

# Towards Justice-Oriented AI: Reflections on a Research-Community Partnership to Co-Develop Swahili AI Tools

Sybillé Légitime<sup>1</sup>, Kirabo Kakopo Atuhurira<sup>2</sup>, Audrey Mugeni<sup>2</sup>, Catherine D’Ignazio<sup>3</sup>, Zawadi Amany<sup>2</sup>, Harini Suresh<sup>1</sup>

<sup>1</sup>Brown University

<sup>2</sup>Femicide Count - Kenya

<sup>3</sup>Data + Feminism Lab, MIT

sybille.legitime@brown.edu

## Abstract

AI development is notoriously centralized, concentrating power in the hands of a few. Justice-oriented frameworks have emerged to challenge this paradigm, but documented examples of AI projects that aim to distribute power and ownership across the full development lifecycle remain rare. In this paper, co-authors from a Kenyan activist organization that monitors femicide in the region join with a research team based in the United States to reflect on our co-development of Swahili AI tools for femicide monitoring. We adapted two existing tools previously developed by the research team to the Kenyan context—through data collection, model training, and interface localization—prioritizing situated technical capacity-building by recruiting a local developer to lead the effort.

Both tools were successfully deployed and are in use by the community partner. Still, we found ourselves asking a question we suspect many justice-oriented technology efforts also face: “Was the project successful?” Beyond producing the technical artifacts, had we succeeded in sharing power and ownership throughout the development process? To explore this question, we turn to reflexivity as a way to foreground the *process* of technology production alongside its outcomes. Through collaborative autoethnography and drawing on Data Feminism and African Data Ethics as analytical frameworks, we examine how structural factors and project design decisions limited or supported justice-oriented principles throughout the development process. We end by distilling broader insights on relationship-building, the roles of access and documentation in distributing control, and how shifting from performance metrics to partner-benefit measures can better ensure project sustainability. Through this work, we propose normalizing reflexive practice as a vital tool for assessing and advancing just AI development, offering both a critical account of one justice-oriented AI project and broader insights that can shape similar efforts.

## 1 Introduction

The dominant AI development paradigm is centralizing power and resources at an accelerating pace. Closed models, opaque development processes, and evaluations that focus on “general-purpose” benchmarks further the divide between technology producers and the communities their

products are designed to serve. Researchers across disciplines have responded by advocating for interventions that minimize harm, such as reducing bias in model outputs, prioritizing transparency, or building accountability processes. While valuable, these approaches have been challenged for focusing primarily on individuals and technical systems rather than on the structural conditions that produce harm (D’Ignazio and Klein 2023). Justice-oriented approaches, by contrast, are rooted in feminist, decolonial, or other critical frameworks that interrogate and challenge power, calling for development practices that redistribute control over data, labor, knowledge, and technical infrastructure (Klein and D’Ignazio 2024; Mohamed, Png, and Isaac 2020; Ricaurte and Zasso 2023; Young, Katell, and Krafft 2022; Birhane et al. 2022a; Barrett et al. 2025).

Community-led and participatory approaches have been proposed as key ways to redistribute power in practice, particularly those that center long-term partnerships, mutuality, and the needs of marginalized communities (D’Ignazio and Klein 2023; Birhane et al. 2022b; Costanza-Chock 2020). Documented examples have engaged communities to shape specific parts of the AI development process, such as data collection, annotation, or model evaluation (Nekoto et al. 2020; Patton et al. 2020; Suresh et al. 2022; Cooper and Zafiroglu 2024). Work on infrastructuring further focuses on building structural conditions for community ownership and control that extend beyond any single project (Dantec and DiSalvo 2013; Karasti 2014).

Yet surveys of the broader landscape of participatory and community-led AI have found that stakeholders are rarely involved in problem formulation, technical development and data analysis, or long-term ownership (Delgado et al. 2023; Cooper et al. 2022; Feffer et al. 2023). These gaps persist in part due to significant structural barriers, including conflicting incentives, resource and time constraints, trust and relationship-building challenges, and the infrastructural changes required to redistribute control over typically centralized material resources (Young et al. 2024; Suresh et al. 2024; Maas and Inglés 2024). Researchers and developers navigating these barriers lack methodological guidance and shared repertoires of experiences and tools to draw on.

This paper responds to that need by documenting and reflecting on a concrete justice-oriented AI project, contributing both a critical account of the effort and method-

ological guidance for others pursuing similar goals. Specifically, we reflect on a yearlong research-community partnership to build Swahili AI tools for femicide monitoring. We adapted two tools previously developed by the research partner organization to the Kenyan context through data collection, model training, and interface localization. Our process was guided by the principles of *Data Feminism* (D'Ignazio and Klein 2023) and *African Data Ethics* (Barrett et al. 2025). While both frameworks emphasize challenging power and engaging in reciprocal, community-based relationships, African Data Ethics further highlights investing in local talent and infrastructure. In our project, we aimed to support situated technical capacity-building by hiring a local developer to lead the development effort. At the same time, we observed tensions that arose as we ran into the structural barriers mentioned above, such as extended timelines and the practical need to centralize technical infrastructure. Reflecting on our approach, we found ourselves asking: “*Was the project successful?*” One answer might stop at evaluating the technical artifacts: both tools were deployed and performing well. But to evaluate the project against its justice-oriented aims, we also needed to ask: **in what ways did we succeed—or fall short—in sharing power and ownership throughout the development process?**

