

Blooming in Muddy Waters

DEI at AI Ethics Conferences

by Emnet Tafesse

**DATA &
SOCIETY**

with Meg Young,
Ranjit Singh, and
Jacob Metcalf



Blooming in Muddy Waters: DEI at AI Ethics Conferences

by **Emnet Tafesse**
with Meg Young,
Ranjit Singh, and
Jacob Metcalf

Contents

01 Executive Summary

03 Introduction

08 Author Note

09 Situating AI Ethics in DEI Work

12 Improving DEI in AI Ethics

14 Taking Stock of DEI Now and Moving Forward

14 Disciplinary and Practitioner Diversity

18 Affinity Groups

20 Increasing Speaker and Organizer Diversity

23 Travel Support and Other Scholarships

25 Conference Locations and Geographic Diversity

28 Conclusion

30 Acknowledgments

31 Appendix: Methods

33 Endnotes

Executive Summary

This report documents challenges faced by BIPOC (Black, indigenous, and people of color) scholars to participate in AI ethics conferences and the need for more equitable and intentional spaces. Despite foundational contributions, BIPOC scholars and the larger community remain underrepresented. Many people I spoke with report feeling disconnected in these conference spaces, reflecting broader disparities within the field as a whole.

Efforts to promote diversity, equity, and inclusion (DEI) in AI ethics conferences to date range from formal support for affinity groups, travel grants, speaker and organizer diversity, and geographic diversity, to multidisciplinary. However, these do not address deeper challenges like tokenization, in which individuals from minoritized groups are forced to stand as universal representatives rather than being recognized for their individual expertise.

A second key challenge in conference spaces is that attendees often find their identity as BIPOC researchers leads to assumptions about their knowledge and experiences as members of affected communities, undermining their multidimensional identities and scholarly contributions. This collapse between identity and representation perpetuates a dynamic that flattens essential insights in discussions on technology development, detracting from the very purpose of these conferences. This report identifies specific opportunities for conference organizers to better champion the work of BIPOC attendees and create more welcoming spaces.

Recommendations

Respondents called on conference organizers to shift peer review processes to better engage other communities of practice, including transnational and non-Western scholarship. They highlighted the value of practitioner and civil society perspectives in discussions of algorithmic justice. They also called for stronger connectivity to communities where conferences are hosted by spotlighting local campaigns and organizations.

Conference organizers have the opportunity to improve affinity group support and take steps to enhance experience of BIPOC through: involving affinity group organizers in curating main conference programs, spreading out affinity groups throughout the schedule for broader engagement, holding affinity group events concurrently with the main conference, and promoting these events to all attendees to highlight BIPOC scholars. These needs will be further addressed by additional social infrastructure for newcomers to the community.

Conference organizers should work to improve speaker and organizer diversity by providing BIPOC scholars with more decision-making power while being mindful of the additional burdens put on the community through service roles. Conference organizers should also include moderator training to improve overall intentionality of the space, and focus more programming on the trailblazing work of BIPOC scholars.

Travel support and inclusive planning practices are an essential component of supporting more equitable conference spaces. Conference organizers have the opportunity to improve existing practices by providing a longer duration between acceptance and the conference to support visa applications. Organizers should continue to invest in DEI work and communicate projected timelines for notification and reimbursement transparently, with the long term goal of shifting away from reimbursement norms in academia, which require individuals to outlay their own funds and potentially wait months to be reimbursed.

Conference organizers must also be intentional about choosing conference locations. Choose host countries with fewer visa barriers. Consider bridging transnational gaps as a long-term investment in relationships with a wider set of literatures, scholars, institutions, and organizations—these take time and effort to build. In every conference location, create opportunities for local attendees and lower barriers to local engagement.

This work was supported by the inaugural DEI Scholars grant from the ACM FAccT Conference.

Introduction

“Remember to imagine and craft the worlds you cannot live without just as you dismantle the ones you cannot live within.”

—Ruha Benjamin

During my time as a graduate student studying public policy, I faced the stark realities of academic spaces. From a curriculum that ignored systemic issues to daily microaggressions, it was clear that my BIPOC peers and I experienced barriers that limited our ability to fully participate in, to simply be in, an academic environment.¹ I saw the work and expertise of my peers and BIPOC professors was often overlooked and discounted. In the midst of these experiences and efforts to build solidarity, we attempted to carve out spaces of support: creating student organizations that challenged the mostly white academic spaces we were in; creating spaces that were by us and for us.

The very first academic conference I attended was Data for Black Lives in 2018, founded by technologist and activist Yeshimabeit Milner. The conference felt like an antithesis to the hurt and disempowerment of my experiences in traditional academic spaces—a breath of fresh air that centered on Black scholarship and community. I remember a wholesome space of deep alignment, laughter, good food, and opportunities to create connections. I made friendships that have lasted through the years. I learned how data and research can harm our communities and, at the same time, ways we can utilize research, advocacy, and activism to undo systemic forms of oppression. That event, and my experiences in many other spaces in the years since (ones that at times were far less welcoming) gave me a personal mission to understand the beauty of spaces that BIPOC communities carve out—how we bloom in the hardest of environments—and how I can continue to follow the lead of Black women scholars whose contributions have been deeply consequential for the field of AI ethics.

This report is a product of my initial steps toward this mission. It presents the experiences of BIPOC attendees at three major AI ethics conferences as they relate to DEI efforts, that is, the policies and practices each conference used to promote diversity, equity, and inclusion for attendees, organizers, and presenters. My participant observation at conferences is buttressed by interviews with conference organizers about the challenges of addressing structural inequalities. It is the result of a year-long study based on surveys, interviews, observation, and conversations in and about AI ethics² academic conference spaces. It focuses on AI ethics conferences at a key moment in their evolution from being about matters of DEI in the development of data and technology, to playing an active role in fostering these values where knowledge production happens.

For the field of AI ethics, and the professional organizations that are growing up around it, issues of DEI are fundamentally important. AI ethics is generally understood as an extension of the computer science domain, which has historically lacked gender and racial diversity.³ According to the Pew Research Center analysis of the 2017–19 American Community Survey of employed adults ages 25 and older, of the population of people working in STEM, Black people make up 7%, Latinx people make up 8%, and Asian people make up 20% of computer scientists.⁴ Only 2.4% of new PhDs in AI in 2021 from US universities were awarded to Black students.⁵

However, numerous pivotal theoretical and empirical contributions to AI ethics have been made by women of color, as well as queer and trans people of color. Pioneering work by Professor Latanya Sweeney on metadata inferences inspired others to look at discrimination in data mining and classification.⁶ Joy Buolamwini and Timnit Gebru’s “Gender Shades” paper demonstrated that accuracy in facial recognition varies with race and gender, further highlighting the need for Black feminist theories like intersectionality to be more fully accounted for in computing and establishing new methodological grounds for AI fairness research.⁷

This work was supported by the ACM Fairness, Accountability, and Transparency (FAccT) Conference’s 2022 DEI Scholars award, which was founded to support the organization’s strategic goal to “innovate and lead on diversity and inclusion.” The ACM FAccT DEI Scholars program seeks to fund scholars to conduct research that deepens our collective understanding of and advances DEI initiatives. The research conducted by DEI scholars is independent and neither ACM nor FAccT has any privileged access to the data collected.

This project focuses on the experiences of BIPOC people along the axes of race and ethnicity only. We acknowledge that “BIPOC” is neither a homogeneous nor an uncontested category of personhood, and all members of this community have other intersectional identities that shape their experiences. The research participants interviewed and surveyed were asked to self-identify with this category. We focus on race and ethnicity here because of its centrality to the AI ethics field and because of practical constraints on the scope of the project; we anticipate others who receive this particular award will conduct complementary research in order to address related issues such as caste, gender, disability, and migration. In presenting our recommendations, our hope is that many of our proposed best practices may have beneficial effects for DEIA efforts beyond race and ethnicity. In addition, we default to the acronym “DEI” in this report to follow the naming convention used by the conferences, but hope that those efforts also grow to include accessibility.

Abeba Birhane and Vinay Prabhu traced the source of racist and misogynistic bias to image labeling practices.⁸ Inioluwa Deborah Raji's research agenda has charted the course on methods for algorithmic auditing and oversight.⁹ These are just a few examples of scholarship among many that illuminate how algorithms reflect and entrench power asymmetries in society. But even as BIPOC scholars have substantially set the field's agenda, they remain underrepresented among participants in its conference spaces.

While the individual paths that scholars take to their research agendas vary, their life experiences can inform their work in ways that are generative for the field more broadly. This idea is theorized in feminist epistemology by Sandra Harding as standpoint theory, that is, that the vantage point a person has in the world is important and gives them unique insight as it relates to their lived experiences.¹⁰ Black Feminist theorist Patricia Hill Collins describes how centering Black women's experiences opens novel conceptual terrain, especially to investigate the operation of power.¹¹ Collins argues that multiple unique vantage points are necessary to get a fuller vision of a problem, and that power asymmetries will make it more difficult for some groups of people to be heard. This resonates with my observations in the field.

For this report, I attended and was a participant observer at the 2022 edition of three conferences: the Conference on Fairness, Accountability, and Transparency (FAccT); Artificial Intelligence, Ethics, and Society (AIES); and the Neural Information Processing Systems Conference (NeurIPS)—all organized within computer science professional organizations, namely the Association for Computing Machinery (ACM), Association for the Advancement of Artificial Intelligence (AAAI), the Institute of Electrical and Electronics Engineers and the NeurIPS Foundation. While my data collection was anchored around the 2022 editions of each conference, I also returned to the 2023 FAccT conference to hold formal and informal conversations about BIPOC people's experiences in AI ethics spaces, where I was able to get another check on the validity of my findings.¹²

While each of these conference communities is distinct in their history, membership, and practices, they share common features: they are all organized within computing professional associations. Of the three, NeurIPS is the oldest and by far the largest by attendance, but all three conferences include contributions to machine learning research and the conceptual and cultural inheritances of computer science. Where FAccT and AIES are venues exclusive to AI ethics and often assert fairness, transparency, and accountability as core values, Birhane, et al. note that the traditional tenets of machine learning research are in practice far more often commitments to efficiency, performance, generalization, and quantitative evidence.¹³ These values, and the epistemological commitments and norms of computer science, vie with other values and ways of knowing in AI ethics research.

