

# Epistemic Power in AI Ethics Labor: Legitimizing Located Complaints

David Gray Widder

david.g.widder@gmail.com

Digital Life Initiative, Cornell Tech  
New York City, New York, USA

## ABSTRACT

What counts as legitimate AI ethics labor, and consequently, what are the epistemic terms on which AI ethics claims are rendered legitimate? Based on 75 interviews with technologists including researchers, developers, open source contributors, and activists, this paper explores the various epistemic bases from which AI ethics is discussed and practiced. In the context of outside attacks on AI ethics as an impediment to “progress,” I show how some AI ethics practices have reached toward authority from automation and quantification, and achieved some legitimacy as a result, while those based on richly embodied and situated lived experience have not. This paper draws together the work of feminist Anthropology and Science and Technology Studies scholars Diana Forsythe and Lucy Suchman with the works of postcolonial feminist theorist Sara Ahmed and Black feminist theorist Kristie Dotson to examine the implications of dominant AI ethics practices.

By entrenching the epistemic power of quantification, dominant AI ethics practices—employing Model Cards and similar interventions—risk legitimizing AI ethics as a project in equal and opposite measure to which they marginalize lived experience as a legitimate part of the same project. In response, I propose *humble technical practices*: quantified or technical practices which specifically seek to make their epistemic limits clear in order to flatten hierarchies of epistemic power.

## ACM Reference Format:

David Gray Widder. 2024. Epistemic Power in AI Ethics Labor: Legitimizing Located Complaints. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3630106.3658973>

## 1 INTRODUCTION

In a volume subtitled “An Anthropologist in the World of Artificial Intelligence” compiled after her untimely 1997 death, Diana Forsythe studies and critiques AI cultures in ways that still resonate today. She notes how the construction of what is considered “knowledge” in AI communities often “deletes the social,” privileging codified technical expertise over social knowledges and context [32], which in her analysis of the development of medical AI systems,

led to their failure to usefully diagnose ailments. But more than two decades on, what were once “toy” systems in Forsythe’s time are now the systems many are subjected to daily, even as both the definition [90] and functionality of AI remain questionable [81]. In response, many companies and governments have AI ethics initiatives [10, 49]. This is not without conflict, as people leading these efforts have been fired: in perhaps the highest profile case, Timnit Gebru was fired from Google [68] while attempting to publish a research paper raising concerns about the environmental and bias impacts of ever larger language models [11]. Microsoft recently laid off one of its AI ethics teams, defending this decision by saying the central ethics team was broken up and individual workers moved “within the individual product teams that are building the services and the software” [85]. For those not fired, critical scholarship points to logics including techno-solutionism that constrain leading ethics programs at technology companies [67].

Integrating empirical interviews together with theory, this paper shows how the work of AI ethics is often understood as a “subjective” endeavor, and thus within the technology cultures in which this work sits, is cast as lower status when compared with the engineering work, seen as objective or rational.<sup>1</sup> This is predicted by Forsythe’s argument about the “deletion of the social” [32, p. 28], which she positioned as a detriment foremost to AI functionality, but in our context also operates to delete social knowledges from AI ethics discussions, and position such discussion as lower status. For women and other minoritized groups doing ethics work and disproportionately filling ethics roles, this interacts with documented ways in which their knowledge and work is devalued and delegitimized [94]. In response, I show how some participants seek to legitimize ethics work by casting it as objective, quantitative and therefore authoritative. However, I show how other participants attempt to speak about ethics from their own situated, embodied perspective, but find their attempts delegitimized.

In some ways, scientific documentation of disparity in AI is crucial, and powerful: Buolamwini and Gebru’s Gendershades project tested and documented gender and skin type disparities in facial recognition systems [20], which have now been cited in testimony before the U.S. Congress [18], and in popular media outlets [60], and is likely now one of the most widely recognized examples of AI ethics failings. Both Model Cards [69] and Datasheets for

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
FAccT '24, June 03–06, 2024, Rio de Janeiro, Brazil

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-0450-5/24/06  
<https://doi.org/10.1145/3630106.3658973>

<sup>1</sup>The definitions of objectivity and subjectivity have taken nearly polar opposite meanings across history [23, p. 31]. The more familiar “new” senses of these words define so-called objective phenomena as not interpreted, and are thus taken to be “external to or independent of the mind” [28], and subjective phenomena are “relating to the thinking subject ... proceeding from or taking place within the individual consciousness or perception; having its source in the mind” [27]. In recognizing that all knowledge is situated and embodied, and thus not independent of the mind [46], I reject this dichotomy even while recognizing it to be a social reality for many, including for many of my participants.

Datasets [38] are widely seen as best practices, and have come to enjoy the authority of scientific “Nutrition Facts” on packaged food, both points recognized in practitioner blogs (“like nutrition labels [one can read a] model card, and quickly understand...” [101]), and via pictographic representations shown on Google’s prominent blog post about Model Cards [39].<sup>2</sup> The examples above are extremely important to the extent to which ethics is taken seriously in technology cultures and organizations, and have become touchstones by which people communicate about AI ethics work.

This push towards standardized processes can also be seen in attempts to codify AI ethics into rules and standards such as the European Union’s AI Act [64], and the US National Institute of Standards and Technology’s “AI Risk Management Framework” [74]. This NIST framework asserts that while many characteristics that would make an AI system “untrustworthy” are “directly connected to societal dynamics and human behavior,” and that indeed, views on “what makes an AI technology trustworthy differ,” it nonetheless asserts that “there are certain key characteristics of trustworthy systems,” such as “fair ... accountable and transparent” [74]. In this, the NIST framework provides for the contextual nature of trust [96] to some degree, but then narrowly circumscribes this in a way perhaps more insidious than its wholesale omission. It acknowledges situatedness in how different people may understand what trust is, only to then reduce trust shallow techno-solutionistic questions of design [40], by asserting “fair ... accountable and transparent” as universally important principles, with a stable and shared meaning, even while other research shows how the political economic context of inter-firm AI development shreds this minimal context [36].

However, if we conceive of ethics as an inherently and productively contested domain, as past work has called for [102], this push for legitimacy through codification and quantitative objectivity is likely to be counterproductive and harmful. This is especially pernicious for members of minoritized groups seeking to raise ethical concerns grounded in their lived experience, yet who are not afforded a “view from nowhere,” an epistemic privilege Timnit Gebru discusses in an AI context in her 2021 lecture [37]. This will also serve to underwrite exclusionary pedagogy by delegitimizing ways of knowing from disciplines beyond computer science, as Raji and coauthors demonstrate in their analysis of AI ethics course syllabi [82]. To begin to interrogate these politics of objectivity and examine epistemic power in AI ethics labor, the theoretical orientation of this paper begins with Lucy Suchman’s notion of Located Accountability [89], which draws on Donna Haraway’s Situated Knowledges [46] to envision a feminist ethic of accountability in technology production enabled by locating oneself in this process, and remaining engaged with—and prepared to intervene in—the effects of this technology. Further, insofar as AI Ethics is cast as, and indeed can often become, a complaint voiced in opposition to existing development trajectories or business imperatives [99], raising ethical issues may be seen as a complaint. For this, even while she writes about harassment complaints in universities rather than

<sup>2</sup>It is interesting to note that this Google blog post is based on a paper by Margaret Mitchell and Timnit Gebru (and others [69]), yet unsurprisingly, I can find no similar blog post bragging about a different paper which included these two authors [11], a paper at the center of their unceremonious firing from Google [68]. More recent Google publications have failed to cite even the Model Cards paper, despite discussing Model Cards [92].

technology ethics, I look towards postcolonial feminist theorist Sara Ahmed’s study of complaint activism [6]. Ahmed shows how without a concerted effort towards “reorientation,” complaints must be expressed in ways legible to institutional logics, in turn further entrenching the epistemological supremacy of the logics that had refused to take these complaints seriously. I explore ways to engage in this reorientation in the later sections of this piece, by hearing intersectional feminist Audre Lorde’s call to not to reach for sources of legitimacy that reinscribe oppressive systems [61], and Black feminist theorist Kristie Dotson’s understanding of the intractability of flattening epistemic power differentials and suggestions for ways forward [29].