To answer these questions, we take a collaborative auto-ethnography approach, synthesizing qualitative data from two focus groups within our team, two written self-reflections, and project traces. Our analysis contributes:

1. A critical examination of how key steps taken in our collaboration process both (a) aligned with justice-oriented principles from Data Feminism and African Data Ethics and (b) shaped the project process, partnership, and outcomes.
2. Insights on specific strategies for AI projects to share design, ownership, and control with community partners in the context of enduring structural barriers.

Through our contributions, we propose normalizing reflexive practice as a vital method for assessing and advancing just AI development, foregrounding the *process* of technology production alongside its outcomes.

We begin by reviewing participatory and community-led approaches to AI, particularly their applications in Global Majority and African countries, as well as prior work on reflexive and autoethnographic practice. Next, we trace the project’s chronology and describe the Swahili AI tools we co-developed. We then outline our methodology. Our analysis reveals how structural barriers, context-specific factors, and our commitment to distributing technical capacity shaped our work. From this, we distill insights on relationality, the many forms of and barriers to ownership, and how to assess and build accountability into justice-oriented projects.

We build on the AIES and broader critical AI community’s work on conceptualizing and implementing shared control and ownership in AI/ML projects (Maas and Inglés 2024; Seger et al. 2023; Feffer et al. 2023; Eslami et al. 2025; Hutchinson et al. 2025; Suresh et al. 2022, 2024; Queerinaï et al. 2023; Tseng et al. 2025) by distilling key strategies—

including the more systematic integration of reflexive practice into researchers’ evaluation toolkits—for pushing the boundaries of justice-oriented AI in action.

## 2 Related Work

### 2.1 Approaches for Increased Community Control in AI Development

Computing research has addressed the involvement of community stakeholders in technology development. Work across HCI and AI mainly draws on participatory traditions (e.g., Participatory Action Research (PAR) or Participatory Design (PD)), with aims of creating more “democratic” development processes, producing technology that works better for its intended users, and/or mitigating default power imbalances in research and design (Cooper and Zafiroglu 2024; Delgado et al. 2023; Birhane et al. 2022a). A subset of this work focuses on “community-based” or “community-led” research, partnering with stakeholders from, e.g., a community-based organization. Cooper et al. (2022) characterize this work under the broad umbrella of “community-collaborative approaches” (CCA): those aiming to “involve collaboration with community stakeholders as co-researchers throughout the research process, while investigating complex problems of concern to the community (e.g., social and health-related), and developing novel technologies or sociotechnical solutions.” CCA approaches to creating AI-based tools have worked with community members to annotate social media posts from the community (Patton et al. 2020), annotate data and evaluate models for femicide detection (Suresh et al. 2022), collect natural language data for machine translation of low-resource African languages (Nekoto et al. 2020), and build AI literacy resources around algorithmic decision making systems (Katell et al. 2019).

While participatory and CCA approaches often share the aim of distributing meaningful decision-making power and involving participants as co-researchers, structural barriers can hinder this. For example, drawing from Sherry Arnstein’s “Ladder of Citizen Participation” (Arnstein 1969) scholars have argued that participants can end up taking more of a “consultation” role (rather than a “partnership” or “ownership” one) (Corbett, Denton, and Erete 2023), and highlight the risks of “participation-washing” (Sloane et al. 2022). Other work has shown that community partners are typically not involved at all stages of the research project (Magaña and Shilton 2025). Delgado et al. (2023) describe how participatory AI approaches tend to involve participants in shaping user interface design or data collection, rather than the design of AI models or the overall purpose of the system, and Cooper et al. (2022) find that community stakeholders in CCA approaches are rarely involved in data analysis or problem definition. These same scholars highlight the issues of organizational timelines, resource constraints, and academic funding and publication incentives, all of which do not support, e.g., outcomes such as building community capacity or decentralizing technical know-how. At the same time, increased community participation in data-centered projects is not necessarily a

panacea and can place an “epistemic burden” on minoritized communities (Pierre et al. 2021). Many scholars, recognizing this reality, have explicitly pushed beyond conventional community involvement, framing development projects as sites for dismantling structural injustice (Costanza-Chock 2020; Kasirzadeh 2022; Leslie et al. 2022; Birhane et al. 2022b; Edenberg and Wood 2023), while others have explored how participatory design can build infrastructure for sustained community ownership and self-determination, instead of narrowly focusing on project scope (Karasti 2014; Dantec and DiSalvo 2013; Agid 2016). Our work explores disparities in community engagement with research: we deliberately aimed to share control over stages of the project, such as AI model development, while reflecting on how the challenges we encountered affected the research-community partnership, and on the project’s overall contribution to long-term community ownership and self-determination.