Computer science is also a starting point for some of the systemic challenges faced in computing research spaces for women and BIPOC people. From 2015–2019, a number of documented incidents of sexism, hostility, and sexual harassment at NeurIPS led to a number of changes to the conference's name, policy, and intentionality around inclusion.¹⁴ For example, affinity groups like Black in AI played a key role in changing how attendees experienced NeurIPS as a space. When asked why to stay and change a space rather than forge a new one, one person told me: "For NeurIPS, I think it's really important to change it because ... this is a space where all the professors for that field go ... that's where all the students go, that's where you have to present your work to graduate, that's where people start collaborations, that's where all the companies go to recruit" (interview participant from NeurIPS and FAccT, June 5, 2023). Now in its 37th year, NeurIPS is a premier venue for machine learning research. While it is not primarily an AI

ethics conference, it has become a site for extensive official workshops on AI ethics and policy topics, and a meeting place for affinity groups that I observed and whose members I interviewed. NeurIPS is also regularly a site of controversy over the social and political dynamics—especially the diversity—of the computing and AI professions.

The other two conferences in my study were created in part based on a desire for publication venues that would be recognized within computer science for work related to bias and ethics. At the time of this writing, both AIES and FAccT are in their sixth year. Both produce published proceedings of empirical and professionally recognized research about the social and ethical consequences of AI technologies that were not previously accepted in venues like NeurIPS and other AI conferences. One respondent described FAccT as a relief from the unwelcoming atmosphere that work like hers had found at computer science conferences.

There are ongoing efforts to encourage diversity and inclusion at all three of these conferences, which vary in the resources allocated to them, visibility, and success.¹⁵ These efforts are understandably focused on increasing the number of BIPOC attendees at the conferences and addressing problems of representation. Notably, many people I spoke with shared positive aspects of their conference experiences and rewarding aspects of these spaces. Each conference is also unique in its evolution over time, strengths, and opportunities to improve.

At the same time, many of the same interviewees expressed ambivalence that, in their experience, the stated values of AI ethics conferences were disconnected from deeper commitments to equity and inclusion in practice. Throughout my fieldwork, I found instances in which AI ethics conferences were at times unwelcoming spaces for BIPOC attendees. These feelings of discomfort were tied to sidelining by other attendees and organizers, assumptions about BIPOC colleagues' identities and expertise, exclusion in micro and macro conference facilitation and planning, and rigidity around what forms of knowledge production and expertise were valorized in AI ethics. The result is that AI ethics conferences are sometimes spaces where much of the discussion is about the experiences of BIPOC people, but do not go far enough to disrupt the dynamic of exclusion¹⁶ characterizing the societal problems that motivate the field.

I offer two major, inter-related challenges that span the experiences of BIPOC scholars I spoke with across all three conferences: 1) the tokenization of BIPOC people inside of majority-white spaces, that is, a symbolic or superficial inclusion without meaningful structural change; and 2) a conceptual collapse between people's identity as BIPOC scholars and the perception of them as representatives of impacted communities. These two challenges can't be addressed by increased representation alone, and in fact increased representation without deeper changes may risk exacerbating these dynamics.

Tokenization is the process of forcing one member of a minoritized group to stand in as a symbolic representative, or "token," for all members of the group, especially in spaces that do not strive for deeper structural changes. One survey respondent warned against "allyship theatrics" and a tendency to "position BIPOC individuals as the one and only representative ... for all BIPOC everywhere (survey respondent #52 on AIES)." Many attendees I spoke with felt isolated at these conferences and singled out at times to represent BIPOC people more broadly. This sense of surface-level inclusion was palpable in experiences some attendees shared of being ignored, having their posters overlooked, or experiencing micro- and macro-aggressions. As one attendee stated, "Only bringing [BIPOC] people here doesn't solve the problem. ... White people who are here also should change their mindset—that needs to change. It's not only [about] providing grants and bringing the Black people here to increase diversity" (interview

participant on NeurIPS, November 30th, 2022). Without the deeper work of addressing these behaviors, surface-level inclusion is hollow.

Furthermore, many attendees felt their identity as BIPOC researchers was conflated by other attendees as also representing “affected” or “vulnerable” communities, regardless of their scholarly expertise. All people occupy multiple, intersectional identities, and a collapse occurs where one feature is assumed to imply other features, expertise or experiences. Occupying a BIPOC identity inside of North American context does not necessarily imply that one is underprivileged or vulnerable, or that one is an expert on race in computing.¹⁷ For example, one respondent reported a panel in which a moderator only asked the BIPOC panelists questions related to race. One attendee who was present later explained to me: “If someone asked me that question ... the first thing that would come out of my mouth [would be that] I’m not an intersectionality expert. And you should treat it as a discipline ... I think the short of it is that there was an unfair amount of pressure on women of color to talk about intersectionality and devaluing ... the work that they’ve done, because they were just viewed as the intersectionality [expert].” (Interview participant on FAccT, July, 27th, 2022). Here, the respondent speaks to the way that racial identity is not coextensive with expertise on race and that BIPOC scholars’ research agenda and expertise can be devalued when enlisted narrowly for this purpose. Put another way, even as BIPOC researchers’ experiences might inform a research agenda, their identities must not be conflated with their scholarly expertise.

While the commitment of AI ethics spaces to the issues of DEI have contributed to recognizable progress in the field, thornier and culturally ingrained barriers to equity remain.

I anchor my analysis of DEI primarily within efforts related to race and ethnicity. However, DEI is not limited to race and ethnicity; the findings of this report should not be read as a comprehensive guide for DEI at AI ethics conferences. On the contrary, a key insight of Kimberle Crenshaw’s theory of intersectionality is that anti-oppression work that does not account for multiple intersections of marginality fails to account for the ways power operates. For conferences, this implies that a commitment to building an inclusive community requires multiple, overlapping efforts to address the ways in which people’s identities correspond with a lack of access—for example, inviting a diversity of speakers means little if they cannot afford caretakers for their dependents back home. By designing an environment that addresses the needs of people from a number of faiths or genders, those who are immuno-compromised or neurodivergent; or those with financial barriers, chronic illness, or mobility challenges, conference organizers have the potential to foster inclusive environments across the contours of race and ethnicity as well.

Ultimately, I argue that the current DEI efforts of AI ethics conferences must go further to act on the same social forces that precipitated the need for the field by more conscientious action within the conference space. This report distills pragmatic changes to conference DEI efforts that organizers can take related to affinity groups, speaker/organizer invitations, scholarships and travel funding, remote conference locations,

and disciplinary diversity that should all be considered in light of the real barriers to participation faced by BIPOC scholars. While the commitment of AI ethics spaces to the issues of DEI have contributed to recognizable progress in the field, thornier and culturally ingrained barriers to equity remain.

To address deeper challenges like tokenism and misperceptions of BIPOC colleagues' experiences, AI ethics conference spaces must commit to forging spaces with a wider range of perspectives, including with funding, programming, and organizing. These changes will better acknowledge each person's individual expertise, experiences, and contributions, and support the field to acknowledge and further act on AI ethics problems at their source.

Author Note

A note about authorship: I (Emnet Tafesse, the first author) conducted all primary research for this project, including interviews, field observations, and collecting survey data. My research was supported by the ACM FAccT 2022 DEI Fellowship. I am a Black, queer, fem, first-generation Ethiopian-American woman born and raised in Seattle, Washington, and now residing on the East Coast of the United States. My identity played a significant role in the connections I made for this research, my ability to solicit frank feedback from other BIPOC scholars, and in my own firsthand experiences of being in AI ethics spaces.

When I write “we” in this report, I refer to my coauthors and teammates on the AI on the Ground team at Data & Society—three other AI ethics scholars from various backgrounds and disciplines who have been instrumental in helping me to parse the large amount of data I collected and to position this work within the organizational context and history of these conferences. I have valued sharing this work, even as I have followed my own sense of where it is important for it to end up. For a more in-depth acknowledgment of how we collaborated, see the Methods appendix at the end of this report.

Situating AI Ethics in DEI Work

Diversity, equity, and inclusion (sometimes more broadly constituted to include accessibility) is a phrase that organizations use to describe the set of policies and practices that form its response to demands for change. Race and ethnic studies scholar Nick Mitchell defines it as a “wide-ranging and often contradictory set of legally regulated and organizationally mediated efforts to reform public and private institutions ... to undo or interrupt historically sedimented, socially entrenched, and culturally normalized practices of discrimination and exclusion,” especially as taken up by elite institutions. Mitchell notes that the rise of DEI as a core method for addressing racism has consequences for what society recognizes as constituting racism and how to respond to it—because DEI takes place within institutional contexts it renders some very real aspects of racism as out-of-scope because they are not tractable to institutional solutions.¹⁸ So, too, do DEI practices vary and are adopted based on an organization’s own discretion. As a result, the purpose of DEI efforts is contested: where proponents interpret them to be efforts for institutional reform, critics highlight their history as arising from within elite white institutions as a strategy for self-preservation and appeasement.¹⁹

Evidence for DEI initiatives as reproducing power relations traces back to their inception. In the United States, Mitchell describes the way US institutions adopted commitments to diversity during the Cold War as an effort to signal to the international community that prosperity is not bounded by race or gender. Meanwhile, as mass movement mobilization led to the passage of the US Civil Rights Act, activists pressured universities as to how the Act’s “nondiscrimination” mandate would be defined in practice—particularly in hiring and admissions. DEI efforts were the eventual result of university action as they scrambled to prevent greater government regulation, defining nondiscrimination on their own terms. While universities took steps to admit additional minority students, student movements asked them to confront their own role in supporting and sustaining imperialist war, neo-colonial exploitation, and epistemologies of domination. American Studies scholar Roderick A. Ferguson traces how student movements of the 1960s were capitalized on and subsumed by universities, which produced “minority difference” as a means to extend their power, which later contributed to the institutional formality of DEI efforts.²⁰ Feminist scholar Sara Ahmed argues that DEI efforts ultimately do more to intervene on the perception of the institution rather than the institution itself; she explores for instance how a diversity policy is produced as evidence that an organization is achieving its goal in the absence of substantive change.²¹

The relationship of DEI efforts to AI ethics is complex. Race and other identity categories are foregrounded in AI ethics as a field; however, the field has also oversimplified race and other identities as computable, when, in reality, they are fluid and contextual.²² Many harmful consequences of AI systems for communities of color are downstream of such misapprehensions about race. Among these is the idea that race is

extricable from other variables in a dataset, when according to philosopher Lily Hu, these effects on data compose what race is and how it operates.²³ Similarly, critical data science scholar Ruha Benjamin understands race as itself a technology: race is a tool for ordering social relations, which treats people differently based on assigned categories. Race operates as both an input to technologies with racialized outcomes, and an output as disparate impacts of automated decisions.²⁴ Ethnographer Sareeta Amrute describes race in the tech industry as akin to an algorithm that governs repeatable behavior, continuously and simultaneously providing problems to address and methods to address them.²⁵

In Harrington et al. (2019), human-computer interaction scholars have argued that harmful BIPOC experiences are part of what propels the field forward, in the sense that they provide fodder for research by scholars who are white or from elite institutions to publish on their own behalf at an “objective” distance without supporting deeper modes of justice in their research practices.²⁶ Multiple interview and survey respondents in this study also pointed to this dynamic, as I share below. In AI ethics, race often serves as an analytic category that generates new measures of algorithmic harm, methods for addressing them, and opportunities for demanding changes to the tech industry. Racism—and therefore also diversity, equity, and inclusion—is a foundational topic of research for the field of AI ethics, and it should be a central, explicit concern as we explore how to build and sustain professional venues so the field does not replicate those very dynamics it seeks to disrupt.