The empirical basis for this paper is 75 interviews with technologists conducted remotely between early 2021 and late 2023, in a multi-sited approach [100]. Purposive and snowball recruitment [77] helped reach participants in a variety of contexts: the majority work in companies ranging from startups to multinationals, many work on open source projects, several are academics, and a smaller group are activists seeking to question local development of carceral technology. Many inhabited and moved between multiple such categories. Participants were in four continents: the majority based in North America and Europe but also including those in Asia and Africa. The majority were doing AI work at some point in its “supply chain” [22, 98]: spanning dataset curation, academic AI research, building new models, to deploying them in consumer or B2B contexts. Corporate participants were often required to engage with corporate AI ethics practices. This includes some who have accepted AI ethics as part of their work, whether through intrinsic motivation or as part of their assigned duties [e. g., 67], and many who have not. Interview questions were tailored to context and did not adopt consistent predetermined structure[87]: some portion of participants spoke to me about their own ethical concerns, often after raising them to their employers and seeing them go unaddressed. This study received IRB or other ethics approval as appropriate.

The variety of contexts in which I investigate AI ethics practices allows us to understand how a variety of epistemological stances interact with the legitimacy of AI ethics claims, and what is rendered possible or not as a result of these stances. Participants were invited to self-identify their gender: 57 identified as men, 13 identified as women, four identified as non-binary, and one declined to answer. Most interviews lasted one hour (though ranging between 25 minutes and 3 hours), and were transcribed and analyzed using inductive thematic analysis [15, 16]<sup>3</sup>, under an *interpretivist* epistemological paradigm [59].<sup>4</sup> My analysis is therefore inflected by my own positionality: I am a male-identified scholar in the United States with formal training in computer science, mentored by feminist STS scholars, with experience organizing against pervasive surveillance [41, 42] and carceral technology [24, 54]. It is from

<sup>3</sup>In an *inductive* thematic analysis, themes are grounded in and emerge from the data, as opposed to *deductive* analysis where themes are decided *a priori*, and is properly a *reflexive* process, where the researcher reflects on their assumptions, training, values, and how this may shape analysis. See also: [15, 17].

<sup>4</sup>Interpretivism, similar to its younger cousin constructivism [86], understands that individuals are experts in interpreting their own experience, asks researchers to deeply understand these interpretations, and with both sharing the task of constructing wider meaning from these. To positivists, this is an unacceptable limitation, as no view from nowhere objectivity is possible. See also [23, 37, 46, 59].

attempts to reconcile differences in language and legitimacy between the first and last of these—computer science and carceral tech activism—that I am motivated to write this paper, but my positionality necessarily limits those I can talk with and inflects the interpretations I arrive at, and others may arrive differently. In some ways, these locations represent the “stepping in and stepping out” characteristic of fieldwork [79] in that some of my nonbinary and woman-identified participants construct their AI ethics practice as self-consciously feminist or even necessarily constructed in femme-only spaces, yet also, the predominantly technical and male backdrop in which AI ethics work sits paints me as part of the ingroup relative to the bulk of participants in this work: men in technical roles.

This paper proceeds as follows: I first show how AI ethics labor is often seen as a chore, in contrast to the “real work” of building AI systems. I then show how attempts to legitimize this work by casting it as a quantitative, authoritative and even an automated project have succeeded to some extent in “elevating” it, by constructing it in the same terms “real” AI engineering work is constructed. I then turn to attempts by some participants to make ethics complaints in ways that eschew quantification and instead come from the insight of their own particular location and experience, and recount how such efforts are often delegitimized. Then, I examine how some participants seek to construct alternative AI ethics practices, often based on their embodied experiences, and largely but not entirely beyond objective institutional logics of legitimacy. I conclude by motivating and explaining humble technical practices as those which specifically seek to make their epistemic limits clear, sketching ways they may be implemented, and placing such practices in dialogue with work on liberatory uses of quantification, critical technical practices, participatory and intersectional AI, and epistemic power relations in AI.

This paper contributes the following:

- This work provides an empirically-based catalog of AI ethics labor conducted from a variety of epistemological stances and contexts, and examines their sources of—and the extent to each is afforded—legitimacy, complementing scholarship which scrutinizes hierarchies of knowledge in AI ethics educational contexts [82], the supremacy of quantification in AI fairness work in corporate contexts [26], and the use of seemingly “objective” algorithms in public sector [80].
- Flowing from this, this work attempts to frame discussion of hierarchies of knowledge between the wide range of academic disciplines (ostensibly) included in AI ethics discourse: from those concerned with “the development and deployment of algorithms” to those “examining human experiences, needs, perceptions” [2], and between these and attempts to include knowledge from activist and lived experiences outside the academy [1], by examining how these different knowledges are (not) afforded legitimacy in practice.
- Excellent exceptions notwithstanding [e.g., 8, 45, 52, 78, 84, 91], reflexive analysis of AI ethics scholarship argues that more space needs to be made for feminist, Black, decolonial and non-Western thought [13, 58, 82]. In its use of theory to interpret empirical data, this paper demonstrates the relevance of feminist STS [32, 46, 89], post-colonial scholarship [6, 93], and Black feminist thought [29, 61] to AI ethics questions.
- Using this theory, this paper demonstrates the role of *epistemic power* in the context of AI ethics labor, joining work bringing an epistemic lens to algorithmic fairness [31]. This responds to reflexive examinations of past AI ethics scholarship calling for greater attention to structural power asymmetries [13, 58] and growing examples of research on other kinds of power in AI ethics, such as in algorithmic decision-making [50], workers raising ethical concerns [99] and engaging in collective action [14], corporate capture co-opting academic critique [102], and between those subjected to algorithms and those developing them [12]. This localization of epistemic power in AI ethics labor is used to motivate and sketch the concept of *humble technical practices*, drawing on related past work on critical technical practices [5, 65].
- Finally, in examining epistemic power hierarchies, this paper provides a deeper diagnosis for the move to technosolutionsistic practices of AI ethics, widely critiqued as problematic and insufficient [3, 40, 58, 67, 70], by locating one reason for this move at the level of epistemic power.

## 2 THE LOWER STATUS OF ETHICS WORK WITHIN AI CULTURES

This section demonstrates how my participants understand ethics work to be lower status relative to “real” AI work, but similar phenomena have been examined before. Past research shows how certain work in technology communities, like writing documentation or building community, is often done by women, and is less valued than writing code [71]. More recent work within AI has even shown differences between high status work on models versus lower status work on the dataset models are built from [83], and recent manifestos have attacked “tech ethics” on the basis that it hinders “progress” on endeavors such as “Becoming Technological Supermen” [9]. Forsythe shows how these status divides are inscribed in the language used by the AI practitioners she studied: they restricted their use of the word “work” to what they think of as ‘real AI’: building systems” [32, p. 26], set in contrast to the other kinds of work such as grant writing, meetings, and also the “knowledge elicitation” interviews seen as boring but necessary to “capture” knowledge to encode into systems.

When asked about whether his community of deepfake content creators had ever discussed ethics, one participant replied: “Uh, no. No. I don’t think so. [...] It’s mostly technical, you know, stuff and, and sharing information about models. Um, another metaphor I use is like, it’s like long exposure photography, but like from like the 19th century where you’re carrying around plates, you know, and you have to be like a pristine, like technician in handling all the parameters to set up your camera and everything.” In this metaphor, he explained how the work deemed worthy of discussion is focused on the technical craft of making realistic deepfakes, not the ethical questions involved in doing so. Notably, these two are seen as distinct and separate, a stark example of the more widespread division between technology and society [57]. This community was overwhelmingly men, in contrast to other communities examining the gendered impact of deepfakes [63].