## 2.2 Community-Centered AI in the Global Majority and Africa

Beyond the structural forces that shape collaboration within participatory and CCA approaches, engagements with communities in countries of the Global Majority are further shaped by histories of imperialism, colonial extraction, and ongoing forms of neocolonialism (Mohamed, Png, and Isaac 2020). Prior work has documented how these dynamics manifest in technology research and development through asymmetries in funding, representation, agenda-setting, knowledge valuation, and ownership of technical artifacts (Correa Lucero and Martens 2025; Septiandri et al. 2023). Community-centered approaches in these settings must contend not only with interpersonal power differentials, but also with transnational inequalities that structure research partnerships and constrain community self-determination (Png 2022).

A range of civil society organizations across Africa, the Americas, and Asia have developed community-rooted approaches to data and AI work that explicitly center local priorities, lived experience, and long-term accountability. Examples include feminist and community-based organizations such as Data Género<sup>1</sup> in Latin America, Pollicy<sup>2</sup> in Uganda, DOHS Cares<sup>3</sup> in Nigeria, and Tattle<sup>4</sup> in India. In the African context, the focus of the current work, prior academic research has described participatory AI efforts to evaluate image generation models (Becker et al. 2025), study the adoption of AI in creative sectors (Ngaruiya et al. 2024), and annotate Swahili datasets (Kamande et al. 2025). Nekoto et al. (2020) describe Masakhane<sup>5</sup>, an expansive participatory project to collect datasets and build language technology for a range of African languages.

Research has also proposed broader feminist and decolonial frameworks to guide such work. *Data Feminism* (DF) (D’Ignazio and Klein 2023) articulates how intersectional

feminist principles can be used to surface and challenge structural inequalities embedded in data-driven systems and practices, and has been adapted by practitioners in African contexts to reflect Afrofeminist perspectives and priorities (Iyer, Chair, and Achieng 2021). *African Data Ethics* (ADE) similarly offers a framework rooted in African philosophy and the realities of data science work on the African continent, articulating principles that foreground “communalism, self-determination, and cultural preservation” (Barrett et al. 2025). While developed from distinct intellectual traditions, DF and ADE share overlapping commitments to challenging power and centering pluralistic, situated knowledge. Our work is guided by both frameworks; in particular, we draw on the shared principle to *Challenge Power*, the ADE principles of *Grounding Design in Community-Engaged Consensus* and *Asserting Self-Determination* (“By Africans, For Africans”), and the DF principles of *Embracing Pluralism* and *Elevating Emotion and Embodiment*.

## 2.3 Reflexivity and Autoethnography

Reflexive practices have long been emphasized within feminist thought, stemming from the recognition that knowledge production is inseparable from the producer’s identity and sociopolitical context (Collins 2022; Harding 2013). Haraway’s conceptualization of *situated knowledge* posits that actively recognizing and reflecting on one’s role in the process of knowledge production leads to more objective claims than claiming neutrality by distancing oneself (Haraway 2013). Extensive work in critical AI has highlighted how researcher, developer, or data worker positionality shapes project agendas, dataset construction, and design decisions, encouraging reflexive approaches that foreground this subjectivity (Walker et al. 2025; Scheuerman and Brubaker 2024; Miceli et al. 2021; Schroeder, Pareek, and Barocas 2025; Cambo and Gergle 2022; Miceli, Schuessler, and Yang 2020). While researchers implementing participatory AI approaches have published about the outcomes of their efforts, fewer papers have provided reflexive accounts of specific AI development efforts (Suresh et al. 2022; Ciolfi Felice et al. 2025). We argue that such reflexive accounts, by attending to researcher positionality, institutional incentives, infrastructural constraints, and relational dynamics, are key to pushing the boundaries of justice-oriented AI.

Autoethnography (AE) is a qualitative research approach that originated in the social sciences in the 1980s as a response to growing recognition that personal experience shapes the research process and that dominant research traditions often privileged a narrow set of perspectives (Ellis, Adams, and Bochner 2011); it has since been adopted in HCI (Rapp 2018). AE treats researchers’ first-person experiences as analytically meaningful data rather than as peripheral context, drawing on sources such as fieldnotes, personal reflections, project artifacts, and correspondence. These materials are analyzed through iterative cycles of writing and reflection to connect lived experience to broader social, cultural, or institutional structures. In this way, autoethnography offers a way to methodologically instantiate reflexivity, making visible aspects of research practice that are of-

<sup>1</sup><https://datagenero.org/>

<sup>2</sup><https://pollicy.org/>

<sup>3</sup><https://dohscares.org/>

<sup>4</sup><https://tattle.co.in/>

<sup>5</sup><https://www.masakhane.io/>

ten excluded from technical or outcome-focused accounts (Kaltenhauser, Stefanidi, and Schöning 2024).