AI ethics as a subfield of computer science is also shaped by an insider–outsider dynamic in which AI ethics researchers necessarily (and sometimes awkwardly) intervene on methods in computer science with questions about fairness. The ACM FAccT conference exists in part because academics in computer sciences and related fields were dissatisfied with the lack of attention to urgent social and ethical questions in mainstream computing publications and conferences. Even now, when issues of AI ethics and policy are regularly front page news, the top-line AI and data science disciplinary conferences, such as NeurIPS and AAAI, tend to relegate ethics scholarship to side program workshops, rather than treating them as a part of the main research tracks. This insider–outsider dynamic is evocative of the experiences of many BIPOC professionals who operate in traditionally white professional spaces.²⁷ This dynamic has complicated connections with matters of diversity and inclusion—one might reasonably expect that AI ethics conferences would be more diverse than mainstream computer science conferences, but that has not necessarily been the case.²⁸

Historically, AI ethics is not unique in facing these problems. For example, anthropology has long discussed the centrality of race and ethnicity to both its epistemic practices and professional conventions. Like AI ethics, anthropology has faced the dynamic of being about racial and ethnic groups without those groups setting the terms of the discussion. With its historical roots in European colonialism, initial anthropological research was concerned with the management of cultural differences in pursuit of exploitative colonial projects.²⁹ Over the past five decades, the American Anthropological Association has made the status and experience of BIPOC anthropologists (typically called “minority anthropologists”) a subject of multiple reports and surveys, particularly focused on how professional conventions and exclusive mentorship practices prevent BIPOC anthropology students and faculty from fully participating in the profession.³⁰ One of the conclusions of these studies was that racist professional practices ultimately limit the topics and methods considered appropriate to the discipline, and recommended multiple practical steps to make the field more inclusive during graduate study and professionalization.

Computer science has also struggled with diversity and inclusion in professional and leadership roles and educational pipelines.³¹ Extensive evidence demonstrates that demographic and social diversity improves scientific progress and technological innovation;³² yet despite the higher rates of innovation among early-career researchers from minoritized backgrounds, they overall have worse career outcomes.³³ STEM fields have consistently struggled with establishing a diverse workforce and educational pipeline, and computer science typically lags other fields.³⁴

The US National Science Foundation has emphasized the importance of integrating DEIA (diversity, equity, inclusion, and accessibility) efforts across undergraduate education, graduate training and early research experiences, and industry and government; professional conferences are one of the places where all of those groups come into contact.³⁵ A study by Acuna and Liang (2021) examined the relative diversity in AI ethics and computer science conferences and found that they were relatively similar with regard to ethnic and racial diversity, but AI ethics conferences were closer to gender parity and were characterized by far fewer contributors from China.³⁶ BIPOC computer science scholars have noted the significant toll of racism in their careers and personal lives, especially in the context of the recent and ongoing protest movements for racial justice in the US.³⁷

These prior conversations set the stage for considering how AI ethics conferences can address DEI efforts within their own spaces. Computer science, like most STEM fields, has also struggled with recruiting and retaining members of racialized groups. In order to make advances on core scientific and theoretical questions of AI/ML fairness and ethics, the professional spaces in which this work occurs necessarily must be welcoming and rewarding for everyone, especially members of BIPOC communities.

Improving DEI in AI Ethics

The findings below describe the experiences of BIPOC respondents in the 2022 editions of the AIES, FAccT, and NeurIPS conferences (which I also attended in person). At the time of this study, each of these conferences had renewed their commitment to diversity, equity, and inclusion in part by collecting more detailed demographic data about the identities of attendees; the composition of DEI efforts across conferences varied in terms of available resources, precedent within the community, and the commitments of organizers. In the following section, we describe the work involved in producing DEI at AI ethics conferences and some key experiences shared by respondents. We close each section with pragmatic guidance to AI ethics conference organizers and attendees.

DEI must ultimately be about community building and intentionality—who builds the space, who is welcomed into it, and who has space in it. DEI work at academic conferences can encompass a number of efforts. As a scholar who has previously co-organized FAccT conference put it:

I have been involved in organizing the conference [in previous years], because it feels like a home conference. I have been trying to make it as diverse as possible, make sure that it's welcoming and inclusive of Black folks, folks from the African continent, through various efforts such as disseminating the call for papers. I've been making those kinds of efforts to make it my conference—to make it feel like home. But on the other hand, it does feel like I don't belong there because there are very few people that look like me. ... Sometimes it feels like it's another world (interview participant on FAccT, March 14, 2023).

The primary DEI efforts of AI ethics conferences are currently programming diverse speakers, fostering intellectual and topical diversity in the main program, providing targeted scholarships and travel grants, establishing clear codes of conduct, or hosting affinity groups. Some of the responsibilities for DEI are diffused across conference leadership roles, while others are administered by designated DEI chairs for that year's conference. The composition of DEI efforts can also change depending on the resources available, edition of the conference, or individual priorities of chairs. In interviews, respondents to this study identified childcare support, wheelchair access, pronoun tags, land acknowledgments, and dietary options as additional types of inclusion work. Other efforts that were not mentioned could include language interpretation, a scent-free environment, rooms for religious observance, and seating options for size diversity, people who are hard of hearing, or who are mobility impaired.

Organizers cited resource constraints as affecting the scope of DEI work, including the limitations of a given conference location, volunteer capacity, and budget. These dynamics highlight a consistent challenge in accomplishing DEI efforts at conferences. While there is an organizational investment in DEI, it is treated as programmatically separate from the core scholarly activities of conferences. There is a declared dedication to DEI on paper that is easy and cost-free, but implementing these commitments is much more fraught and challenging in practice.

DEI must ultimately be about community building and intentionality

Participants also shared an expectation that conferences like FAccT and AIES be held to their stated values of fairness, accountability, transparency, and ethics in the practice and organizing of the conference itself. One person shared a sense of disappointment at harmful interpersonal experiences at FAccT to a greater degree than at other conferences: “FAccT was the most shocking one because it is supposed to be addressing fairness and accountability, and I had the most racist interactions I’ve had at that conference ... I would expect this from NeurIPS because nobody’s claiming that that’s about fairness and transparency, but when these people write about these things and act like this, it’s even more infuriating ... because you’re benefiting from that image and from that kind of work. You get your tenure, you publish your papers—and this is what you’re doing, you know?... Because what they think is they’re the good ones.” (Interview participant on NeurIPS and FAccT, June 5, 2023). The respondent went on to add that while the conference is changing over time, it still has work to do. Many respondents shared a desire to see the topics and research areas of the conference be actively reflected interpersonally and in the conference organization.

Taking Stock of DEI Now and Moving Forward

In this section, I describe five aspects of DEI efforts in conference organizing: disciplinary diversity, affinity groups, speaker and organizer diversity, travel support, and geographic reach. For each aspect of DEI work, I describe what NeurIPS, FAccT, and AIES conferences did in 2022 and the relevant experiences of BIPOC attendees with whom I spoke. In each section, I synthesize the feedback I received into recommendations for conference organizers.

Disciplinary and Practitioner Diversity

While fostering disciplinary and practitioner diversity is not a conventional facet of DEI initiatives, many respondents described this as a priority. This can look like adding additional tracks to the call for papers, which requires expanding the forms of expertise available within the program committee, area chairs, and via peer review. It can also look like supporting additional forms of contribution to the conference via practitioner-specific tracks, keynotes and keynote panels, professional development workshops, or other curated programs. These avenues create more opportunities for folks to be involved who may not know how to or be interested in submitting papers for a blind review, but whose work nonetheless can move the field forward. Expanding the forms of expertise that are considered relevant in the space recognizes the broad reach of AI ethics to a number of areas of practice, as well as the insights possible when established understanding in one field (e.g., sociologists' understanding of the social construction of race) is translated to another area where the state of the art may not yet account for them. Beyond academic disciplines themselves, participants argued for the relevance of “activism, political action, [and] community work, etc. ... [as] a meaningful contribution” (survey respondent #17, ACM FAccT & NeurIPS). These comments argued that meaningful intellectual work around AI ethics is done outside of traditional academic spaces by communities who face conditions of data-driven marginalization on an everyday basis.

Of the three conferences I attended, two communities had an explicit commitment to interdisciplinary and multidisciplinary work (FAccT and AIES), whereas NeurIPS was primarily focused on computer science (even as its workshops often featured other disciplines, such as the Resistance AI Workshop at NeurIPS 2020).³⁸ The FAccT conference states a commitment to a “big tent” in its strategic plan;³⁹ this growth in the number of

disciplines accepted is visible in the tracks announced in its call for papers (CFP) over time. In 2018, the CFP asked for submissions across four primarily computational categories (including natural language processing, or NLP; databases, and security), as well as a human-computer interaction track and another for law/policy/social science. By 2023, the total number of tracks in the call grew to 10, of which four were related to computer science, and others spanned law, historical and cultural critique, philosophy, and human factors. As the “big tent” philosophy of FAccT has grown, a few respondents mentioned a fear that computer scientists would not stay present for this cross-pollination—this concern was especially cited by respondents with respect to the Symposium for Foundations of Responsible Computing, whose fourth edition was held in 2023 at Stanford University and which primarily focuses on theoretical and formal approaches.

Expanding the forms of expertise that are considered relevant in the space recognizes the broad reach of AI ethics

These three conferences also varied with respect to the share and type of practitioner presence. While all three are academic conferences, there are a number of methods organizers can use to increase practitioner presence, including keynote invitations, curated tracks, and permitting organizers to program affinity groups alongside the conference. Co-located events such as industry recruitment events can also promote practitioner presence. AIES included the fewest practitioners and was primarily for academic work. While NeurIPS was not as interdisciplinary, it featured a strong industry presence across organizers, presenters, and attendees. It does not actively foster the presence of civil society organizations. FAccT was the most diverse in this respect: since the beginning, its tutorials track has featured a strong proactive orientation. In 2020, it started a curated portion of the program called “CRAFT”—Critiquing and Rethinking Accountability, Fairness, and Transparency—that actively fostered civil society presence.