I often asked participants what work they were most excited about, and what work they saw as a chore. The most common task

seen as a chore was “paperwork” needed to release new systems, which notably included AI ethics questionnaires, often included filling in model cards. One described how each of the steps needed to release code could be a Byzantine “trap,” and that each task was meted out randomly to teammates as necessary and undifferentiated homework. Another referred to this as “light work,” and another said that this work did not “feel like I was doing something that I need to do,” but merely that it was a “have to” do.

One woman machine learning engineer described how she had once previously gone through these release steps, and now she was the one pigeon holed into doing these tasks evermore: “It’s mainly me on my team who’s going through the [release steps...] I had to figure it out for one product and after that it has always been, [name] you know how to do it. You go ahead and do it.” Forsythe wrote about how women and their work in AI labs were marginalized: they were administrators or secretaries, and by excluding their work from the category of the “real work” of building AI systems, they were not included as members of “the lab.” In one case, she recounts how they were “rendered invisible through the systematic bracketing out of people who perform work that is gendered as female” [32, p. 172]. While our previous participant had a masters degree and an AI engineering job title, she was nonetheless the only woman on her team and the one told to do the work coded as a “chore.” Taken together, these examples show how AI ethics work is seen as a lower status “chore.”

### 3 AUTOMATED MODEL CARDS: LEGITIMACY VIA QUANTIFIED OBJECTIVITY

Whereas in the last section, we saw how doing ethics work is often accepted as an unwelcome but necessary chore, here we see how others seek ways to legitimize its value. In her study of sexual harassment in universities, Sara Ahmed wrote how students’ complaints were perceived as “legible” to their institution only when expressed in a certain way, using “evidence that took the right form” [6, p. 5], and often faced threats to their career when raising these complaints. In an AI ethics context, Madaio and coauthors found that advocating for AI fairness issues can bring perceived career risks, but that checklists can empower individuals to raise these issues [62].

For some, pushing for legitimacy for AI ethics involved seeking to make ethical evaluation into an “objective” process through the use of automated tools. One AI engineer, on reviewing his organization’s model card template, which already sought to walk through the kinds of questions which should be asked when thinking about ethics, lamented that “A lot of filling out the model card is subjective, right, because it’s based on my understanding of how we use the model and my understanding of the data, the data set.” When I asked why this might be an issue, he said “because, if I give the card and the data set and the model to another person to answer the question, his answer might be different.” To him, ethics ought to be decidable and rational, and therefore he believed model cards ought to be able to conclusively “answer” questions of ethics, rather than open up such questions for contestation.<sup>5</sup> To help achieve

<sup>5</sup>While the original paper proposing Model Cards “specifically aims to *standardize* ethical practice and reporting” (emphasis mine), the paper does so in order to enable questions to be asked “across different institutions, contexts, and stakeholders,” and explicitly positions it as one technique among many needed interventions [69].

these ends, he said how on his team “we try to automate a lot of that now” by building “tooling to run the data and the model [...] that’s actually, objectively looking at it,” calculating numbers to generate a model card. Later in the interview, he reflected “I haven’t really figured out if you actually need a person to answer these questions, or the tooling can answer these questions for you.” This was also apparent in how he spoke of ethical issues, for example talking of ethics issues as if they could be straightforwardly fixed as soon as better tools were available to enable model explainability, saying that “The de-bug-ability of it so that at least you can fix the problem” is most pressing to him. Finally, when asked what came to mind when he heard “ethical AI,” his first response was “fairness,” saying that it is “based on a lot of definitions that we have around nondiscrimination ... policy”, attending to definitions and policy as first order concerns.

Automating model cards was one way to cast ethics as quantitative and objective, but others sought to base their claims to AI ethics by citing evidence from outside authorities. In one group discussion about AI ethics I witnessed in particular, participants persistently backed up points they raised by citing a variety of outside sources, including: the book *Thinking Fast and Slow*, the book *Bowling Alone*, the notion of the Dunbar Number, and “a study done way back in 1990s, 1998 or 96, something like that.” In a shorter follow up conversation one-on-one, a participant brought up “Checklist Manifesto. Have you read that book? It’s a fantastic book. It’s by Atul Gawande,” and “this website called [the] AI incidents” database. In this way, ethical arguments were seen as especially valid when they were based on an external, and often scholarly, reference point. The reference to the AI Incidents Database [66] is particularly interesting: it is a resource which helps to expand the frame of discussion by providing evidence for the contextuality of AI harm, but its invocation suggests the need for such evidence to be particularly structured and collated to be seen as especially legitimate.

Finally, the participant in a professional community committed to making “quality” deepfakes (discussed above), when asked about ethical issues, said that he is “bothered” that most discussions of ethics of deepfakes center on geopolitical misinformation concerns (eg, faking politicians) rather than the predominant case of deepfakes: non consensual porn [7] (this discrepancy examined using a feminist lens in prior work [63]). To him, “objective” debate was needed: “like let’s have an objective, you know, conversation about the benefits and the disadvantages [of deepfakes...]” At the same time, he reasoned that most deepfake porn may not be high quality enough to nor common enough to be concerning in his desired “objective” debate: “I know it’s like 90% is like pornography, but it’s also, you can have a debate about, well, how many users are doing that? No one really knows, you know? And, and to be honest, those fakes aren’t [...] at] the level that the [professional] people that I work with.”

### 4 GREY SKIN AS TECHNOFIX: FAILURES TO LODGE LOCATED COMPLAINTS

It is analytically notable that both Ahmed and Suchman use the word “location” Suchman calls us to acknowledge the located nature of our “vision”—that our perspective on the world is not a “view

from nowhere” but is instead based in our particular social, political, geographical, and embodied worlds [89]. Ahmed speaks of location to denote the misdirection of blame: “to become a complainer is to become the location of a problem” [6, p. 3], in how a complainant might be told that they are the problem for misinterpreting things or raising a fuss, instead of treating the harassment or harasser as the problem. Both Suchman and Ahmed use “location” to denote the particular partial perspective from which one’s views are generated, in Ahmed’s case denoting the particular cues by which a student is attuned to a pattern of harassment mistakenly read as innocent to the “objective” admin who would adjudicate the matter. Those “too close” to the harassment were seen as unable to be objective.

Many participants report that when raising ethics complaints that are based on their own particular location, these are often delegitimized. For example, a woman of color working on virtual reality tech (the only woman of color on her team) asked for small steps to make their tech more inclusive for people of different skin colors, or those missing a hand. She reports: “And virtual reality by itself, I will say is not very accessible, right? [...] So a really famous VR software at the time, had done inclusivity, in terms of the color of the skins of hands, right, and allowing for people with one hand to operate it. So if you didn’t want to use white hands or stuff like that, and I brought it up as an option, because we were working with folks who could be helped by that.”

In line with the previous section, she had first tried citing “papers on why diversity helps on how people will come back to your application,” attempting to legitimize her concern using both outside scholarly authority and the valued business imperative of attracting more users. The work to legitimize this complaint by using the kind of evidence and imperatives seen as legitimate by her company constituted labor, which fell to her. But this approach failed, so this participant noted to her team how this concern was rooted in her own usage of the system, seeing white hands projected in VR in place of her color: “I don’t have white skin and the device did and I was like: Wait, if I’m developing it, if people who are using it, [they] might feel even more intensely about it.” On one hand, she noted how “it was easier to sort of bring up those diverse opinions” by virtue of her identity, but on the other, she noted how she was concerned that others would ask whether her “perspective [is] coming because I am a unique person [...] or is it only coming because I’m a woman and then a woman of color?” As a result, she said that she didn’t want to “rock the boat too much” for fear these rocks would be dismissed by virtue of being “quirks” of her identity. Indeed, she said “I think [that] sometimes I could be considered intense.”

Ultimately, the accessible and inclusive VR features she asked for were not implemented, “And then I was sort of told, Well, nobody asked for it.” Here, “somebody” would refer exclusively to paying users, meaning she was included in the empty group of “nobody.” Instead of customizable skin colors, the skin was made grey, which she named as a “colorblind” outcome, and disliked. This sidestepped situated discussions of skin color, and instead, adopted grey skin as a kind of objective skin from literal nowhere.