Collaborative autoethnography (CAE) shifts from individual analysis to collective sensemaking and can enable deeper insights to emerge through dialogue (Chang, Ngunjiri, and Hernandez 2016). Work in HCI has framed (collaborative) autoethnography as a feminist method that views the “personal as a valid site of knowledge production” (Garcia and Cifor 2019) and allows for attention to topics, such as internal power dynamics or points of “failure,” that are typically left out of research publications (Howell, Desjardins, and Fox 2021). CAE is particularly well-suited to reflecting on work that was itself collaborative, allowing team members to learn from and support each other, share power by collectively offering and examining each other’s experiences, and surface points of convergence and tension (Chang, Ngunjiri, and Hernandez 2016). Our use of this method builds on the work of Ciolfi Felice et al. (2025), who use CAE to provide an account of implementing a feminist AI project in Latin America, highlighting tensions and strategies, and making visible the “feminist work” that went into the effort.

### 3 Background: Collaboratively Developing Swahili AI-Based Tools

Here we describe our main case study: the collaboration between a research partner and a community partner to support the community partner’s data collection workflow for counting femicide cases in Kenya. We provide context on the project partners, the tools developed through the collaboration, and the project chronology.

#### 3.1 The Project Partners

**Data Against Femicide** Data Against Femicide (DCF) is a South-North participatory action research and design project that supports the work of activists, human rights monitors and civil society organizations that produce data on gender-related violence and its most lethal expression: femicide/femicide<sup>6</sup>. Since 2019, DCF has sought to understand and share the practices, challenges, and strategies of data activism against femicide in different contexts, to foster an international community of practice around the production of data on femicide, and to co-design and develop technological tools to support data production and social communication about femicide. Leadership of DCF is transnational and inter-institutional, with core contributors based in Argentina, the United Kingdom, the United States, and Uruguay. The DCF network, events, and educational offerings bring together a wider network of partners across Latin America, North America, Europe, and Africa.

**Femicide Count - Kenya** Femicide Count - Kenya (FC-K) is a feminist research, documentation, and prevention ini-

<sup>6</sup>While there are some key differences in the origins and definitions of these two terms (femicide/femicide), for the purposes of this paper we use the terms that our partners use, “femicide” in the case of our research partner and “femicide” in the case of our community partner in Kenya (D’Ignazio 2024).

| ID | Affiliation            | Role           |
|----|------------------------|----------------|
| A1 | Data Against Femicide  | Developer      |
| A2 | Femicide Count - Kenya | Lead developer |
| A3 | Femicide Count - Kenya | Co-founder     |
| A4 | Data Against Femicide  | Co-founder     |
| A5 | Femicide Count - Kenya | Team member    |
| A6 | Data Against Femicide  | Collaborator   |

Table 1: Research Team Composition. A1, A2, ... refer to the first, second, and subsequent authors, respectively.

tiative established in 2022, building on the project Counting Dead Women Kenya, launched in January 2019. The project emerged in response to the systemic invisibilisation of femicide in Kenya, where killings of women and girls are frequently underreported, misclassified, or framed in ways that blame victims rather than interrogate perpetrators and enabling structures. FC-K was founded to address this gap by generating publicly accessible, survivor-centered data on femicide in Kenya, while simultaneously challenging the cultural, legal, and political conditions that allow such violence to persist.

FC-K employs a mixed-methods, community-anchored approach that combines longitudinal media monitoring, case documentation, qualitative analysis, and participatory research. The project counts femicide cases through systematic review of print and digital media, analyses broader patterns of violence, and works with institutions and community groups to improve response mechanisms. The FC-K team is interdisciplinary and African- and feminist-led, grounded in the principle that those most affected by violence must shape the knowledge produced about it.

**The Research Team** Our research team consists of team (co-)leads from both DCF and FC-K, the lead developer hired as part of the collaboration, and researchers/team members from both organizations. Roles and author identifiers are summarized in Table 1. The project also involved additional contributors, including US-based DCF members who helped with technical support, former FC-K team members, and volunteer translators.

#### 3.2 Supporting the Data Activist Workflow: the Email Alerts and Data Highlighter Tools

The core development work involved adapting two existing AI-based tools developed by DCF to the Kenyan/Swahili-language context. Both tools were initially co-designed with activist organizations monitoring femicide in primarily Latin American and US-based contexts to support their data collection processes. Below, we describe each tool and the development work required to adapt them for FC-K’s context.

**Data Against Femicide Email Alerts** Email Alerts is a digital tool that delivers relevant news articles (i.e., about femicide in a particular region) to activists monitoring femicide or other forms of gendered or racialized violence in a particular context. The tool first fetches articles from a global, open-source media collection platform using tailored

queries, then filters them using a machine learning model trained to predict relevance for a specific geographic region, language, and/or type of violence.