Experiences of Respondents

Disciplinary and practitioner exchange was an overwhelming concern for survey respondents and the BIPOC attendees spoke about wanting researchers to go beyond theories and publishing papers. Some reported concerns with gatekeeping around what material was considered relevant to AI ethics scholarship, such as in the peer review process. They shared concerns about cultural homogeneity in the peer review pool reproducing conditions in which work that draws on other intellectual traditions is discounted. Speaking about the FAccT conference, one person shared:

On the one hand, you want to encourage diverse research. On the other hand, if your benchmark is Western research, that system ends up systematically excluding people that do not want to adhere to the status quo ... traditional academic presentation [is] constrained by all these measures that have been put in place to create an academic standard, but this standard is a double-edged sword, because it also serves to filter out ideas that are critical—that don’t fit the template. (Interview respondent on FAccT and NeurIPS, April 26, 2022).

Here, the path-dependence in constituting the boundaries of the AI ethics field was understood to be keeping critical, but nonetheless high-quality, work at bay.⁴⁰

Respondents also noted that the work presented at conferences on AI ethics has a tendency to be far more theoretical than applied work, with one survey respondent stating, “The topics themes and discussions of the accepted programs are disproportionately about topics that admit interesting algorithms/ML rather than on-the-ground needs. Honestly it can be very self-serving” (survey respondent #86 on ACM FAccT and NeurIPS). Some respondents felt that the topics and types of work presented at conferences were not radical enough, lacked advocacy, and did not sufficiently center the experiences of BIPOC communities or those who were significantly impacted by data-driven technologies: “It does feel like I don’t belong there because there are very few people that look like me. Very few people that work on critical perspectives, critical work as I do” (Survey respondent #02 April 26, 2023). A pair of conference organizers noted that one way to address this concern would be to have a track of featured advocacy work in the main program, as FAccT began to do in 2022 with its “community keynote” function and which it does in part through its CRAFT track. Similarly, one organizer of both FAccT and NeurIPS suggested that NGOs and civil society groups could have a featured or consultative role in developing conference programming. Another respondent noted that having a wide diversity of conference-goers is a benefit to traditional academics by creating opportunities to find new collaborators and impactful research sites.

Even as many of the respondents in my study discussed the value of multiple disciplines and practitioner perspectives, I also heard from organizers who expressed concern about the consequences of this expansion. One person shared a fear that if the conference became too interdisciplinary, computer science researchers would not find it to be as relevant. Another framed the problem as one of ensuring that those from other fields have a minimum amount of data expertise to meaningfully engage computational researchers; at the same time, they recognized that it is not necessary to understand the underlying mathematical properties to know when something is causing harm.

I note a tension that many respondents felt between personal and professional incentives. While these conferences are largely academic spaces, attendees have been vocally advocating for more inclusion not only across disciplines but to also include advocacy and policy groups as well as affected communities beyond what conferences attempt to do through speakers and workshops. In other words, part of the draw among many BIPOC attendees to AI ethics research is the recognition that this technology is being used to harm their communities or people like them. They expressed a distaste for empty theorizing about harms, and wanted to see more applied work that spoke to their values and experiences.

For others, the draw to participate in these conferences is intertwined with academic credentialing, a radically different incentive. Additionally, while AI ethics is inherently interdisciplinary, computer science as a discipline (or sub-areas like machine learning) does not always recognize AI ethics as contributing to its own evolution as a field. Insofar as the conference only addresses the need for credentialing and professional development, for those who are not in the same field or are not academics at all, there are fewer incentives to participate.

Recommendations

Shift peer review processes to better support engagement with other communities of practice, including transnational and non-Western scholarship.

- Make explicit calls for research that centers BIPOC ways of knowing and citational justice.
- Ensure more diversity in peer reviewers by disciplinary perspectives, domain area, and epistemological commitments.
- Train peer reviewers to assess for rigor and quality of scholarship, and to a lesser extent on ideas of relevance which may be overly path-dependent.
- Consider that a rebuttal period + revision process is also a means to support new disciplinary and geographic perspectives.
- Feature practitioners from advocacy and other applied work via curated tracks or by novel CFPs with their own submission and acceptance criteria.
- Ensure that Area Chairs can spot when papers are being rejected for irrelevant geographical or cultural reasons, rather than reasons of scholarly quality.

Listen to practitioners and invite those from multiple domains.

- Proactively conduct and sustain outreach to policy, advocacy, and civil society groups in recognition of their ability to speak to the real-world consequences of AI ethics.
- Formally invite practitioners to share what would be valuable to them from participating in an AI ethics conference and support those goals (e.g., meeting technologists for collaboration opportunities, connecting with interns and collaborators).
- Encourage translation as a desired norm for interdisciplinary conversations during the conference (e.g., defining key terms but also scaffolding discussions about key concepts and debates in a way that is accessible across disciplinary orientation).
- Waive registration fees for advocates or civil society groups who hold service positions or support conference organizing.

Create long-term and genuine connections with local communities while recognizing that AI ethics work cannot be done without them.

- Create genuine connections with local community members around the conference venue and engage with their concerns.
- Conceive of conferences as community bridges. Promote the conference locally, early, in order to support more community building with the local organizers to fortify connections with local and national organizations that are working toward the same goals.
- Recognize that community members and their advocates have different incentives to participate in these conferences than credentialed academics, and work to make the relationship intellectually and professionally reciprocal.

Affinity Groups

Affinity groups are spaces organized for sub- or alter-communities and can take on many forms, such as a dinner, a half- or full-day workshop, or even group sessions held by formal groups. These spaces provide networking opportunities for individuals who have similar experiences and a safe harbor from otherwise homogeneous conferences. Events held by groups like Black in AI, founded by computer scientists Timnit Gebru and Rediet Abebe, feature panels, mentoring sessions, and presentations that allow community members to connect and share their work. Queer in AI aims to center queer BIPOC people and provides scholarships, emergency aid, and graduate application support.⁴¹

NeurIPS hosted a full day of workshops for affinity groups a day prior to the official conference proceedings, including groups like Indigenous in AI, founded by Michael Running Wolf; Black in AI; LatinX in AI, founded by Laura N. Montoya; Queer in AI, founded and developed by a group collectively known as Organizers of Queer in AI⁴²; and Women in Machine Learning, founded by Hanna Wallach, Jenn Wortman, Lisa Wainer, and Amy Greenwald. Programming included paper presentations, poster sessions, panels, mentoring sessions, and more. FAccT had some workshop sessions during the critical CRAFT track that attendees referred to as affinity spaces-adjacent. While Queer in AI may not have met in the 2022 edition of FAccT, they held a meetup in other editions, including the 2023 conference in Chicago, Illinois. AIES did not have specific affinity group spaces. Notably, the bylaws of many conferences do not allow invite-only events to be part of the official schedule or happen at the formal conference venue, and thus affinity groups may sometimes meet as a para-conference gathering or not at all.

Experiences of Respondents

Affinity groups highlight members' work and help create community. Affinity groups also have been an important, dedicated, and visible route for funding BIPOC scholars to attend conferences. Attendees felt a greater sense of community when they attended a conference with affinity group support. Conference spaces can have hundreds if not thousands of participants; respondents shared that being able to find specific communities is very important to them. Many found themselves feeling at ease in affinity spaces: "I came with Black in AI, and I came with friends that I like. I didn't have any worries with NeurIPS" (Interview participant on NeurIPS, November, 30th, 2022). This person compared the affinity group experience favorably to experiences in other conferences, where discrimination had at times been palpable. Affinity spaces felt like spaces where BIPOC attendees could bring their whole selves:

This is the thing about Black women and Black culture, whether it's African American or African. When we go into a conference, it's not just our academic work we bring, it's our selves. It's our background, it's our culture, it's our vibe. When you go to NeurIPS you go to the main conference, it's like 95% white men—it just feels like you are out of place. And then you go to the Black in AI workshop and you just feel like home because it's familiar faces, it's your people ... Black in AI is slightly different than NeurIPS; it's like we bring our lifestyle, we have parties, we have Black music, we have Black joy. It's the same with FAccT... when we feel like FAccT is our home, that's also what we bring...that shapes FAccT that shapes the conference into what it is...So when Black women come into the conference, it's all this that that they bring that becomes part of the conference, that creates

that welcoming comfort zone for other Black people. (Interview participant on NeurIPS and FAccT, April 26, 2023).

Here, affinity groups foster a sense of home and belonging within conferences that can at times be overwhelmingly white or male-dominated. Another person described being at NeurIPS over time and how much better it became as organizers cultivated more affinity groups and workshops: “[Organizers] put a lot of effort into changing it... And I think it’s better now. It’s much better than before, because there’s like a lot of different kinds of [affinity groups and] workshops.” (Interview participant on NeurIPS and FAccT, June 5, 2023). This person explained that these groups allow attendees to create a space within a space to find community and people interested in the same things.

When we go into a conference, it’s not just our academic work we bring, it’s our selves. It’s our background, it’s our culture, it’s our vibe.

Some respondents enjoyed mentorship sessions and opportunities that affinity groups offered and wanted to see these opportunities expanded. One person expressed a desire for those more senior in the field to offer mentorship and dedicated time to younger scholars through affinity groups. “One thing I would love is if researchers could honestly just say, ‘Yes, there’s a problem.’ Part of it involves that we don’t mentor enough, [so] I will take it upon myself ... to mentor three kids” (Interview participant on NeurIPS, November 30, 2022). This quote highlights the opportunity for more senior scholars in the field to lead on intentional engagement with BIPOC students and early-career scholars.

While affinity group spaces were a great way for community building in uplifting scholars, many attendees experienced a lack of support from the wider conference community. For example, at NeurIPS, affinity group paper presentations and poster sessions were held separately from the rest of the conference, the day before many of the conference activities started. This led to very few attendees from the wider conference community participating. One conference attendee whose work was featured in an affinity group meeting noted that it might be more successful to schedule affinity group sessions in a more prominent part of the conference, stating, “Instead of making the affinity group workshops separately, before or after the main conference, it might be good to make it in the middle ... so the work of [women and BIPOC scholars] will be also visited by other communities” (Interview participant on NeurIPS, November 30, 2022). Multiple respondents felt sidelined as their work was not presented alongside other attendees, which led to feelings of disconnection. Holding all affinity groups at the same time also does not make space for conference-goers whose intersectional identities or allyship defies choosing a single event; allowing affinity groups to meet at different times throughout the program will allow attendees to participate in more than one identity group during a conference.

This highlights another consistent challenge in aligning two important goals of DEI work at conferences: establishing inclusive spaces where people can find community is sometimes in tension with making the conference as a whole more diverse and equitable. There is simultaneously a need to highlight work done by

and inside of BIPOC scholarly communities, and a need for integrating that work into the main program. One way through this challenge is for non-marginalized scholars to attend affinity-group-led sessions as frequently as they do the main program.