Ahmed writes that “we realized that to be heard we had to make our experiences legible [...] For some of us, the institution expected

emotion and hurt to be expressed. For others [...] emotional expressions were viewed as irrelevant or even detrimental to complaint.” [6, p. 267-268]. While the context of sexual harassment in universities is different than the context of this participant, she was caught in a similar double-bind in choosing how to base her claim: on one hand, she had tried to cite academic evidence on why inclusion would help product success. On the other hand, she had attempted to speak from her own situated perspective, as a woman of color, for why she did not want VR to white or grey-wash her hands, but instead was implicitly called “nobody” — outside of the role of the customer, her own ask for more inclusive features did not count.

## 5 ALTERNATIVE AI ETHICS: SPACE FOR EMBODIED COMPLAINTS

Upon thinking back to an AI ethics conversation she had convened with exclusively femme work colleagues, one participant reflected how this changed the tenor of discussion, creating a shared “feeling like it’s [...] a part of your lived experience to be harmed by technologies, versus feeling like it’s an abstract intellectual conversation.” She went on to say how shared experience may make people feel comfortable to “give more of yourself,” because she can “feel more comfortable when you’re sort of assuming that the people in the room that are there with you are on the same page” because of shared experience of marginalization, even while particulars may differ. In this way, she felt that discussing AI ethics with other marginalized people legitimizes speaking from “lived experience” instead of conceiving it as an “abstract intellectual” conversation. Here, gaps left by the latter are filled by the former, and it is important to note that even those engaging in “abstract intellectual” conversations do so from lived experience, but are simply more limited in what they recognize that this experience teaches them. In a recent discussion, she raised “foreign intervention” as a concern for discussion, based in part on the fact that “I’m from Latin America, my family’s from Latin America originally. And I think a lot about the impacts of technology on a global scale.” She contrasted this recent discussion with past experiences of group discussions at her workplace which included men: “it’s also worthwhile to like, point out that we’re all women. [...] there’s a different way these conversations happen [...] when there’s like, male genders participating in activities like this. [...] men take up space in particular ways.”

One participant, who identified as a “grumpy feminist anthropologist” told of her past experience taking an AI ethics class, where student groups would sign up to present and lead discussions of tech ethics issues in the news. She mentioned how she “was paying attention to [the] gender component of those discussions.” Noting that there were “much more men than women” in the class generally, and that in particular: “the people who signed up for discussion on [the article] ‘Google will know I’m pregnant before I [do]’ were men. And they were talking about periods and being pregnant and all these things without like, any personal experience.” She discussed how space was not created to discuss this topic from particular personal experience, instead treating all discussants as having equal—detached and supposedly objective—perspectives on the issue. She goes on: “There was no acknowledgment of [the] embodiment of the experience when it comes to ethics and AI,”

mentioning that this was “pretty memorable [...] as a woman and as somebody who gave birth and was pregnant.” Speaking of a more recent AI ethics conversation, she suggested the way it was constructed limited discussion in “what kind of terrain, what kind of habits, what kind of people” fit and did not, giving an example of a local street vendor selling juice at what used to be a rich and busy gathering spot, but whose livelihood has been taken through the uptake of delivery apps. Speaking how the “example of [the] juice maker doesn’t fit in here,” she characterized the way the conversation was constructed as an “epistemic kind of violence” in how it adopted and tacitly enforced a “certain imagery of the culture.” A different participant spoke similarly, reflecting on how she often put forward her experience as an immigrant when discussing technology: “Sometimes I do it very consciously to like, put my foot in the room or like to remind [of] this context, that it affects everything. Like, everything that happens and gets built in the US is going to get exported as an idea if not an actual product into the rest of the world and into the global south.” She spoke of how she does this even when conspicuously “conscious about how I, in this country, I’m talking about this ‘other place’ like I’m often the only person bringing up an ‘other place’,” aware of how this may ‘other’ her but seeing this as important nonetheless.

Others spoke concretely about how embodiment and situated experience played a role in their own development process, and how this intersected with corporate logics. One participant was developing a body scanning health product, mentioning: “body scanning in particular is sensitive information, and so we recognize from the beginning that if you show somebody a body scan, number one, different people are going to react to it differently. There is great potential to have someone get motivated or track progress. However, there’s also great potential to exacerbate body image issues. So that’s something we’re very sensitive to in the way we designed the app and the way we present data.” He described the ways his team had sought to design the interface to account for these differing reactions, including a “multi screen tutorial that walks the user through; how should you think about this information? You might be shocked by this information. That’s OK. This is about you achieving your goals. And for each piece, each measurement, for example, we would provide in that tutorial, an explanation. Here’s your waist circumference, here’s why it’s important, and here’s a bit of research around it.” While echoing discourses of legitimacy through scholarly medicine as seen above, this participant spoke of how as part of the “venture incubation process” they would themselves volunteer to serve as each other’s pilot participants, recognizing how much of it is “a personal experience having [one’s body] evaluated [by] 3D body scanning in a research context.” He spoke of how they had developed mutual vulnerability through “the opportunity to get to know each other. And so honestly, we have a pretty good team rapport. [...] we check in on each other outside of work, and so having that familiarity is definitely helpful, in talking about this.” However, he also spoke about how a shared “passion for the user, for the customer” motivated the attention to the myriad ways people may relate or react to this technology, interestingly, suggesting attention to various experiences can be cast as a market imperative.

## 6 CONCLUSIONS: TOWARDS HUMBLE TECHNICAL PRACTICES

“The master’s tools will never dismantle the master’s house. They may allow us temporarily to beat him at his own game, but they will never enable us to bring about genuine change” [61]. Audre Lorde delivered these famous words to a conference of white feminist academics, pointing out how they were ostensibly advocating feminism but doing so with and within the patriarchal structures of academia. Her words are read to “challenge reformist feminists to become radical feminists” [75], imploring them to lay down the allegorical tools of the master that may give them academic credibility while simultaneously holding down others who do not have access to these tools. Lorde asks: “What does it mean when the tools of a racist patriarchy are used to examine the fruits of that same patriarchy?” I ask an analogous question: if AI ethics relies on quantified neutral objectivity for its legitimacy, what will that foreclose? And Lorde responds: “It means that only the most narrow perimeters of change are possible and allowable.” The simple point of my piece is analogous: we reach a limit on the legitimacy that AI ethics can achieve while understanding itself as neutral and objective, which will thereby delegitimize practices of AI ethics that authentically integrate lived, located and embodied experiences. Put even more directly, by entrenching the epistemic power of quantification, the wide use of Model Cards and their ilk risk legitimizing *AI ethics* as a project in equal and opposite measure to the extent that they delegitimize and marginalize embodied and lived experiences as legitimate parts of the same project.

Using the metaphor of sight, Haraway asks us to reject the “deadly fantasy that feminists and others have identified in *some versions of objectivity*, those in the service of hierarchical and positivist orderings of what can count as knowledge” [46, p. 580, emphasis mine]. Crucially, she does not give up on objectivity wholesale, and instead asks us to construct a feminist objectivity that is grounded in “situated and embodied knowledges and an argument against various forms of unlocatable, and so irresponsible, knowledge claims” [46, p. 583].<sup>6</sup> In this way, feminist objectivity makes clear how objectivity is stronger when it acknowledges that all knowledge is socially situated, and the ways in which lived experience builds and inflects knowledge. And further, it is not quantification, *per se*, that are the “masters tools,” instead it is the doctrinal positivist objectivity which disguises a hegemonic perspective as a “view from nowhere,” but which often adopts quantification as its tool.<sup>7</sup> Indeed, while examples of limited or oppressive uses of quantification abound [19, 73, 95], we can also find examples of interpretive quantification practices calculated toward better understanding one’s embodied experience [72], documenting wider trends of oppression [30, 91], phenomena which affect our lives

<sup>6</sup>Haraway’s notion of situated knowledges is similar to the Harding discussion of strong objectivity [47]. Haraway also emphasizes that she has no patience for relativism, calling it equally in denial of the “stakes in location, embodiment, and partial perspective” as both relativism and “view from nowhere” objectivity “make it impossible to see well” [46, p. 584]. Instead, she suggests how leaning in to the “particularity and embodiment of all vision” thus “allows us to construct a usable, but not innocent ... objectivity” (582).