To train a new Swahili model for the Kenyan context, A2 (with support from A1 and A5) collected and annotated 515 articles from Kenyan and Tanzanian Swahili news sources, using FC-K's definition of femicide as a guide. The final dataset was evenly split between femicide and non-femicide articles. Model development was led by A2, with the support of A1 and A6. The model was trained using code adapted from DCF's codebase, which combined a pre-trained embedding model with a logistic regression classifier. With  $P@20=0.89$ ,  $AUC=0.87$ ,  $FPR=0.21$ , and  $FNR=0.16$  on 5-fold cross-validation, the classifier was deployed to the Email Alerts platform and monitored by the FC-K team for 2 weeks.

**Data Against Femicide Highlighter** Data Highlighter is a Chrome extension that automatically detects names, locations, dates, and custom-defined terms of interest in a news article to support information extraction. Data Highlighter also allows activists to link spreadsheets of collected articles to streamline data collection.

Prior to this work, the Data Highlighter relied on spaCy models for named entity recognition (NER) in English, Spanish, and Portuguese. Since SpaCy does not support Swahili, A2 and A1 explored and tested alternative NER models, ultimately incorporating a pretrained XLM-ROBERTa model finetuned on the MasakhaNER dataset, which includes named entities annotated by Masakhane for 10 African languages (Adelani et al. 2021). The tool was deployed to the Chrome Web Store and tested by the FC-K team on sample Swahili articles.

The user interfaces for both tools were translated into Swahili. Translations were co-developed by A2 and A1, with assistance from native speakers. Swahili versions of both tools were launched at a broader community event organized by DCF, and FC-K presented them at a feminist AI conference in Nairobi, Kenya.

### 3.3 Project Chronology and Main Phases

**Establishing the Partnership.** The collaboration began with DCF's initial outreach to FC-K in 2021, followed by a second round of discussions in the summer of 2024 and project kickoff in the fall of 2024 after DCF secured funding and assigned internal points of contact for the project. The collaboration stemmed from FC-K's interest in adapting the tools previously developed by DCF to support their work in the African context (with a particular focus on Swahili-language data). DCF secured grant funds from its affiliated US-based institution that covered FC-K staff's estimated time on the project, from initial project meetings through tool deployment and evaluation.

**Lead Developer Recruitment.** Recruitment began by drafting the interview process and job posting, followed by cold outreach through DCF and FC-K's networks to feminist, activist, and women-in-tech spaces across the African continent. Screening involved an initial interview assessing both experience in AI development and alignment with

FC-K's mission, conducted jointly by FC-K and DCF team members. Candidates advancing with FC-K's approval completed a coding challenge modeled on expected work for one of the Swahili tools, followed by a solution walkthrough interview with DCF interviewers. Finally, the interviewers deliberated on the candidates' overall performance to make their final decision and extend offers.

**Lead Developer Onboarding.** Once the offer letter was accepted, we created and shared a statement of work outlining the project's expected deliverables, payment rates, and terms for the candidate to review. We also produced an onboarding document, which included the key repository and documentation resources for A2. Finally, A2 was included in all-team biweekly meetings, and developer-specific biweekly meetings were established with A2, A1, and A6.

**Data Collection and Model Development.** The primary development work extending the Email Alerts and Highlighter to Swahili happened over the course of February–August 2025. We extended the timeline to accommodate finding and collaboratively annotating Swahili-language news articles for training data. We extended the timeline to accommodate finding and collaboratively annotating Swahili-language news articles for training data. Weekly meetings continued throughout, with code review and co-development sessions held both on schedule and as needed. Throughout the project, A2 was compensated at a rate of \$50/hr using funding allocated by A4 and A6.

**Results Dissemination.** We launched the Swahili Email Alerts and Data Highlighter tools at a community event hosted by DCF, then A2 and A3 presented the project as a keynote at the Gendering AI Conference in Nairobi, Kenya.

## 4 Methods

We aim to use reflexivity as an intentional feminist practice that foregrounds researcher subjectivity and relational dynamics in the research process. Specifically, we use a collaborative autoethnography methodology, drawing in particular from the work of Ciolfi Felice et al. (2025). As part of the CAE, subsets of the core collaborators participated in two focus groups, A3 and A2 contributed to dialogic self-reflections, and we collected project traces for analysis. The process enabled us to engage in a collaborative, critical retrospective on the project and the extent to which project design decisions supported or limited justice-oriented aims. The study did not involve recruiting participants outside the partnership or collecting sensitive information.

**Focus Groups.** The initial focus group consisted of a 90-minute Zoom session, facilitated by a volunteer moderator. Participants (A1, A2, A3, A4, A6) summarized and critically reflected on their involvement in and experiences throughout the project: recounting its "origin story", highlighting novelties, challenges, lessons learned, and celebrations during its key phases, and speculating about opportunities for regional collaboration. The discussion guide included open-ended questions on project initiation, overall reflections on project successes and challenges, reflections on collaboration, and potential interventions to foster a liberatory, community-based research methodology. The second

focus group, moderated by A1, was a 60-minute Zoom session in which participants A1, A2, and A5 reflected on their involvement in the project, their experiences participating in model development, and their impressions of the knowledge transfer process.