Recommendations

Promote wider access and engagement with affinity groups.

- Space affinity groups out across the conference program to encourage wider engagement and cross-pollination and to afford opportunities for this work to be championed by the entire conference community.
- Hold affinity groups events during the same days as the main conference program.
- Promote public affinity group events to all attendees.

Use curated tracks in the main program to feature affinity group work.

- Invite affinity group organizers to have the opportunity to curate a part of the main conference program.

Address experiences of disconnection by creating social infrastructure for newcomers.

- Revisit any official conference policy discouraging social gatherings, such as the FAccT policy on “Co-located events hosted by third-parties.”
- Design social events for all attendees that more intentionally foster social connection among newcomers.
- Create opportunities for mentorship between career stages.

Increasing Speaker and Organizer Diversity

One facet of DEI work is to increase the diversity of organizers and invited speakers. This can look like the organizing committee inviting significant academics and practitioners who are also BIPOC to keynote the conference or join conference leadership. Organizing committees and invited speakers are the most public “face” of the academic community, so deliberate representational diversity at these levels can have significant downstream effects. Speaker programming honors significant contributions to the space, and can be curated to address topical or thematic gaps in the rest of the program. While each of the three conferences studied had some diversity in speaker and organizing committee programming, one interview respondent shared that closely related AI, machine learning, and natural language processing conferences were not making reflective choices in this regard: “There weren’t any Black speakers, speakers were a lot of white people [and people] from big tech companies....All the Black people were in the workshop presentations. So when it comes to keynotes or guest speakers, I didn’t see any Black person speaking...which I’m still trying to comprehend.” (Interview participant on Empirical Methods in Natural Language Processing⁴³)

(EMNLP), March 14, 2023). This respondent's experiences in closely related conferences underscores what others shared about the dynamics in the AI field more broadly.

Many conference organizers stated the need for diversity at the level of both conference committees and invited speakers. However, few organizers spoke about program and organizing committee curation as a way of recognizing the substantive impact of BIPOC standpoints on the intellectual trajectory of the field. One organizer of AIES asserted that keynote programming requires organizers to be very reflective and intentional about which perspectives have previously been spot lit at the conference and which would make an important contribution to the community. Keynotes can be an important opportunity to cross-pollinate specific conference communities with highly relevant and essential contributions from closely related fields. For example, Media Studies scholar, Professor Andre Brock of Georgia Tech was one of the keynote speakers for the 2022 FAccT conference in Seoul; his work on "weak tie racism" provided multiple generative methodological and theoretical directions for the field and brought to light a novel form of algorithmically generated harm.

Experiences of Respondents

Respondents considered speaker and organizer diversity to be an important part of DEI. Multiple BIPOC attendees cited keynotes at AIES 2022 by computer scientists Inioluwa Deborah Raji and Rediet Abebe as playing a role in making them feel more welcome and included.

At the same time, some respondents considered proactive invitations to BIPOC scholars to be tokenizing. One survey respondent said: "Each time I present my work, author a book about it, discuss it... people that tokenize and objectify me, my work, [and] my community, eat it up, regurgitate it as their own discoveries and go to great lengths to erase, replace, extract, exploit and abuse" (Survey respondent #50 on AIES). Another survey respondent stated "Harms toward BIPOC communities are just sidelined, noted just as how certain applications are harmful and that's about it. The particular absence of liberatory approaches to tech that are [also] indigenous deeply concerned me. Keynote speakers are not the only ways to have that scholarship [be] present at conferences" (survey respondent #47 on AIES). Here, the respondent notes a narrowness to the inclusion of BIPOC scholarship to take note of harms, but not liberatory or alternative approaches.

Some participants, for instance at the AIES conference, noted a disconnect between the presence of BIPOC scholars among keynote speakers and their absence in conference leadership positions. One survey respondent said, "I want more extensive representation of marginalized groups at the organizational level. Again, I think AIES's leadership has been too white for the past few years" (Survey respondent #56 on ACM FAccT & AIES). Another person observed "the sheer volume of whiteness around and the positions of power that the white folks occupy" (Survey respondent #63 on AIES). Another participant shared the experience of joining an organizing committee for FAccT that was almost all white.

At the same time, conference organizers who have extended invitations to BIPOC scholars to co-organize note that their invitations compete with many other demands on these scholars' time:

There is a lot of research that especially those who are coming from underrepresented backgrounds, they're in high demand, meaning that they probably will be asked a lot more by a lot of different organizations and events to serve. So they might be actually

even more overloaded and overburdened. So it's a dilemma. (Interview respondent on FAccT, June 22, 2022).

This quote underlines that there is an ongoing tension between centering historically excluded voices as a push toward inclusion and respecting the limits (in time and labor) of prominent subgroup scholars.

BIPOC attendees highlighted the impact of organizer diversity on multiple aspects of their conference experience beyond the highly public level of keynotes—from panel discussions, to chairing a session, to moderating Q&A at the end of talks. In one case, a panel moderator directed questions about equity solely to a BIPOC panelist. Another session chair was visibly unfriendly to a non-native English speaker. Multiple respondents in the survey noticed opportunities to ask questions being directed solely to white men in Q&A segments. This indicates that inclusion should be an effort across the variety of settings in which a conference program unfolds from the highly visible and public keynotes to more mundane and semi-public conversations during paper presentations.

Recommendations

Consider how work is allocated between existing and new organizers.

- Be aware that some invitations to increase organizer diversity place additional burdens of time commitments and unpaid labor on BIPOC participants.
- Ensure that work assigned to local organizers includes work beyond operations support, including creative and high-profile opportunities like keynote programming.
- Recognize foundational work in the field by BIPOC scholars through distinguished keynote invitations and by programming panels to address questions raised by their work.

Uplift BIPOC scholars and their work.

- Spotlight liberatory and trailblazing work by BIPOC scholars and communities, rather than a sole focus on vulnerability and harm.
- More deeply engage BIPOC scholars beyond topics of race and equity.

Improve capacity for DEIA work and accessibility.

- Create a DEIA chair position in conferences like AIES that do not already have one.
- Provide training and guidance on inclusive moderation for session chairs and panel moderators.

Travel Support and Other Scholarships

Travel support is an important access need and is a form of equity. All three of the conferences provide the opportunity to apply for different levels and types of travel support to all attendees and presenters, and some affinity groups offer separate travel funding opportunities. Honoraria are typically also offered to keynote speakers and presenters, which become particularly important when inviting keynote speakers who do not have faculty or research funds available. Having funding to cover things like hotels, food, and travel make it possible for many people to afford high travel costs and conference fees.

NeurIPS each year has a financial assistance program to apply for travel support; the conference also has a strong presence of affinity groups, which are self-supported to provide their own travel funding to accepted participants.⁴⁴ In AIES 2022, financial support was provided solely for people accepted into the student program by waiving the conference registration and a need-based financial stipend, depending on the student's access to funding in their home institution. FAccT also has registration fee waivers and travel grants by application that are administered by the financial chair. In addition to need-based funding, in 2022 FAccT had surplus funds from its online-only conferences from the year before that were allocated toward the DEI fellows' travel expenses and fellowships.

Experiences of Respondents

Attendees were often excited to attend and be in conference spaces. Respondents expressed the desire for more scholarships and expanded travel support. Some of them have also had to contend with frictions around “reimbursement culture,” in which they are expected to have adequate credit and savings in order to pay for expensive travel upfront and wait to be reimbursed.⁴⁵ Large bureaucratic institutions, such as universities or the scholarly societies that host conferences, typically have long lag times between reimbursement claims and payment, creating challenging circumstances for attendees who struggle with financial precarity. For them, the cost of attending such events is already a significant burden, with expenses ranging from hundreds to thousands of dollars. As one interviewee put it, “We are institutionally funded, where probably when you get in, you have your institution paying for you. So first of all, even getting people to pay for your registration and booking the ticket for you, is a challenge—and it was a challenge.” (Interview respondent on EMNLP, March 14, 2023). These barriers are even more acute for independent scholars who lack the support of affiliated organizations and have to bear these costs by themselves. Other respondents noted that acceptance notices for conference papers often come too close to the conference itself, compounding the problem with more expensive flights. Many participants and organizers acknowledged that acceptance delays create tremendous issues around funding and visas which greatly hinder participation and representation in the conference space.

Increasing the types of conference participation opportunities available would also support students' ability to secure institutional resources for their attendance. One survey respondent expressed wanting to have different participation opportunities available to them that were not contingent on publishing papers—such as a pre-candidacy colloquium. They stated, “As an early researcher, it is so difficult for me to work on a project on time to submit a paper so as to help me get the scholarship to attend the conference. Can organizers also look out for some of us who are doing our best to stay strong in the field? I have so many ideas, and it's difficult to actually work toward these ideas to papers due to lack of resources and support.”

(Survey respondent #11 on NeurIPS & EMNLP). Another person suggested that increasing opportunities to formally participate in the conference as first- and second-year students would help foster collaborations for future paper submissions to the conference.

The gap in resources and opportunities, however, extends beyond securing travel support. One respondent suggested other methods of investment might be better than travel funds:

We came here, we presented the paper and we go back. [...] It doesn't bring any change to our community. [...] And if the people who provide grants ... travel grants and so on... if they're really concerned about the development of Black people or underrepresented people in general, they should have invested [in our communities]. Actually, my dream is NLP. I would love to create an NLP lab to help students in my country. (Interview participant on NeurIPS, November 30, 2022).

While travel funds are a necessary step in creating equitable conference participation, other forms of investment—guided by the scholars they are meant to support—will be necessary for long-term investment in diverse scholarship.

Recommendations

Provide resources and adequate timelines for those obtaining visas.

- Provide an earlier staggered acceptance process to support those with visa requirements.
- Designate a fast-response team with support letter materials and share their contact information with visa applicants in case of challenges.

Expand and advertise funding for DEIA efforts.

- Continue to invest and expand funding for DEIA efforts.
- Create a working group to assess how the conference's sponsorship policy and practices affect funding toward DEI initiatives and whether changes would be beneficial.
- Waive registration costs for PhD students, local students, and those experiencing hardship.
- Conduct outreach so more people are aware that funding is available.
- Direct funds and opportunities to labs doing BIPOC-led work.

Foster greater transparency and responsiveness in funding processes.

- Be clear about the projected amount of funding available, anticipated likely level of individual support, notification and reimbursement timeline, and point of contact in the event of problems.
- To the extent possible, reduce the amount which an attendee must float the cost of attendance.
- Advance cash payments to scholars who receive travel funding support so that there is no reimbursement process.
- If professional association (e.g., ACM) membership includes a discount, apply this discount to all attendees receiving any level of formal funding support.