<sup>7</sup>Others have pointed out that an objectivity rooted in positivism is not commensurable with interpretivism given deep epistemological disagreements, even while their frequently chosen methods (ie, quantitative vs qualitative) may be commensurable [59, p. 174].

but occur at a scale too large to be appreciated directly [76], as well as other uses of quantification to disrupt power relations or advocate for the marginalized [4, 44, 51, 55]. This warning against entrenching the epistemic supremacy of quantification applies to these examples too: not to suggest that quantification must be eschewed, but instead to argue that steps must be taken to limit its epistemic supremacy in order to make space for other ways of knowing.

Therefore, and perhaps unsurprisingly to feminist epistemologists, I label the “view from nowhere,” with its false guise of neutral objectivity, and the epistemic power that comes with it, as the master’s tools. In her 2021 talk, Gebu talks about how epistemologies of AI adopt this “view from nowhere,” and use it to construct hierarchies of knowledge in AI. She relays how some receive who heard her proposal to include sociocultural understanding as part of ML data collection [48] responded to with the dismissive question “What was the algorithmic component of her work?”—revealing the practice of improving an algorithmic technique on a quantifiable dataset to be the esteemed “pinnacle” [37, 43]. This question does much to disclose what is, and isn’t, valued in the knowledge system of the dominant technical practices of today, and demonstrates how “view from nowhere” algorithms are given epistemic power. Relatedly, in an educational context, Raji, Scheuerman, and Amironesei demonstrate how courses teaching “AI ethics” often construe ethics using the language and techniques of computer science, reinforcing the discipline’s dominance in hierarchies of knowledge, and thus resulting in an “exclusionary pedagogy” not open to other ways of knowing [82].

Black feminist philosopher Kristie Dotson reminds us how difficult of a road it will be to undo these hierarchies, on the road towards feminist objectivity that includes embodied and situated experiences. In her paper, *Conceptualizing Epistemic Oppression* [29], she expands beyond work on epistemic injustice [34] to remind us that the nature of epistemic power means that those holding it are very rarely confronted with the need to reevaluate their epistemic frames, and are therefore able to “absorb extraordinarily large disturbances without redefining [their] structure.” On the other hand, those using non-dominant epistemic systems may be “rejected as nonsensical; [or as] a deceiver with dangerous ideas; or [their] conclusions [...] might even invoke ridicule and laughter,” or simply “deep, awkward silences” when scientists are confronted with indigenous knowledge claims they believe are objectively false [93, p. 143]. Dotson suggests that while “Attempts to convince those who are relatively more epistemically powerful to relinquish some of their power might be warranted as a first line of address,” one might also seek to lessen the effect of faulty credibility assessments.

So, this difficult road implies the need for epistemological pluralism within quantitative or technical practice, in order to enable the incorporation and privileging of other kinds of knowledge, a project related to those others have attempted. Notably, we have Philip Agre’s *Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI* in which he relays his awakening out of dominant AI practices of “write programs that solve problems better than anybody else’s” as the only valid proof of the veracity of a critique [5], and thus the only path to new knowledge, which nowadays usually means quantifying improvement on narrow benchmarks [43]. He critiques AI, but argues for the need to reform AI, into a practice

that is critical and reflexive: a critical technical practice, asking us not to let go of the technical, but to be reflexively critical of its limits. More recently, Malik and Malik theorize “Critical Technical Awakenings” using Paulo Freire’s theory of critical consciousness [33] and emphasize the role of community, especially communities including non-technical and non-positivist thinkers, to enable more such awakenings [65]. This aligns with calls to resist solitary and epistemically homogeneous approaches to teaching tech ethics [82]. Among other issues, though a process Ahmed calls “atomization” [6, p. 180], solitary approaches reduce the ability of a community to fashion and legitimize alternative epistemic practices: seen recently in how Microsoft disbanded a central ethics team and moved the few ethics practitioners who remained to report to technical product teams [85]. Also, notably, there are common pitfalls on the road to epistemological pluralism, such as “participatory” AI approaches. Participatory AI falls short when those asked to “participate” are not given a say in the terms (including epistemic terms) of their participation nor a commitment to their ongoing inclusion [88], when careful participation is sacrificed in service of system scale [25], and when “intersectional” approaches to AI reduce the rich lived experience of intersectional oppression to subgroup parity fairness metrics [52].

Taking Agre’s cogent cultural critique into an understanding of epistemic power, I ask those practicing “AI,” and other related quantified sciences widely practiced as value neutral, to help level epistemic power differentials by underscoring the partiality of their work [89], constructing what I call *humble technical practices*, which specifically seek to make their epistemic limits clear. Beyond the minimum technical humility sometimes demanded in “threats to validity” or “limitations” paper sections, which seek to enumerate and answer for possible technical flaws from within the epistemic frame of the paper, this would include gesturing towards (as doing more would likely be seen as nonsensical, until such practices become routine [93]) the epistemic limits of quantification, and what other epistemic frames may have to offer. Beyond written products, this may be easier in interactive settings such as research talks or conversations with peers, and especially effective when given on prominent platforms. This would involve acknowledging the incomplete view that large-scale data analysis offers, and directing readers or the audience to other forms of knowledge on similar problems. This can also blend well with approaches to share epistemic power: elevating and platforming those with lived experience to serve as the natural expert on the nature of the problems they face and the roots of these problems. This should involve not only attention to citational justice by seeking to give epistemic power to marginalized academics and marginalized academic disciplines through citational practices [56], but crucially must also seek to more broadly recognize knowledge sources not represented in tech companies or the academy. When lived experience is segmented off into its own conference track or side meetings [1], it can literally be difficult to see them as worthy of being in the same room [35].

Epistemic power is bound up with, and often flowing from, other kinds including economic and political power. For example, Forsythe writes how the (poor) appropriation of ethnographic methods in positivist ways by untrained AI researchers had economic outcomes: “my own research funds have dried up” [32, p. 143]. It is therefore those with other kinds of power, along lines of race,

gender, or wealth, that should first share or cede epistemic power. Notably, people harmed by AI are often not paid nor recognized for sharing their lived experience of this harm in the ways that academics may be paid to collect it [88]. Therefore, I call for more ways to legitimize the sharing of lived experience in and on its own terms as valid AI ethics practice, and indeed, the preeminent and most true form of such practice. This must be listened to, and AI Ethics must be seen as a project to legitimize and amplify these on their own terms.

Technology audits led by advocates and activists, instead of those occurring as “closed-door compliance” exercises locked within technology companies and their epistemic frame, are a way to open the doors to other kinds of knowledge and thereby flatten epistemic power hierarchies. Peaks Krafft and coauthors propose a flowchart, questionnaire, and other tools to be wielded by “those with the lived experience of being differentially targeted by surveillance technologies,” specifically positioning this as *expertise* of the sort particularly integral to equity in public sector algorithmic systems [53]. They conclude their work with the sendoff: “The technofuture we project through this work is defiance rather than compliance,” fitting given that their audit framework is designed to enable such defiance to be taken seriously. I echo this call for defiance: we will build the power we need when we speak from our experience, building “collective knowledge of the specific locations of our specific visions” [89], without hesitation and without reaching for other sources of legitimacy.