**Self-Reflections.** An initial review of the focus group transcripts indicated that aspects of the project warranted further exploration. Therefore, in addition to the focus groups, A2 and A3 contributed written self-reflections. The written reflections raised the FC-K team's impressions on the usefulness of the deployed Swahili tools, the lead developer recruitment process, transparency in collaboration, and the project's role in advancing FC-K's mission.

**Project Artifacts.** Reviewing the project artifacts, both generated during the partnership and inherited from the research partner's existing infrastructure, provided further insight into how DCF's organizational and technical infrastructure influenced the project's unfolding. Meeting notes provided traces of the actors, discussions, and decisions involved in the project; process documents and contracts enshrined the partners' principles and missions in organizational structures and protocols; and onboarding documents, setup guides, technical documentation, and code repositories informed capacity-building and knowledge transfer throughout the project.

**Analysis.** In autoethnography, data generation and analysis/interpretation may overlap (Chang, Ngunjiri, and Hernandez 2016; Ciolfi Felice et al. 2025). For example, we discussed salient points of convergence and tension during the focus group, forming an initial set of ideas about key themes. A1 subsequently conducted a round of close reading and open coding of the focus group transcripts and self-reflections, supported by relevant project traces. A1, A2, A3, and A6 discussed and iterated on emergent themes, which were further solidified and articulated during the writing of the narrative.

**Limitations** Our reflexive assessment process was not without limitations. First, while we created multiple avenues for participants to share their honest opinions about the project, underlying power dynamics within the team (e.g., A2's funding originating from DCF) may still have influenced how these views were expressed. Second, the CAE process was conducted after tool deployments. Embedding reflections throughout development could have yielded actionable insights in real time. Finally, incorporating a broader range of perspectives from collaborators beyond the core research team—such as FC-K's affiliated translators or maintainers of DCF's technical infrastructure—might have further enriched our insights.

## 5 Findings

Research-community collaborations often operate within structures that limit shared power and ownership. While we could not entirely escape these constraints, we actively worked to challenge them, particularly in the technical domain. In this section, we draw on the guiding principles of *Data Feminism* (DF) and *African Data Ethics* (ADE) to frame and contextualize this effort, and we evaluate the

project's "success" through reflexivity and the CAE process. Our analysis focuses on three interrelated dimensions: the specific steps we took that instantiated principles from DF and ADE; the barriers we faced in pursuing justice-oriented AI and how we navigated them as a team; and the ways in which our goals were problematized or reshaped through practice. Throughout, we foreground the process of technology production alongside its outcomes.

### 5.1 Setting the Project Foundations

**Project Initiation** Recognizing at the outset the need for this project to extend the Femicide Count - Kenya's control and agency into the technical domain signaled a departure from the normative, colonial mode of partnership and built trust for our cooperation. By the time the partnership began, Data Against Feminicide had been in its fifth year of co-designing AI tools with human rights organizations across the Americas. This extensive experience informed its view of centrality and the concentration of resources by a single party. Notes the DCF co-founder:

*We're making tools for activists who are largely in Latin America, but then all of our servers and all the technical expertise are still centered with us. And so this was our experiment with trying to think about that and think about, well, how do we not just centralize. That feels like a very colonial dynamic, so how do we decentralize the server infrastructure [and] the technical know-how? (A4).*

These centralized structures and the power imbalances they reflect speak to a colonial dynamic predicated on extraction and exploitation. The FC-K co-researchers, acutely aware of this history, contrasted their experience at the initiation of this partnership with initial conversations with a different prospective technical partner:

*[The other prospective partner] doesn't talk to us anymore, but what they'd wanted was to question: how do you know that the data you are collecting is correct? How do you know that what you are doing is what is supposed to be done? (A3)*

Subverting those normative expectations for our partnership reflected a continued desire to enact the principles focused on *Challenging Power* within the DF and ADE frameworks. Where DF calls for naming and challenging power differentials at the source, ADE introduces the colonial lens, highlighting the distinct dynamics that arise in a partnership between Kenya-based and US-based organizations. However, some structural features of the project still reproduced centralization and surfaced power imbalances. For example, securing funding was a precondition for initiating a partnership, and the research partner's capacity to finance and allocate resources, constrained by the funding agenda, created a clear impediment to dismantling colonial power. When our timeline stretched due to developer recruitment and collaborative data annotation, A4's funding source ran out; continuing the project relied on A6 reallocating funding from another source so we could continue compensating A2 for her work. Institutional bureaucracy further exacerbated this,

imposing administrative burdens on the community organization (e.g., leading to delayed payments) and revealing enduring inequalities.

We therefore acknowledge that decentralizing technical expertise was a meaningful step toward dismantling power asymmetries, and naming it explicitly within the partnership gave it additional weight. Yet, this effort still unfolded against the persistent backdrop of colonial infrastructures and practices that continue to shape research-community collaborations (Leslie et al. 2022). True equity-based collaboration extends beyond the objectives and dispositions of individual collaborators; it requires funding structures that are distributed among, and accountable to, the communities whose interests the project serves.