Conference Locations and Geographic Diversity

The three conferences we examine for this report were founded in the US (NeurIPS and FAccT) and UK (AIES), but are imagined as global in their reach and concern. After all, AI systems and their consequences do not remain confined to the context of their development; they are transnational in scope and scale. The desire to feature more non-US/UK work has led to many conferences rotating their annual destinations. In theory, this creates opportunities for the conference to engage with local communities of scholars engaged in similar research pursuits who would otherwise find it difficult to participate in the conference, given the difficulties of obtaining a visa and the resources required. Western nations like the US are notoriously challenging for some foreign professionals to travel to, with the length of time necessary to secure a visitors visa often extending past the gap between acceptance to the conference and the event itself.⁴⁶

The three conferences I studied reflected this US-UK centrality but also a desire to connect with scholars globally. NeurIPS 2022 was hosted in New Orleans, Louisiana, in the United States

FAccT 2022 took place in Seoul, South Korea, which was its first edition outside of the US or Europe. AIES 2022 was held in Oxford, United Kingdom.

Experiences of Respondents

Multiple respondents noted the importance of rotating conference geography for visa purposes:

They always hold it in countries where most people can't attend anyway...if you call it international, it needs to be international. Right? If you're clear about it being like the US or whatever, sure. If it's supposed to be international, it's supposed to be international. And international is not the US and EU, or Canada. CVPR is the largest Computer Vision Conference, and it's supposed to be a North American Conference, and it is only held in the US and Canada. Never in Mexico. Not once in Mexico. "North America," right? (Interview participant on NeurIPS and FAccT, June 5, 2023).

To date, most editions of AIES and FAccT have been held in the US or Europe, where people in many countries face barriers to visa access. In 2019, FAccT published a strategic plan outlining a goal to lead in the geographic inclusivity of the conference, "across [the] country, across [the] globe (especially non-West and non-Northern Hemisphere)."⁴⁷ Since publishing its strategic plan, the FAccT conference has adapted its practices, establishing a pattern of hosting the conference in the United States and in other countries by alternating years; FAccT 2024 will be in Latin America. Many interview and survey respondents share this broader goal and underlined the extra work required to forge a genuinely international scholarly community.

While moving the conference to new locations does provide greater accessibility for some academics at the expense of others, respondents expressed other challenges. Given that conference leadership is still largely white, and US/UK-based, the labor to relocate proceedings to a new culture or host country is placed on local organizers. These local organizers are often not experienced event planners, but AI ethics scholars themselves, and are required to do exceptional amounts of cultural translation work between the main conference organizers and local vendors. For instance, it can be difficult for these local organizers to find spaces that can accommodate large conferences and the specific needs of the largely Western attendees. When conferences are held in the Global South or other non-Western locales, meeting these many needs falls to local

organizers and volunteers. This coordination also comes along with the literal translation labor—translations for contracts and fellow organizers.

When conferences are held in the Global South or other non-Western locales, meeting these many needs falls to local organizers and volunteers.

In addition to the factor of extra burden on local organizers, relocated conferences often do not change their focus or content to engage with local communities. The conferences may remain US-centric, or accept and promote the same style of scholarship regardless of physical location. As one respondent put it, “I still feel that the majority of the work at FAccT or AIES targets the topics of interest from Western countries and the Global South, [focus] is far lagging behind.” Respondents across the board described experiencing this gap between location and conceptual focus: “So it’s a bit of an oxymoron. On the one hand you want to encourage diverse and different research. On the other hand, if you end up measuring it—if your benchmark is Western research, that system, that way of doing reviews, ends up systematically excluding people that do not want to add to the status quo.” Efforts to increase geographic diversity, then, do not dislocate deeper barriers experienced by those offering non-Western perspectives and frictions they face in terms of the review criteria by which work is evaluated.

Recommendations

Be mindful of the visa issues for attendees.

- Choose countries with fewer visa barriers to entry for scholars from Global Majority countries.
- Host AI ethics conferences in countries in the African continent (c.f. The 2020 edition of the International Conference on Learning Representations, which was scheduled to be held in Addis Ababa but went virtual for COVID).

Build in a longer runway toward building more durable connections to scholars transnationally.

- Immediately on choosing a transnational conference location, begin building ties to relevant local scholars and civil society groups.
- Hold a bi-monthly transnational conference call to build community leading up to a new host country context and to build awareness.
- Declare conference dates years in advance so that local organizers have much longer notice on when the conference will be held in their location.
- Allow for a longer runway for negotiations with local vendors in new geographic areas to allow for more mindful and intentional planning.
- Anticipate and provide for translation expenses, especially for contracts, instead of expecting local organizers to do it.

Mitigate the extent to which moving locations places some of the most challenging conference organizing labor on likely-newcomers to the community.

- Hire a local event planner to avoid overburdening local scholars when planning conferences in the Global South.
- Consider traveling affinity spaces geared toward highlighting local topics, communities, and scholars.
- Evaluate the paper review processes to promote familiarity and engagement with a wider range of country-contexts and with non-Western scholarship.
- Build for relationships with new community members to endure beyond that year's edition.

Foster more opportunities for local scholars, organizations, and other interested attendees.

- Waive registration for everyone within a two-hour travel radius of the conference to support local community access.
- When in the United States, host conferences at or in partnership with historically Black colleges and universities, Hispanic-serving institutions, and tribal colleges.

Conclusion

A lot of people when they talk about why DEI is failing, or why a certain institution is not managing to hire or retain Black folks—unfortunately, a lot of the discussion revolves around putting the onus on Black people themselves. We should stop asking why aren't Black women coming here or why aren't Black people applying here or there. It should be: 'How do we change the environment? How do we ...attract Black women, and retain them?' ... The onus has to be on institutions on creating warm, welcoming environments, as opposed to putting the onus on Black women and Black folk. This is a lot of the misunderstanding I often see around initiatives around DEI." (Interview respondent on FAccT and NeurIPS, April 26, 2022).

A year after starting the fellowship and research for this work, I attended FAccT 2023. It was based in Chicago. I attended in person and got to follow up with some of the participants who had agreed to speak with me the previous year. One of the coauthors on this report was also a co-chair for CRAFT sessions. In the very drafting of this report, its findings have been profound; this year's FAccT conference featured a much wider array of advocacy, policy, and community members than previous ones. I even had the opportunity to present some of the findings from a draft of this report during a CRAFT workshop. Furthermore, the conference organizers made efforts to include local Chicago-based organizations such as the Invisible Institute, and Chicago-based community members were present at the event. There has, indeed, been some progress; however, it was disheartening to discover that many BIPOC individuals still shared similar experiences like the ones I have shared in this report. Being there in person, I not only saw the rougher edges of inclusion efforts, but heard new stories of cultural and institutional failings. Perhaps it was people's familiarity with my work on the topic, but this year, in arguably the most visibly inclusive FAccT ever, I had numerous participants confide in me with the most egregious stories of experiencing exclusion. Fellow BIPOC attendees were still experiencing similar burdens; they still shared feelings of discomfort. In terms of organizational support—while in the year prior there was a significant pot of money available for DEI funding (such as the fellowship that allowed for this research to happen) this year there was no budget allowance for similar research efforts.

The goal of this report was to analyze the effectiveness of specific, existing DEI efforts at AI ethics conferences—specifically from the point of view of BIPOC participants. And clearly, the experiences of these folks bears out many ways that these efforts could be improved. But I've also been motivated during this work by

the memory of finding truly welcoming and empowering spaces—just like the Data for Black Lives conference I walked into in 2018. And so, in addition to granular improvements and recommendations, I wanted to see if there are glimpses of that community, that connection, in other people’s experiences. And there are. People described to me the excitement of connecting with another BIPOC scholar in a room, or the joy people felt when meeting up for dinners, drinks, or dancing after hours.

Being there in person, I not only saw the rougher edges of inclusion efforts, but heard new stories of cultural and institutional failings.

Most of the work of conference DEI committees is focused on the problem of representation, or getting a broader demographic of people in the room. However, the experiences of people in the room as reported to me are not centered on this issue, but rather on issues of tokenization and collapse of their identity with their scholarship. By pointing this out, I do not intend to suggest that representation does not matter or does not deserve further material commitment and investment. Rather, my call is to *also invest in cultural and intellectual change that can address the experiential issues that make BIPOC attendees feel unwelcome once they are present*. This is why I opened with the quote from Ruha Benjamin, a reminder to craft the worlds we need—addressing the gaps in DEI work cannot only be about limiting the harms of past cultures or dismantling the biases of existing structures. We need to enact improvements with a vision of and excitement for the new spaces and communities we are building. This is, admittedly, the harder path for anyone pursuing DEI work, and a significant ask for people occupying temporary seats on a DEI conference committee—but it is nonetheless the way forward.

Inviting BIPOC scholars to participate in AI ethics spaces, after all, is only the first step. While at this moment AI ethics is taking off in the scope of its reach and impact, it is important to remember that AI as an approach to engineering and as a pillar of computer science is also growing at an incredible rate. The concerns and scholarship of BIPOC scholars invested in the impact of these technologies on real communities needs to find a path to inclusion, recognition, and ultimately substantive change beyond the halls and meeting rooms of AI ethics conferences. It is, after all, these deeper institutional sites of exclusion that create the problems that our scholarly community sees reflected and amplified by AI.

BIPOC scholars are participating in these conferences despite the experiences that I have recounted in this report. I imagine a world where BIPOC scholars join the field because of their meaningful experiences in conference spaces.

Acknowledgments

I would like to acknowledge and give my warmest thanks to the interviewees, attendees, and survey respondents who not only were open to having these important, candid conversations, but welcomed me into their community. Thank you to all those who have supported the paper including Lauren Thomas Quigley, Patrick Davison, Yeshimabeit Milner, Serena Oduro, Jenna Burrell, Lindsey Cameron, Evani Radiya-Dixit, and the many Data & Society colleagues who have spent their time ideating with me in a feedback session. Thanks as well to the Communications Team at Data & Society—Alessa Erawan, Eryn Loeb, Gloria Mendoza, Sona Rai, and Chris Redwood—for their patient production, promotion, and design work. Finally, I would like to thank the whole AIGI team for their amazing work and support of this project.

This work was supported by the inaugural DEI Scholars grant from the ACM FAccT Conference.