## ACKNOWLEDGMENTS

This paper began as an answer to an incisive question that Dawn Nafus asked during my dissertation defense. Versions of this draft benefited from her generous comments and those of Blair Attard-Frost, Sireesh Gururaja, Benjamin Laufer, Lucy Suchman, Michael Madaio, and Os Keyes. Walks discussing these ideas with Maria Ryabova, Harini Suresh, Emily Tseng, and Frank Pasquale helped sharpen and encourage this work. I thank the Visual AI and HCI labs of Princeton CS for the opportunity to present early versions of this work, particularly Sunnie Kim, Andrés Monroy-Hernández, and Olga Russakovsky. The early stages of writing this article were coincident with writing a forthcoming position paper written with Claire Le Goues on a similar topic in Software Engineering [97]. Portions of data presented here were collected in collaboration with Derrick Zhen, Laura Dabbish, and Jim Herbsleb. My thinking has benefited from the intellectually diverse atmospheres of the Digital Life Initiative at Cornell Tech, the Feminist Political Economy reading group (open to all, join us!), and in my work with activist groups in Pittsburgh including Against Carceral Tech, the Coalition Against Predictive Policing, and #NoTechForICE. May Carnegie Mellon never again deploy racist predictive policing [21] nor propose IoT algorithmic management of custodial workers [42], and learn from the lived experience of those such systems harm and the activists who fight against them.

## REFERENCES

- [1] 2023. ACM FAcCT 2023 Call for CRAFT Proposals. <https://facctconference.org/2023/cfpcraft>.
- [2] 2024. ACM FAcCT - 2024 Call for Papers. <https://facctconference.org/2024/cfp>.
- [3] J. Khadijah Abdurahman. 2021. Calculating the Souls of Black Folk: Predictive Analytics in the New York City Administration for Children’s Services. *Columbia Journal of Race and Law* 11 (2021), 75–110. <https://journals.library.columbia.edu/index.php/cjrl/article/view/8741>
- [4] Rediet Abebe, Solon Barocas, Jon Kleinberg, Karen Levy, Manish Raghavan, and David G. Robinson. 2020. Roles for computing in social change. In *Conference on Fairness, Accountability, and Transparency (FAT\* ’20)*. Association for Computing Machinery, New York, NY, USA, 252–260. <https://doi.org/10.1145/3351095.3372871>
- [5] Philip Agre. 1997. Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI. In *Bridging the Great Divide: Social Science, Technical Systems, and Cooperative Work*. Erlbaum. <https://pages.gseis.ucla.edu/faculty/agre/critical.html>
- [6] Sara Ahmed. 2021. *Complaint!* Duke University Press, Durham, NC.
- [7] Henry Ajder, Giorgio Patrini, Francesco Cavalli, and Laurence Cullen. 2019. The State of Deepfakes: Landscape, Threats, and Impact. *Amsterdam: Deeptrace* (2019).
- [8] Lameck Mbangula Amugongo, Nicola J. Bidwell, and Caitlin C. Corrigan. 2023. Invigorating Ubuntu Ethics in AI for Healthcare: Enabling Equitable Care. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Chicago IL USA, 583–592. <https://doi.org/10.1145/3593013.3594024>
- [9] Marc Andreessen. 2023. The Techno-Optimist Manifesto. <https://a16z.com/the-techno-optimist-manifesto/>
- [10] Blair Attard-Frost, Andrés De los Ríos, and Deneille R. Walters. 2022. The ethics of AI business practices: a review of 47 AI ethics guidelines. *AI and Ethics* (April 2022). <https://doi.org/10.1007/s43681-022-00156-6>
- [11] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Conference on Fairness, Accountability, and Transparency*. ACM, Virtual Event Canada, 610–623. <https://doi.org/10.1145/3442188.3445922>
- [12] Garfield Benjamin. 2022. #FuckTheAlgorithm: Algorithmic Imaginaries and Political Resistance. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 46–57. <https://doi.org/10.1145/3531146.3533072>
- [13] Abeba Birhane, Elayne Ruane, Thomas Laurent, Matthew S. Brown, Johnathan Flowers, Anthony Ventresque, and Christopher L. Dancy. 2022. The Forgotten Margins of AI Ethics. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAcCT ’22)*. Association for Computing Machinery, New York, NY, USA, 948–958. <https://doi.org/10.1145/3531146.3533157>
- [14] William Boag, Harini Suresh, Bianca Lepe, and Catherine D’Ignazio. 2022. Tech Worker Organizing for Power and Accountability. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 452–463. <https://doi.org/10.1145/3531146.3533111>
- [15] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- [16] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (Aug. 2019), 589–597. <https://doi.org/10.1080/2159676X.2019.1628806>
- [17] Virginia Braun and Victoria Clarke. 2021. One Size Fits All? What Counts as Quality Practice in (Reflexive) Thematic Analysis? *Qualitative Research in Psychology* 18 (2021), 328–352. <https://www.tandfonline.com/doi/full/10.1080/14780887.2020.1769238>
- [18] Meredith Broussard. 2019. *Statement before the Task Force on Artificial Intelligence of the Committee on Financial Services, U.S. House of Representatives*.
- [19] Simone Browne. 2015. *Dark Matters: On the Surveillance of Blackness*. Duke University Press. <https://doi.org/10.1215/9780822375302>
- [20] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability, and transparency*. Proceedings of Machine Learning Research, 77–91.
- [21] Pittsburgh Coalition Against Predictive Policing. 2018. Predictive Policing in Pittsburgh: A Primer. (2018). [https://capp-pgh.com/files/Primer\\_v1.pdf](https://capp-pgh.com/files/Primer_v1.pdf)
- [22] Jennifer Cobbe, Michael Veale, and Jatinder Singh. 2023. Understanding Accountability in Algorithmic Supply Chains. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Chicago IL USA, 1186–1197. <https://doi.org/10.1145/3593013.3594073>
- [23] Lorraine Daston and Peter Louis Galison. 2007. *Objectivity*. Zone Books, New York.
- [24] Tom Davidson. 2020. Pittsburgh City Council Told Facial Recognition Bill Should Be Stronger. *TribLIVE* (Sept. 2020).
- [25] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO ’23)*. Association for Computing Machinery, New York, NY, USA, 1–23. <https://doi.org/10.1145/3617694.3623261>
- [26] Wesley Hanwen Deng, Nur Yildirim, Monica Chang, Motahare Eslami, Kenneth Holstein, and Michael Madaio. 2023. Investigating Practices and Opportunities for Cross-functional Collaboration around AI Fairness in Industry Practice. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM,