**Designing and Carrying out Developer Recruitment** A defining feature of our project was recruiting a developer from FC-K's context. Designing and conducting developer recruitment enabled us to prioritize Afrocentrism, create new structural resources, and exemplify horizontal governance and decision-making. Decisions regarding recruitment networks, job-posting language, and interview structure were informed by FC-K's stated core values. Recruitment was explicitly Afrocentric, with a strong preference for women candidates, targeting African Women's tech networks in addition to FC-K's network. The FC-K co-founders were instrumental in shaping the process by participating in the non-technical interview round, reviewing job postings and interview materials, and evaluating candidates to ensure values alignment. This process embodied the ADE principle *For Africans, By Africans* under the broader aim to *Assert Data Self-Determination*, which calls for "local African data practitioners [to] lead data science work so its development is properly situated in the communities in which the technology will be deployed" (Barrett et al. 2025). Giving the developer recruitment process the care it required, however, extended the project timeline far beyond what we had anticipated, and in ways that required us to adjust later project stages:

*When I look back at the original timeline that I had proposed. [...] I think I wrote in 2 weeks for us to find somebody, and that was so unrealistic. It was actually more like 4 months or something (A4).*

In hindsight, this delay was not surprising; many community-research collaborations struggle to function under tight timelines (Cooper et al. 2022). At the time, though, we had underestimated what the process would demand.

Our different positionalities as co-researchers also revealed the limits of some of our perspectives, shaping how we engaged with this principle. The DCF-affiliated researchers came from diverse cultural backgrounds (A1 is Haitian; A4 and A6 are American with lived experiences and family origins in Latin America and India, respectively), yet were all trained and operated within US-based private university systems. Those of us affiliated with FC-K, by contrast, were embedded in East Africa: A2 is Ugandan and worked from Uganda, A3 is Kenyan and worked from Kenya, and A5, who had joined FC-K while in Kenya, is Kenyan and collaborated from the US. In particular, A2, our

lead developer, shared a strong interest in FC-K's values but had limited proficiency in Swahili. This situation confronted the DCF side of our team with the particularities of Africa's ethnic, linguistic, and cultural diversity—a reality the FC-K partners already knew intimately. It pushed us from DCF to examine what "local context" meant for this project, and to ask: did East African regional proximity constitute sufficient local context? Did the lead developer need to be Kenyan, or more specifically, from Kenya's Coast Province? Rather than a fixed category, "local" emerged as layered and relational. Regional proximity fostered situated knowledge and trust among the FC-K members, but limited linguistic proximity constrained A2's direct engagement, requiring mediation by volunteer translators. Those of us in DCF had not anticipated this need. Centering FC-K's values and vision reflected our team's justice commitments, but it also surfaced real tensions in a collaboration between a non-Kenyan developer and a Kenyan partner. Had DCF been based at a Ugandan or Kenyan institution, questions of recruitment timeline and geographic scope might have been defined more precisely from the start.

## 5.2 Collaboration and the Development Process

**Making Time Across Space** Geographic and organizational distance created early coordination challenges, as is common in decentralized projects. In addressing these, we established shared secure communication channels and consistent team meetings to prioritize transparency, open communication, and horizontal collaboration throughout the project. Beyond information sharing, this allowed us to find our team rhythm and cultivate trust. Through Signal groups, biweekly developer and all-team meetings, we shared progress updates, discussed impediments, requested feedback, and collectively made decisions that shaped the project's direction. FC-K's team leader summarizes:

*Sometimes our work, our research, is taken away from us, or it is taken advantage of all the time but we got into such a good rhythm, we managed to work really well and to work like partners instead of working, [...] how organizations work when you're working with organizations from the Global North (A3).*

Despite the time difference (necessitating early morning meetings for some researchers) and changes to FC-K's team configuration over the duration of the project, our commitment to commonly implemented structures of collaboration anchored our partnership (Drahota et al. 2016), and helped us fully instantiate the ADE principle to *Ground Design in Community-Engaged Consensus*.

**Centering Capacity-Building, and Meeting Uncertainty with Flexibility** We recognized that this work would necessitate bidirectional knowledge exchange, so we anticipated and planned for these moments while remaining open to spontaneous learning opportunities. Our approach began with structured onboarding: we provided the A2 with a guide (including technical documentation, application user guides, and points of contact), and organized team meetings. We then encouraged A2 to share her acquired system knowledge with the broader FC-K team.