Suggested Citation: Emnet Tafesse, Meg Young, Ranjit Singh, and Jacob Metcalf.
“Blooming in Muddy Waters: DEI at AI Ethics Conferences.” Data & Society Research Institute.
<http://dx.doi.org/10.2139/ssrn.4570694>

Appendix: Methods

This research was conducted over one year using interviews, surveys, and observational data from three conferences broadly concerned with AI ethics in 2022. I attended: (1) the Fifth ACM Conference on Fairness, Accountability, and Transparency, which took place on June 21–24 2022 in Seoul, Republic of Korea, (2) the Fifth AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES) held on August 1–3, 2022 in Oxford, United Kingdom, and finally, (3) the Thirty-sixth Neural Information Processing Systems (NeurIPS) organized between November 28–December 9, 2022, in New Orleans, Louisiana, United States. Interviews were conducted during and after conferences with participants and conference organizers and spanned topics of conference events and organization, personal experiences in conference spaces, and the field more broadly.

These interviews were complemented with a survey to elicit input from a broader number of attendees than my interviewees. Survey questions asked about respondents' identities, which conferences they had attended, what forms of DEI and other relevant efforts they had observed, their experiences, and their input on how to make each conference more inclusive. Recruitment was conducted on-site by distributing fliers and online via Twitter; both methods directed to a SurveyMonkey page. The survey received 102 responses, of which 53 respondents fully completed the survey. Of the respondents, 26% described their employment as academic, 48% as graduate students, 15% as industry, 7% as civil society/nonprofit, 2% as government, and 2% as other. Of the respondents, 39% identified themselves as African American/African/Black, East Asian 18%, and Hispanic/Latinx 2%, Middle Eastern/North-African 2%, South Asian 10%, Southeast Asian 8%, and others 20%. The survey and interview invitation was promoted via tweet and email announcement at the same time as each key conference; this broadcast function resulted in one interview subject who had not attended FAccT, AIES, or NeurIPS, but who is an AI ethics researcher anchored in relevant NLP-specific conferences.

There were 16 interviews conducted with both attendees and organizers. Our interview results were qualitatively coded using an iterative, inductive approach to describe and cluster themes from respondents. Interviews were re-coded by two members of the team using the stabilized codebook. This project was overseen by an external Institutional Review Board, which classified the research as exempt under the US Common Rule. We have anonymized the identity of all respondents to preserve privacy.

Our data collection left the question of organizing DEI efforts open-ended, given that it can take a number of forms. The semi-structured interview protocol allowed respondents to provide their own perspective on these efforts. The survey questions were less open-ended and provided a list of common aspects of conference organizing that often constitute DEI and asked participants to indicate which they had observed at a given conference, including diverse keynotes, travel support, affinity groups, child care, equity-themed papers and workshops, anti-harassment policies, or other efforts (with an invitation to write a comment in an open text box).

We note that the commitments of DEI efforts extend far beyond racial equity measures to include aspects like accessibility for wheelchairs and people who are hard of hearing, childcare support, and private rooms for breastfeeding or religious observation; however, this project is scoped more narrowly around the experiences of people who are Black, indigenous, and POC and the salience of race and ethnicity. This choice is not to limit a broader understanding of what constitutes DEI, but to follow the first author's research agenda of exploring joy and solidarity in BIPOC spaces.

First author Emnet Tafesse conducted all interviews, designed and distributed the surveys, and led initial data analysis. During coding, all members of the team collaborated on identifying themes and parsing transcripts. Emnet and team member Meg Young created the code book against which all of the data was re-coded. Emnet and Meg worked to produce an initial outline and shared the drafting of each section, with all team members helping to contribute writing, revisions, and citations on individual sections. Following internal review of an initial draft, Emnet and Meg took primary responsibility for finishing the draft, corresponding with external experts, and carrying out rounds of revisions.

Endnotes

- 1 **Daphne Barretto, Julienne LaChance, Emmanuelle Burton, and Soohyum Nam Liao**, “Exploring Why Underrepresented Students Are Less Likely to Study Machine Learning and Artificial Intelligence,” *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education* *V* 1, USA, ACM, June 2021, *Crossref*, doi:10.1145/3430665.3456332.
- 2 The term “AI ethics” is highly contested. **Abeba Birhane, et al.** note that the term “ethics” is too abstract, and not adequately focused on harms, justice, and equity. Nonetheless, we are using “AI ethics” in this report as a widely recognizable shorthand while acknowledging its philosophical limitations. See **Abeba Birhane, Elayne Ruane, Thomas Laurent, Matthew S. Brown, Johnathan Flowers, Anthony Ventresque, and Christopher L. Dancy**, “The Forgotten Margins of AI Ethics,” *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (2022): 948–58.
- 3 **Daniel E. Acuna and Lizhen Liang**, “Are AI Ethics Conferences Different and More Diverse Compared to Traditional Computer Science Conferences?” *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (July 2021), *Crossref*, doi:10.1145/3461702.3462616.
- 4 **Richard Fry, Brian Kennedy, and Cary Funk**, “STEM Jobs See Uneven Progress in Increasing Gender, Racial and Ethnic Diversity,” (Washington DC: Pew Research Center, April 1, 2021), www.pewresearch.org/science/2021/04/01/stem-jobs-see-uneven-progress-in-increasing-gender-racial-and-ethnic-diversity/.
- 5 **Daniel Zhang, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, Deep Ganguli, Barbara Grosz, Terah Lyons, James Manyika, Juan Carlos Niebles, Michael Sellitto, Yoav Shoham, Jack Clark, and Raymond Perrault**. *The AI Index 2021 Annual Report*, <https://aiindex.stanford.edu/ai-index-report-2021/#individual-chapters> (2021).
- 6 **Latanya Sweeney**, “Discrimination in online ad delivery.” *Communications of the ACM* 56, 44–54 (2013); **Latanya Sweeney**, “Discrimination in online ad delivery: Google ads, black names and white names, racial discrimination, and click advertising,” *Queue* 11 (2013): 10–29.

- 7 **Joy Buolamwini and Timnit Gebru.** “Gender shades: Intersectional accuracy disparities in commercial gender classification.” In *ACM Conference on Fairness, Accountability and Transparency* 77–91 (2018); **Inioluwa Deborah Raji and Joy Buolamwini.** “Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial AI products.” In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* 429–435 (2019).
- 8 **Abeba Birhane and Vinay Uday Prabhu,** “Large image datasets: A pyrrhic win for computer vision?” In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)* 1536–1546 (2021). doi:10.1109/WACV48630.2021.00158.
- 9 **Inioluwa Deborah Raji,** “From Algorithmic Audits to Actual Accountability: Overcoming Practical Roadblocks on the Path to Meaningful Audit Interventions for AI Governance.” In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* 5 (Association for Computing Machinery, 2022). doi:10.1145/3514094.3539566; **Inioluwa Deborah Raji and Joy Buolamwini.** “Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial AI products.” In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* 429–435 (2019); **Inioluwa Deborah Raji, Peggy Xu, Colleen Honigsberg and Daniel Ho,** “Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance,” Preprint at <https://doi.org/10.48550/arXiv.2206.04737> (2022).
- 10 **Sandra Harding,** ed., *Feminism and methodology: Social science issues*, (Bloomington: Indiana University Press, 1987).
- 11 **Patricia Hill Collins,** “Black feminist thought in the matrix of domination.” *Black feminist thought: Knowledge, consciousness, and the politics of empowerment*, 138 (1990), 221–238.
- 12 I worked with my collaborators Lauren Quigley, Adriana Alvarado Garcia, Juana Becerra Sandoval, Ranjit Singh, and Tsige Tafesse to hold a CRAFT session (that is, a curated conference track for “Critiquing and Rethinking Accountability, Fairness, and Transparency”) dedicated to fostering a space for BIPOC attendees, discussing how we see the conference, and how we wish it would shift.
- 13 **Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao.** “The Values Encoded in Machine Learning Research,” in *ACM Conference on Fairness, Accountability, and Transparency* 173–184 (Association for Computing Machinery, 2022). doi:10.1145/3531146.3533083.
- 14 **Holly Else,** “AI Conference Widely Known as ‘NIPS’ Changes Its Controversial Acronym,” *Nature*, (November 2018), <https://doi.org/10.1038/d41586-018-07476-w>. **Anima Anandkumar,** “#ProtestNIPS: NIPS Acronym Encourages Sexism and Is a Slur. Change the Name.” Change.Org. October 25, 2018. <https://www.change.org/p/members-of-nips-board-protestnips-nips-acronym-encourages-sexism-and-is-a-slur-change-the-name>. **Jennings Brown,** “NIPS’ AI Conference Changes Name Following Protests Over Gross Acronym,” Gizmodo. (November 19, 2018), <https://gizmodo.com/nips-ai-conference-changes-name-following-protests-ov-1830548185>.
- 15 **Isabelle Hupont, Emilia Gomez, Songul Tolan, Lorenzo Porcaro, and Ana Freire.** “Monitoring Diversity of AI Conferences: Lessons Learnt and Future Challenges in the DivinAI Project.” *arXiv preprint arXiv:2203.01657* (2022).

- 16 **James Charlton**, *Nothing About Us Without Us: Disability, Oppression and Empowerment*, (University of California Press, 2000); **Sasha Costanza-Chock**, *Design Justice: Community-Led Practices to Build the Worlds We Need*, (MIT Press, 2020); **Safiya Umoja Noble**, *Algorithms of Oppression: How Search Engines Reinforce Racism*, (NYU Press, 2018).
- 17 See also **Olúfẹ́mi Táíwò**, “Being-in-the-Room Privilege: Elite Capture and Epistemic Deference.” *The Philosopher* 1923. (November 20, 2021), <https://www.thephilosopher1923.org/post/being-in-the-room-privilege-elite-capture-and-epistemic-deference>; **Olúfẹ́mi Táíwò**, *Elite Capture: How the Powerful Took Over Identity Politics* (Chicago: Haymarket Books, 2022).
- 18 **Nick Mitchell**, “14. Diversity.” In *Keywords for African American studies*, pp. 68–74. New York University Press, (2018).
- 19 **Sara Ahmed**, *On being included: Racism and diversity in institutional life*, Duke University Press, (2012); **Roderick A. Ferguson**, *The Reorder of Things: the University and its Pedagogies of Minority Difference*. (2012); **Nick Mitchell**. “14. Diversity.” In *Keywords for African American studies*, pp. 68–74. New York University Press, (2018).
- 20 **Roderick A. Ferguson**, “The reorder of things: The university and its pedagogies of minority difference.” (2012).
- 21 **Sara Ahmed**, *On being included: Racism and diversity in institutional life*, Duke University Press, (2020).
- 22 **Lily Hu**, “What is ‘Race’ in Algorithmic Discrimination on the Basis of Race?” *Journal of Moral Philosophy* (2022).
- 23 **Lily Hu and Issa Kohler-Hausmann**, “What’s sex got to do with machine learning?” In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (2020, January).
- 24 **Ruha Benjamin**, “Innovating inequity: if race is a technology, postracialism is the genius bar.” *Ethnic and Racial Studies* 39, 1–8 (2016); **Ruha Benjamin**, *Race After Technology: Abolitionist Tools for the New Jim Code*, (John Wiley & Sons, 2019).
- 25 **Sareeta Amrute**. “Bored Techies Being Casually Racist: Race as Algorithm.” *Science, Technology, & Human Values* 45, 903–933 (2020).
- 26 **Christina Harrington, Sheena Erete, and Anne Marie Piper**, “Deconstructing Community-Based Collaborative Design: Towards More Equitable Participatory Design Engagements.” *Proc. ACM Hum.-Comput. Interact.* 3, 216:1–216:25 (2019).
- 27 **Elijah Anderson**, “The White Space.” *Sociology of Race and Ethnicity* 1, 10–21 (2015); **David L. Brunson, Nathaniel G. Chapman, Joong Won Kim, J. Slade Lellock, Megan Underhill, Erik T. Withers, and Jennifer Padilla Wyse**. “The Culture of White Space: On the Racialized Production of Meaning.” *American Behavioral Scientist* 64, 2001–2015 (2020).
- 28 **Daniel E. Acuna and Lizhen Liang**, “Are AI Ethics Conferences Different and More Diverse Compared to Traditional Computer Science Conferences?” in *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* 307–315 (Association for Computing Machinery, 2021). doi:10.1145/3461702.3462616; **Ali Akbar Septiandri, Marios Constantinides, Mohammad Tahaei, and Daniele Quercia**, “WEIRD FAccTs: How Western, Educated, Industrialized, Rich, and Democratic is FAccT?” in *2023 ACM Conference on Fairness, Accountability, and Transparency* 160–171 (2023). doi:10.1145/3593013.3593985.
- 29 **Peter Pels and Oscar Salemink**, *Colonial Subjects: Essays on the Practical History of Anthropology*. (University of Michigan Press, 2000).

