accountability isn't sustainable without effective external incentives.

8. Prioritize equity through regulation and governance

Put simply: institutions cannot enforce accountability on their own. We also need regulation and governance that prioritizes the administration of equity and decency. In the United States, we have some form of data privacy legislation in most states, although it remains uneven (Lively 2022). There is also no shortage of new policy ideas, including the proposed American Data Privacy and Protection Act (2022), the state-level California Consumer Privacy Act (2018), the revised proposed Algorithmic Accountability Act (2022) at the federal level, and legal frameworks for algorithmic destruction (Li 2022). These policy ideas follow from the EU's General Data Protection Regulation, first introduced in 2016; its Artificial Intelligence Act, proposed in 2022; and the European Commission's Ethical Guidelines for Trustworthy AI (2021), which are among the broadest and most visible examples of what government-enforced regulation can provide.

But the reality is that corporations will not "do better" without additional external incentives and consequences. Irresponsible, unethical, inequitable corporations must be penalized by taking their profits, destroying their algorithms, *and* deleting their ill-collected data. Marshall makes several specific recommendations along these lines: crafting a legal framework for companies to be prosecuted for malcoding and serving as bad-faith data brokers, creating actionable methods for algorithmic destruction, and enforcing change through impact assessment.

Even still, governance is not enough on its own. Policies must themselves remain accountable to the people they aim to protect. There will always be a role for community engagement, political organizing, and formal and informal protest in keeping governments accountable to themselves and to all of the communities they vow to serve.

9. Acknowledge emotional labor as a key component of data-scientific work

Finally, it is important to acknowledge that data science, like all work in the world, depends upon the labor of numerous people--and not only technical labor but emotional and affective labor as well. One example that underscores the range of forms of labor involved in data-scientific work is the work currently being undertaken by feminist activists around the world to compile the otherwise missing datasets that can document the issue of feminicide–or, gender based killings of women and girls. This work takes time–for example, María Salguero, a Mexico City-based activist, has been working several hours per day since 2016 to compile her own dataset of feminicides in Mexico. Salguero's efforts have resulted in the single largest dataset of these types of crimes in that country–a dataset that has been employed to help families locate loved ones, as well as journalists, NGOs and the Mexican Congress.

But Salguero's labor extends beyond narrow definitions of data work. The act of compiling the data on such traumatic events involves emotional labor as well, as the Data Against Feminicide project learned after interviewing Salguero and other feminist activists across Latin America, the Caribbean, and beyond who are engaged in similar efforts (D'Ignazio et al. 2022). This emotional labor exacts a cost, as has been reported in accounts of the psychological and physical toll on people who work in content moderation (Newton 2019). In the case of the anti-feminicide activists, this emotional labor also holds positive value, as it enables the activists to honor the women and girls who have been killed as they register their lives and deaths as data. But in order for these activists to continue this crucial work, it must be named as such and valued accordingly, both culturally and monetarily.

More broadly, these forms of cultural and commemorative data work have become the forms of labor on which data science increasingly depends, especially data science aimed at achieving gender equity. As such, it is our work as data scientists and data activists, researchers and policymakers, to ensure that this labor is properly respected, compensated, and named.

Recommendations

We are not the only ones to have observed that mainstream data practices require significant revision if they are to be leveraged to achieve gender justice. In the guidelines outlined above, we have attempted to distill the core components of what a revised approach to data science might entail. In closing, we offer this set of recommendations to guide future work:

- Prioritize equity over equality, justice over ethics
- Acknowledge and account for how structural power impacts the creation of datasets and data systems
- Include impacted community members as co-designers in any data science project

- Ensure meaningful and inclusive categories of data collection, aggregating and disaggregating categories to protect vulnerable populations as warranted
- Acknowledge the context surrounding any dataset through documentation and other qualitative forms of information gathering
- Codify transparency through meaningful audits, impact assessment, and individual and collective reflexivity
- Hold institutions accountable for the failures and harms of data systems through forceful legal, financial, and technical consequences
- Prioritize equity and justice through regulation and governance
- Credit and compensate the range of forms of labor involved in data work

These recommendations, together with meaningful community engagement, sustained public advocacy, and political organizing and protest, will enable us to chart a course towards achieving gender equity through data science.