- Chicago IL USA, 705–716. <https://doi.org/10.1145/3593013.3594037>
- [27] Oxford English Dictionary. 2023. objective, adj., sense 1.3.b. <https://doi.org/10.1093/OED/3908623815>
- [28] Oxford English Dictionary. 2023. subjective, adj., sense 3. <https://doi.org/10.1093/OED/1098712908>
- [29] Kristie Dotson. 2014. Conceptualizing Epistemic Oppression. *Social Epistemology* 28, 2 (April 2014), 115–138. <https://doi.org/10.1080/02691728.2013.782585>
- [30] Catherine D'Ignazio. 2024. *Counting Feminicide: Data Feminism in Action*. MIT Press. Google-Books-ID: 2STOEAAAQBAJ.
- [31] Elizabeth Edenberg and Alexandra Wood. 2023. An Epistemic Lens on Algorithmic Fairness. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. ACM, Boston MA USA, 1–10. <https://doi.org/10.1145/3617694.3623248>
- [32] Diana E. Forsythe. 2001. *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence*. Stanford University Press, Stanford.
- [33] Paulo Freire. 1974. *Education for critical consciousness*. Continuum, London.
- [34] Miranda Fricker. 2007. *Epistemic injustice: power and the ethics of knowing*. Oxford University Press, Oxford; New York.
- [35] Seeta Peña Gangadharan and Seda Gürses. 2020. *Crafting CRAFT*. <https://craft-workshop.gitlab.io/>
- [36] Ben Gansky and Sean McDonald. 2022. CounterFAcCTual: How FAcCT Undermines Its Organizing Principles. In *Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 1982–1992. <https://doi.org/10.1145/3531146.3533241>
- [37] Timnit Gebru. 2021. Hierarchy of Knowledge in Machine Learning & Related Fields & Its Consequences. <https://scs.hosted.panopto.com/Panopto/Pages/Viewer.aspx?id=70f6edd7-de91-464e-ae94-acbb011ba2c7>
- [38] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for Datasets. *Commun. ACM* 64 (Dec. 2021), 86–92.
- [39] Google. [n. d.]. Google Cloud Model Cards. <https://modelcards.withgoogle.com/about>
- [40] Daniel Greene, Anna Lauren Hoffmann, and Luke Stark. 2019. Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial Intelligence and Machine Learning. *Hawaii International Conference on System Sciences 2019 (HICSS-52)* (Jan. 2019). [https://aisel.aisnet.org/hicss-52/dsm/critical\\_and\\_ethical\\_studies/2](https://aisel.aisnet.org/hicss-52/dsm/critical_and_ethical_studies/2)
- [41] Anne Gulland and Fayth Tan. 2023. Campus Surveillance: Students and Professors Decry Sensors in Buildings. *Nature* 623, 7986 (Oct. 2023), 444–445. <https://doi.org/10.1038/d41586-023-03287-w>
- [42] Eileen Guo and Tate Ryan-Mosley. 2023. Computer Scientists Designing the Future Can't Agree on What Privacy Means. <https://www.technologyreview.com/2023/04/03/1070665/cmu-university-privacy-battle-smart-building-sensors-mites/>
- [43] Sreesh Gururaja, Amanda Bertsch, Clara Na, David Widder, and Emma Strubell. 2023. To Build Our Future, We Must Know Our Past: Contextualizing Paradigm Shifts in Natural Language Processing. In *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Singapore, 13310–13325.
- [44] Seda Gürses, Rebekah Overdorf, and Ero Balsa. 2018. POTs: the revolution will not be optimized? *11th Hot Topics in Privacy Enhancing Technologies (HotPETs)* (2018), 2.
- [45] Leif Hancox-Li and I. Elizabeth Kumar. 2021. Epistemic Values in Feature Importance Methods: Lessons from Feminist Epistemology. In *Conference on Fairness, Accountability, and Transparency*. ACM, 817–826. <https://dl.acm.org/doi/10.1145/3442188.3445943>
- [46] Donna Haraway. 1988. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies* 14, 3 (1988), 575. <https://doi.org/10.2307/3178066>
- [47] Sandra Harding. 1992. Rethinking Standpoint Epistemology: What Is “Strong Objectivity?”. *The Centennial Review* 36, 3 (1992), 437–470.
- [48] Eun Seo Jo and Timnit Gebru. 2020. Lessons from archives: strategies for collecting sociocultural data in machine learning. In *Conference on Fairness, Accountability, and Transparency*. ACM, Barcelona Spain, 306–316. <https://doi.org/10.1145/3351095.3372829>
- [49] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (2019), 389–399.
- [50] Maximilian Kasy and Rediet Abebe. 2021. Fairness, Equality, and Power in Algorithmic Decision-Making. In *Conference on Fairness, Accountability, and Transparency*. ACM, Virtual Event Canada, 576–586. <https://doi.org/10.1145/3442188.3445919>
- [51] Os Keyes and Kathleen Creel. 2022. Artificial Knowing Otherwise. *Feminist Philosophy Quarterly* 8 (2022).
- [52] Youjin Kong. 2022. Are “Intersectionally Fair” AI Algorithms Really Fair to Women of Color? A Philosophical Analysis. In *Conference on Fairness, Accountability, and Transparency (FAcCT '22)*. Association for Computing Machinery, New York, NY, USA, 485–494. <https://doi.org/10.1145/3531146.3533114>
- [53] P. M. Krafft, Meg Young, Michael Katell, Jennifer E. Lee, Shankar Narayan, Micah Epstein, Dharma Dailey, Bernese Herman, Aaron Tam, Vivian Guetler, Corinne Bintz, Daniella Raz, Pa Ousman Jobe, Franziska Putz, Brian Robick, and Bissan Barghouti. 2021. An Action-Oriented AI Policy Toolkit for Technology Audits by Community Advocates and Activists. In *Conference on Fairness, Accountability, and Transparency (FAcCT '21)*. Association for Computing Machinery, New York, NY, USA, 772–781. <https://doi.org/10.1145/3442188.3445938>
- [54] Lilly Kubit. 2022. On CMU's Campus, Protesters March against Facial Recognition Tech. *PublicSource* (Sept. 2022).
- [55] Bogdan Kulynych, Rebekah Overdorf, Carmela Troncoso, and Seda Gürses. 2020. POTs: protective optimization technologies. In *Conference on Fairness, Accountability, and Transparency*. ACM, Barcelona Spain, 177–188. <https://doi.org/10.1145/3351095.3372853>
- [56] Diana Kwon. 2022. The rise of citational justice: how scholars are making references fairer. *Nature* 603, 7902 (March 2022), 568–571. <https://doi.org/10.1038/d41586-022-00793-1>
- [57] Bruno Latour. 1993. *We have never been modern*. Harvard University Press, Cambridge, Mass.
- [58] Benjamin Laufer, Sameer Jain, A. Feder Cooper, Jon Kleinberg, and Hoda Heidari. 2022. Four Years of FAcCT: A Reflexive, Mixed-Methods Analysis of Research Contributions, Shortcomings, and Future Prospects. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 401–426. <https://doi.org/10.1145/3531146.3533107>
- [59] Yvonna S Lincoln and Egon G Guba. 2000. Paradigmatic controversies, contradictions, and emerging confluences. *The Sage handbook of qualitative research* 4 (2000), 163–188.
- [60] Steve Lohr. 2018. Facial Recognition Is Accurate, if You're a White Guy. *The New York Times* (Feb. 2018). <https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html>
- [61] Audre Lorde. 1984. The Master's Tools Will Never Dismantle the Master's House. In *Sister Outsider: Essays and Speeches*. Crossing Press, Berkeley, CA, 110–114. [https://collectiveliberation.org/wp-content/uploads/2013/01/Lorde\\_The\\_Masters\\_Tools.pdf](https://collectiveliberation.org/wp-content/uploads/2013/01/Lorde_The_Masters_Tools.pdf)
- [62] Michael A Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-designing checklists to understand organizational challenges and opportunities around fairness in ai. In *Conference on Human Factors in Computing Systems*. ACM, 1–14.
- [63] Sophie Maddocks. 2020. 'A Deepfake Porn Plot Intended to Silence Me': exploring continuities between pornographic and 'political' deep fakes. *Porn Studies* 7, 4 (2020), 415–423.
- [64] Tambiama Madiaga. 2023. *Artificial intelligence act*.
- [65] Maya Malik and Momin M. Malik. 2021. Critical Technical Awakenings. *Journal of Social Computing* 2, 4 (Dec. 2021), 365–384. <https://doi.org/10.23919/JSC.2021.0035>
- [66] Sean McGregor. 2021. Preventing Repeated Real World AI Failures by Cataloging Incidents: The AI Incident Database. *AAAI Conference on Artificial Intelligence* 35, 1717 (May 2021), 15458–15463. <https://doi.org/10.1609/aaai.v35i17.17817>
- [67] Jacob Metcalf, Emanuel Moss, and danah boyd. 2019. Owning Ethics: Corporate Logics, Silicon Valley, and the Institutionalization of Ethics. *Social Research: An International Quarterly* 86, 2 (2019), 29.
- [68] Cade Metz and Daisuke Wakabayashi. 2020. Google Researcher Says She Was Fired Over Paper Highlighting Bias in A.I. *The New York Times* (Dec. 2020). <https://www.nytimes.com/2020/12/03/technology/google-researcher-timnit-gebru.html>
- [69] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Conference on Fairness, Accountability, and Transparency*. ACM, Atlanta GA USA, 220–229. <https://doi.org/10.1145/3287560.3287596>
- [70] Brent Mittelstadt. 2019. Principles Alone Cannot Guarantee Ethical AI. *Nature Machine Intelligence* 1, 11 (Nov. 2019), 501–507. <https://doi.org/10.1038/s42256-019-0114-4>
- [71] Dawn Nafus. 2012. 'Patches don't have gender': What is not open in open source software. *New Media & Society* 14, 4 (2012), 669–683.
- [72] Dawn Nafus and Jamie Sherman. 2014. This One Does Not Go Up to 11: The Quantified Self Movement as an Alternative Big Data Practice. (2014), 11.
- [73] Arvind Narayanan. 2022. The Limits of the Quantitative Approach to Discrimination. (2022). <https://www.cs.princeton.edu/~arvindn/talks/baldwin-discrimination/baldwin-discrimination-transcript.pdf>
- [74] NIST. 2022. AI Risk Management Framework: Second Draft - August 18, 2022. (2022).
- [75] Lester C. Olson. 2000. The Personal, the Political, and Others: Audre Lorde Denouncing “The Second Sex Conference”. *Philosophy & Rhetoric* 33, 3 (2000), 259–285.
- [76] Gwen Ottinger, Kelly Bronson, and Dawn Nafus. 2023. Clicks and particulates: Value, alienation, and attunement as unifying themes in big data studies. *Big Data & Society* 10, 1 (Jan. 2023), 20539517231184891. <https://doi.org/10.1177/20539517231184891>