- 30 American Anthropological Association. “1973 Minority Report - Participate & Advocate.” <https://www.americananthro.org/ParticipateAndAdvocate/Content.aspx?ItemNumber=1514> (1973); American Anthropological Association. “The Minority Experience in Anthropology: Report of the Committee on Minorities and Anthropology.” <https://www.americananthro.org/ParticipateAndAdvocate/Content.aspx?ItemNumber=1842> (1996); **Janis Hutchinson and Thomas C. Patterson**, “Final Report: Commission on Race and Racism in Anthropology.” (CCRA). <https://www.americananthro.org/ParticipateAndAdvocate/Content.aspx?ItemNumber=1910&RDtoken=27348&userID=6944> (2010); **Karen Brodtkin, Sandra Morgen, and Janis Hutchinson**, “Anthropology as White Public Space?” *American Anthropologist* 113, 545–556 (2011).
- 31 **Charlton D. McIlwain**, *Black software: The Internet and racial justice, from the AfroNet to Black Lives Matter*, (Oxford University Press, 2019).
- 32 **Suzanne T. Bell, Anton J. Villado, Marc A. Lukasik, and Larisa Belau, and Andrea L. Briggs**, “Getting Specific about Demographic Diversity Variable and Team Performance Relationships: A Meta-Analysis.” *Journal of Management* 37, 709–743 (2011); **Mahijs de Vaan, David Stark, and Balazs Vedres**, “Game Changer: The Topology of Creativity.” *American Journal of Sociology* 120, 1144–1194 (2015); **Mathias Wullum Nielsen, Sharla Alegria, Love Börjeson, Henry Etkowitz, Holly J. Falk-Krzesinski, Aparna Joshi, Erin Leahey, Laurel Smith-Doerr, Anita Williams Woolley, and Londa Schiebinger**, “Gender diversity leads to better science,” in *Proceedings of the National Academy of Sciences* 114, 1740–1742 (2017); **Scott E. Page**, *The Diversity Bonus: How Great Teams Pay Off in the Knowledge Economy*, (Princeton University Press, 2019).
- 33 **Bas Hofstra, Vivek V. Kulkarni, Sebastian Munoz-Najar Galvez, Bryan He, Dan Jurafsky, and Daniel A. McFarland**, “The Diversity–Innovation Paradox in Science,” *Proceedings of the National Academy of Sciences* 117, 9284–9291 (2020).
- 34 **Richard Fry, Brian Kennedy, and Cary Funk**, “STEM Jobs See Uneven Progress in Increasing Gender, Racial and Ethnic Diversity.” Pew Research Center Science & Society. <https://www.pewresearch.org/science/2021/04/01/stem-jobs-see-uneven-progress-in-increasing-gender-racial-and-ethnic-diversity/> (2021); **Gallup, Inc. and Google, Inc.** “Diversity Gaps in Computer Science.” <http://goo.gl/PG34aH>; **Maggie Johnson**, “The Computer Science Pipeline and Diversity: Part 1 - How did we get here?” https://ai.googleblog.com/2015/07/the-computer-science-pipeline-and_8.html (2015); **Maggie Johnson**, “The Computer Science Pipeline and Diversity: Part 2 - Some positive signs, and looking towards the future.” <https://ai.googleblog.com/2015/07/the-computer-science-pipeline-and.html> (2015); 1. **Susan Colby, Helen Ma, Kelsey Robinson, and Lareina Yee**, “What It Will Take to Make the Tech Industry More Diverse,” in *Harvard Business Review* (2016).
- 35 **Marie desJardins and Ran Libeskind-Hadas**, *Growing and Diversifying the Domestic Graduate Pipeline*, (2021). https://www.nsf.gov/cise/advisory/CISE_GraduatePipeline_2021.pdf.
- 36 **Daniel E. Acuna and Lizhen Liang**, “Are AI Ethics Conferences Different and More Diverse Compared to Traditional Computer Science Conferences?” In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* 307–315 (Association for Computing Machinery, 2021). doi:10.1145/3461702.3462616.

- 37 **Sheena Erete, Yolanda A. Rankin, and Jakita O. Thomas**, “I Can’t Breathe: Reflections from Black Women in CSCW and HCI.” *Proceedings of ACM Human-Computer Interaction* 4, 234:1-234:23 (2021); **Christina N. Harrington**, “A Call to Action for the ACM,” Medium, <https://medium.com/@cnicoleharrington/a-call-to-action-for-the-acm-a3274675de03> (2020).
- 38 “Resistance AI Workshop” created by William Agnew, Abeba Birhane, Micah Carroll, Elliot Creager, Agata Foryciarz, Pratyusha Ria Kalluri, Sayash Kapoor, Suzanne Kite, Raphael Gontijo Lopes, Manuel Sabin, Marie-Therese Png, Maria Skoularidou, Mattie Tesfaldet, Mariya Vasileva, Ramon Vilarino, and Rose E. Wang, held December 11, 2020 at NeurIPS Conference. <https://sites.google.com/view/resistance-ai-neurips-20/home>.
- 39 “ACM FAccT Strategic Plan,” Revised January 25, 2020. https://facctconference.org/static/docs/strategic_plan.pdf.
- 40 This method of exclusion by defining what constitutes relevance to the conference recalls Sara Ahmed’s (2012) idea of “kinship technology,” by which an organization reproduces the same membership and social relations over time. See **Sara Ahmed**. *On being included: Racism and diversity in institutional life*. Duke University Press, 2020.
- 41 **Anaelia Ovalle, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J. Sutherland, Davide Locatelli, Eva Breznik, Filip Klubička, Hang Yuan, Hetvi J, Huan Zhang, Jaidev Shriram, Kruno Lehman, Luca Soldaini, Maarten Sap, Marc Peter Deisenroth, Maria Leonor Pacheco, Maria Ryskina, Martin Mundt, Milind Agarwal, Nyx McLean, Pan Xu, A Pranav, Raj Korpan, Ruchira Ray, Sarah Mathew, Sarthak Arora, ST John, Tanvi Anand, Vishakha Agrawal, William Agnew, Yanan Long, Zijie J. Wang, Zeerak Talat, Avijit Ghosh, Nathaniel Dennler, Michael Noseworthy, Sharvani Jha, Emi Baylor, Aditya Joshi, Natalia Y. Bilenko, Andrew McNamara, Raphael Gontijo-Lopes, Alex Markham, Evyn Döng, Jackie Kay, Manu Saraswat, Nikhil Vytla, Luke Stark**, “Queer In AI: A Case Study in Community-Led Participatory AI.” In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (2023, June), 1882–1895.
- 42 Organizers of Queer in AI is the formal name used to refer to the following collective of people: Anaelia Ovalle, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J. Sutherland, Davide Locatelli, Eva Breznik, Filip Klubička, Hang Yuan, Hetvi J, Huan Zhang, Jaidev Shriram, Kruno Lehman, Luca Soldaini, Maarten Sap, Marc Peter Deisenroth, Maria Leonor Pacheco, Maria Ryskina, Martin Mundt, Milind Agarwal, Nyx McLean, Pan Xu, A Pranav, Raj Korpan, Ruchira Ray, Sarah Mathew, Sarthak Arora, ST John, Tanvi Anand, Vishakha Agrawal, William Agnew, Yanan Long, Zijie J. Wang, Zeerak Talat, Avijit Ghosh, Nathaniel Dennler, Michael Noseworthy, Sharvani Jha, Emi Baylor, Aditya Joshi, Natalia Y. Bilenko, Andrew McNamara, Raphael Gontijo-Lopes, Alex Markham, Evyn Döng, Jackie Kay, Manu Saraswat, Nikhil Vytla, and Luke Stark.

- 43 While EMNLP is not one of the primary conferences I studied, this respondent was selected for their relevant experience in AI ethics conferences and workshops in natural language processing (NLP) communities.
- 44 “NeurIPS 2022 Financial Assistance” Conference on Neural Information Processing Systems, <https://neurips.cc/Conferences/2022/FinancialAssistance>.
- 45 **Ana Rosado**, “Reimbursement Culture Is Keeping Graduate Students in Debt.” *Medium* <https://medium.com/@arosad/reimbursement-culture-is-keeping-graduate-students-in-debt-f10ef76af12c> (2019); **Sarah Thomson**, “Reimbursement Culture and Widening Participation in Academia,” *Impact of Social Sciences*, <https://blogs.lse.ac.uk/impactofsocialsciences/2019/03/27/reimbursement-culture-and-widening-participation-in-academia/> (2019).
- 46 **Bloomberg Editors**. “Coming to the US Shouldn’t Be This Hard.” *Bloomberg.com* (2022); **Ari Daniel**, “Why visas are a hot topic right now at the International AIDS Conference,” *NPR* (2022).
- 47 “ACM FAccT Strategic Plan,” FaccT website, (2020). https://facctconference.org/static/docs/strategic_plan.pdf.