Works Cited

Algorithmic Justice League. (2022). https://www.ajl.org/

Benjamin, Ruha. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity.

Birhane, Abeba. (2020). "Algorithmic Colonization of Africa." *SCRIPTed* 17.2. Web. https://script-ed.org/article/algorithmic-colonization-of-africa/

Birhane, Abeba, Pratyusha Kalluri, Dallas Card, William Agnew, Ravid Dotan, and Michelle Bao. (2022). "The Values Encoded in Machine Learning Research." *FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 173-184. https://dl.acm.org/doi/abs/10.1145/3531146.3533083

Birhane, Abeba, William Isaac, Vinodkumar Prabhakaran, Mark Díaz, Madeline Clare Elish, Iason Gabriel, and Shakir Mohamed. (2022). "Power to the People? Opportunities and Challenges for Participatory AI." *EAAMO 2022: Equity and Access in Algorithms, Mechanisms, and Optimization*. https://arxiv.org/abs/2209.07572

Broussard, Meredith. (2018). *Artificial Unintelligence: How Computers Misunderstand the World*. MIT Press.

Buolamwini, Joy and Timnit Gebru. (2018). "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification Systems." *Proceedings of Machine Learning Research* 81, pp. 1-15. https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf

California Consumer Privacy Act. (2018). https://oag.ca.gov/privacy/ccpa

Chun, Wendy Hui Kyong. (2021). *Discriminating Data: Correlation, Neighborhoods, and the New Politics of Recognition*. MIT Press.

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

Costanza-Chock, Sasha, Inioluwa Deborah Raji, and Joy Buolamwini. (2021). “Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem.” *FAccT ‘22: 2022 ACM Conference on Fairness, Accountability, and Transparency*. https://dl.acm.org/doi/abs/10.1145/3531146.3533213

Criado-Perez, Caroline. (2019). *Invisible Women: Data Bias in a World Designed for Men*. Abrams Press.

D’Ignazio, Catherine, Isadora Cruxên, Helena Suárez Val, Angeles Martinez Cuba, Mariel García-Montes, Silvana Fumega, Harini Suresh, and Wonyoung So. (2022). “Feminicide and Counterdata Production: Activist Efforts to Monitor and Challenge Gender-Related Violence.” *Patterns* 3.7. Web. https://doi.org/10.1016/j.patter.2022.100530

D’Ignazio, Catherine and Lauren F. Klein. (2020). *Data Feminism*. MIT Press.

Field, Hayden. (2020). “An A.I. Training Tool Has Been Passing Its Bias to Algorithms for Almost Two Decades.” *Medium*. https://onezero.medium.com/the-troubling-legacy-of-a-biased-data-set-2967ffdd1035

Fields, Robin and Joe Sexton. (2017). “How Many American Women Die From Causes Related to Childbirth? No One Knows.” *ProPublic*a. Web. https://www.propublica.org/article/how-many-american-women-die-from-causes-related-to-pregnancy-or-childbirth

Garamvolgyi, Flora. (2022). “Why US women are deleting their period tracking apps.” *The Guardian*. https://www.theguardian.com/world/2022/jun/28/why-us-woman-are-deleting-their-period-tracking-apps

Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. (2021). Datasheets for Datasets. *Communications of the ACM* 64. 12, pp. 86-92. https://doi.org/10.1145/3458723

Gitelman, Lisa and Virginia Jackson. (2013). *“Raw Data” Is an Oxymoron*. MIT Press.

Goodman, Rachel. (2018). “Why Amazon’s Automated Hiring Tool Discriminated Against Women.” *ACLU*. https://www.aclu.org/news/womens-rights/why-amazons-automated-hiring-tool-discriminated-against

Gray, Mary L. and Siddharth Suri. (2019). Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass. HarperCollins.

Gurumurthy, Anita. (2018). Keynote at *Data Justice 2018*. Cardiff University.