- [77] Lawrence A. Palinkas, Sarah M. Horwitz, Carla A. Green, Jennifer P. Wisdom, Naihua Duan, and Kimberly Hoagwood. 2015. Purposeful Sampling for Qualitative Data Collection and Analysis in Mixed Method Implementation Research. *Administration and Policy in Mental Health and Mental Health Services Research* 42, 5 (Sept. 2015), 533–544. <https://doi.org/10.1007/s10488-013-0528-y>
- [78] Marie-Therese Png. 2022. At the Tensions of South and North: Critical Roles of Global South Stakeholders in AI Governance. In *Conference on Fairness, Accountability, and Transparency (FAcCT '22)*. Association for Computing Machinery, New York, NY, USA, 1434–1445. <https://doi.org/10.1145/3531146.3533200>
- [79] Hortense Powdermaker. 1966. *Stranger and Friend: The Way of an Anthropologist*. W. W. Norton & Company.
- [80] Dasha Pruss. 2023. Ghosting the Machine: Judicial Resistance to a Recidivism Risk Assessment Instrument. In *ACM Conference on Fairness, Accountability, and Transparency (FAcCT '23)*. Association for Computing Machinery, New York, NY, USA, 312–323. <https://doi.org/10.1145/3593013.3593999>
- [81] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. The Fallacy of AI Functionality. In *Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 959–972. <https://doi.org/10.1145/3531146.3533158>
- [82] Inioluwa Deborah Raji, Morgan Klaus Scheuerman, and Razvan Amironesei. 2021. You Can't Sit With Us: Exclusionary Pedagogy in AI Ethics Education. In *Conference on Fairness, Accountability, and Transparency*. ACM, Virtual Event Canada, 515–525. <https://doi.org/10.1145/3442188.3445914>
- [83] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *Conference on Human Factors in Computing Systems*. ACM, 1–15.
- [84] Princess Sampson, Ro Encarnacion, and Danaë Metaxa. 2023. Representation, Self-Determination, and Refusal: Queer People's Experiences with Targeted Advertising. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Chicago IL USA, 1711–1722. <https://doi.org/10.1145/3593013.3594110>
- [85] Zoe Schiffer and Casey Newton. 2023. Microsoft lays off team that taught employees how to make AI tools responsibly. <https://www.theverge.com/2023/3/13/23638823/microsoft-ethics-society-team-responsible-ai-layoffs>
- [86] Thomas Schwandt. 1994. Constructivist, Interpretivist Approaches to Human Inquiry. *Handbook of Qualitative Research Thousand Oaks, California: Sage* (1994).
- [87] Irving Seidman. 2006. *Interviewing as Qualitative Research: A Guide for Researchers in Education and the Social Sciences* (3rd ed.). Teachers College Press.
- [88] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2022. Participation Is not a Design Fix for Machine Learning. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. ACM, Arlington VA USA, 1–6. <https://doi.org/10.1145/3551624.3555285>
- [89] Lucy Suchman. 2002. Located accountabilities in technology production. *Scandinavian Journal of Information Systems* 14 (2002), 16. Issue 2.
- [90] Lucy Suchman. 2023. The uncontroversial “thingness” of AI. *Big Data & Society* 10, 2 (July 2023). <https://doi.org/10.1177/20539517231206794>
- [91] Harini Suresh, Rajiv Movva, Amelia Lee Dogan, Rahul Bhargava, Isadora Cruzen, Angeles Martinez Cuba, Guilia Taurino, Wonyoung So, and Catherine D'Ignazio. 2022. Towards Intersectional Feminist and Participatory ML: A Case Study in Supporting Femicide Counterdata Collection. In *Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 667–678. <https://doi.org/10.1145/3531146.3533132>
- [92] Gemini Team. 2023. *Gemini: A Family of Highly Capable Multimodal Models*. [https://storage.googleapis.com/deepmind-media/gemini/gemini\\_1\\_report.pdf](https://storage.googleapis.com/deepmind-media/gemini/gemini_1_report.pdf)
- [93] Helen Verran. 2013. *Engagements between disparate knowledge traditions: Toward doing difference generatively and in good faith*. HSRC Press, Cape Town, South Africa, 141–161.
- [94] Sarah Myers West, Meredith Whittaker, and Kate Crawford. 2019. *Discriminating Systems: Gender, Race, and Power in AI*. AI Now Institute. <https://ainowinstitute.org/publication/discriminating-systems-gender-race-and-power-in-ai-2>
- [95] Meredith Whittaker. 2023. Origin Stories: Plantations, Computers, and Industrial Control. *Logic(s) Magazine* 19-supra dupa skies (move slow and heal things) (May 2023). <https://logicmag.io/supa-dupa-skies/origin-stories-plantations-computers-and-industrial-control/>
- [96] David Gray Widder, Laura Dabbish, James D. Herbsleb, Alexandra Holloway, and Scott Davidoff. 2021. Trust in Collaborative Automation in High Stakes Software Engineering Work: A Case Study at NASA. In *Conference on Human Factors in Computing Systems*. ACM, 1–13. <https://dl.acm.org/doi/10.1145/3411764.3445650>
- [97] David Gray Widder and Claire Le Goues. 2024. What Is a “Bug”? On Subjectivity, Epistemic Power, and Implications for Software Research. <https://doi.org/10.48550/arXiv.2402.08165> arXiv:2402.08165 [cs]
- [98] David Gray Widder and Dawn Nafus. 2023. Dislocated accountabilities in the “AI supply chain”: Modularity and developers' notions of responsibility. *SAGE Big Data & Society* 10, 1 (Jan. 2023), 20539517231177620. <https://doi.org/10.1177/20539517231177620>
- [99] David Gray Widder, Derrick Zhen, Laura Dabbish, and James Herbsleb. 2023. It's about power: What ethical concerns do software engineers have, and what do they (feel they can) do about them?. In *Conference on Fairness, Accountability, and Transparency*. ACM, Chicago IL USA, 467–479. <https://doi.org/10.1145/3593013.3594012>
- [100] Amanda Williams, Silvia Lindtner, Ken Anderson, and Paul Dourish. 2014. Multitised Design: An Analytical Lens for Transnational HCI. *Human-Computer Interaction* 29, 1 (Jan. 2014), 78–108. <https://doi.org/10.1080/07370024.2013.823819>
- [101] Michelle Yi. 2021. Model Cards: A Nutrition Label for Predictive Models. <https://www.michelleyi.ai/blog/model-cards-a-nutrition-label-for-predictive-models>
- [102] Meg Young, Michael Katell, and P.M. Krafft. 2022. Confronting Power and Corporate Capture at the FAcCT Conference. In *Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 1375–1386. <https://doi.org/10.1145/3531146.3533194>