Hara, Kotaro, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, Jeffrey P Bigham. (2108). "A data-driven analysis of workers' earnings on Amazon Mechanical Turk." *CHI '18: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1-14. https://doi.org/10.1145/3173574.3174023

Hill, Kashmir. (2012). "How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did." *Forbes*. https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/?sh=2eb3ee116668

H.R. 6580: Algorithmic Accountability Act. (2022). https://www.congress.gov/bill/117th-congress/house-bill/6580/text

Hodshire, Soren. (2022). "Understanding How Transgender People Fit Into Medical Research and Clinical Trials. *Healthline*. https://www.healthline.com/health/transgender/transgender-medical-research-clinical-trials

Iyer, Neema, Chenai Chair, and Garnett Achieng. (2021). Afrofeminist Data Futures. *Pollicy*. https://archive.pollicy.org/wp-content/uploads/2021/03/Afrofeminist-Data-Futures-Report-ENGLISH.pdf

J.E.D.I. Collaborative. https://jedicollaborative.com/

Keyes, Os. (2018). "The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition." *Proceedings of the ACM on Human-Computer Interaction, Vol. 2, No. CSCW*, Article 88, pp. 1-22. https://doi.org/10.1145/3274357

Li, Tiffany. (2022). Algorithmic Destruction. *SMU Law Review*, forthcoming. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4066845.

Liu, Katherine A. and Natalie A. DiPietro Mager. (2016). "Women's involvement in clinical trials: historical perspective and future implications." *Pharmacy Practice* 14.1. https://doi.org/10.18549/PharmPract.2016.01.708

Lively, Taylor Kay. (2022). "US State Privacy Legislation Tracker." https://iapp.org/resources/article/us-state-privacy-legislation-tracker/

Mangurian, Christina, Eleni Linos, Urmimala Sarkar, Carolyn Rodriguez, and Reshma Jagsi. (2018). "What's Holding Women in Medicine Back from Leadership. *Harvard Business Review*. https://hbr.org/2018/06/whats-holding-women-in-medicine-back-from-leadership

Marshall, Brandeis. (2022). *Data Conscience: Algorithmic Siege on Our Humanity*. Wiley.

Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, Timnit Gebru. (2019). "Model Cards for

Model Reporting." *FAT* '19: Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 220-29. https://doi.org/10.1145/3287560.3287596

Muller, Michael. (2011). "Feminism Asks the 'Who' Questions in HCI," *Interacting with Computers* 23.5, pp. 447–449.

Newton, Casey. (2019). "Bodies In Seats." *The Verge*. https://www.theverge.com/2019/6/19/18681845/facebook-moderator-interviews-video-trauma-ptsd-cognizant-tampa

Noble, Safiya Umoja. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press.

Ọnụọha, Mima. (2018). *Library of Missing Datasets*. https://github.com/MimiOnuoha/missing-datasets

SisterSong, National Latina Institute for Reproductive Health, and Center for Reproductive Rights. (2014). *Reproductive Injustice: Racial and Gender Discrimination in U.S. Health Care*. Center for Reproductive Rights.

Raji, Inioluwa Deborah and Joy Buolamwini. (2019). "Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products." *AIES '19: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 429-435. http://doi.org/10.1145/3306618.3314244

Raji, Inioluwa Deborah, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. (2020). "Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing." *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 33-44. https://doi.org/10.1145/3351095.3372873

Raji, Inioluwa Deborah, Emily M. Bender, Amandalynne Paullada, Emily Denton, Alex Hanna. (2021). "AI and the Everything in the Whole Wide World Benchmark." *35th Conference on Neural Information Processing Systems (NeurIPS 2021) Track on Datasets and Benchmarks*. https://doi.org/10.48550/arXiv.2111.15366

Ricaurte, Paola. (2019). "Data Epistemologies, the Coloniality of Power, and Resistance." *Television and New Media* 20.4. https://doi.org/10.48550/arXiv.2111.115366

Toxtli, Carlos, Siddharth Suri, and Saiph Savage. (2021). "Quantifying the Invisible Labor in Crowd Work." *Proceedings of the ACM on Human-Computer InteractionVolume 5Issue CSCW2*. https://doi.org/10.1145/3476060

Young, Alison. (2021). "Deadly Deliveries: A USA Today Investigation." *USA Today*. https://www.usatoday.com/series/deadlydeliveries/.