In reality, knowledge exchange happened continuously through impromptu coding sessions, collaborator meetings, and both group and direct messages. These exchanges included system walkthroughs and debugging sessions led by DCF's dedicated application maintainers, as well as peer coding sessions among the Swahili model development team members. Knowledge transfer occurred between organizations, with A2 serving as a key bridge: DCF researchers shared technical support, while FC-K researchers provided domain knowledge crucial for data collection and annotation. Data annotation in particular became a relational and collaborative process, reflecting DF's call to *Embrace Pluralism* rather than impose a universal or predefined annotation scheme. As A2 recounts:

*[Knowledge sharing] really went both ways, because prior to joining the team, I had maybe third-hand [knowledge of femicide] [...] I knew of the femicide crisis that was happening, at least locally where I was, and within the context of East Africa. But I didn't know how to define femicide. [During data annotation] I had to learn what is a mark of femicide, what isn't a marker, and that was a bit of a process where I had to speak to [A3] and [A5], and they sort of gave me a benchmark to look at and say, this article is an instance of femicide, this article isn't, and there were definitely a lot of grey areas, so I'd have to say, hi, do you think this is an instance of femicide. And it was a whole learning process (A2).*

Leading the data collection and annotation process required A2 to learn FC-K's methods and transfer that knowledge to the DCF researchers. In the process, we gained not only practical skills but also empathy for the emotionally and mentally demanding work of reading stories about women and girls whose lives have been tragically scarred or ended by violence: all of us annotators (A1, A2, and A5), despite expressing a similar heightened awareness of cases of gender-based violence in our respective home countries, shared different emotional responses, reporting feelings of "numbness" and "hopelessness" at the alarming and tragic state of gender-based violence in Kenya. Still, this offered us a renewed appreciation for FC-K's work. The collaborative nature of the data collection process, built on mutual learning and a shared solidarity around the problem of femicide, supported the DF principle to *Elevate Emotion and Embodiment*, which recognizes that "this form of collective, community-engaged, intentionally emotion-laden work can also remind us of our interdependence, a key idea from disability studies that emphasizes how we rely on each other, and our actions can have consequences on others" (Klein and D'Ignazio 2024).

**Recognizing Dependencies and Overcoming Bottlenecks** Completing this project successfully required "a village." Beyond the core design and development team, we relied on system maintainers and volunteer translators. This highlights an important reality: successful university-community partnerships often demand more than core research teams alone can provide. Universities and community organizations frequently have shifting structures and members juggling

multiple commitments, which means projects typically require additional consultation, volunteer support, feedback, and resources. Recognizing the value of resources and processes inherited, reused, or repurposed by either partner can help us identify opportunities to revise, refine, and extend them for current or future partnerships. By identifying what can be reproduced and what enables decentralization, while resisting the urge to apply these elements uniformly across all projects, we can better distinguish between useful templates and rigid boilerplate approaches. Equally present in our work, and highlighted in the Project Initiation section (5.1.1), were reminders of FC-K's dependence on DCF. The Email Alerts and Data Highlighter application infrastructure is hosted on DCF's servers, and access to the systems' APIs and codebases is restricted to DCF's team members. Within DCF, a small group of developers holds key technical knowledge and system permissions. While such access control ensures efficiency (e.g., load balancing) and security, it also creates bottlenecks. More significantly, it centralizes key technical components, establishing a publish-subscribe (pub-sub) dynamic rather than a peer-to-peer relationship. A1 remarked:

*[Often] our development kind of depended on [the system maintainers] to look into something, or give us access to something. There are some pieces there that slowed down development just a little bit, [...] but they were also very available, and open to my questions. I would say that was a nice highlight to be able to learn the system through these people (A1).*

However, we reflected that dependence in and of itself does not necessarily undermine the ADE principle of *challenging colonial power*. There is a fine line between sovereignty and the burden of full ownership or responsibility. We found that the trust we built between DCF and FC-K allowed FC-K to benefit from the system's value without bearing the costs of ownership and operation. Because we operated from a norm of making decisions together, FC-K felt comfortable with DCF maintaining the system's technical infrastructure. At the same time, this dynamic stops short of meeting the ADE principle to *Invest in Physical & Organizational Infrastructures*. Shifting material infrastructure without imposing an undue maintenance burden on FC-K, however, would require significant investment in local support and funding beyond a single, part-time developer.

### 5.3 Assessing and Disseminating Project Results

**Possibilities for Improvement** Reflecting on knowledge sharing and project execution, FC-K team members identified areas for process improvement. These included highlighting and describing key infrastructure components in the onboarding documentation and formalizing the mentorship and knowledge-transfer process from the hired developer (A2) to FC-K by dedicating specific time for skills transfer. Responding to what she would have liked to know prior to her engagement in the project, A5 said:

*The knowledge sharing on my end was kind of more of, like, 'I didn't know what I didn't know' until I had*

*to find out. I came to learn about the goals of the entire partnership [...] a couple of months into the partnership. So maybe just, like, being aware of it, knowing this is what we are working towards (A5).*

In a similar vein, A2 recommended that documentation be provided highlighting the system's key infrastructure components, on which the Swahili tools would be built. Both impressions reveal tensions in applying the ADE principle of *grounding design in community-engaged consensus*, highlighting the need to allocate time, space, and resources to continuously communicate the development project's process and impact to any project newcomers, particularly when they may be geographically and institutionally separated from the research team.

A2 and A5 also proposed extensions of the Swahili tools, ranging from localized datasets to train the NER model and classifier for different Swahili-speaking regions, to a multi-class variant of the Email Alerts model to capture more context-specific variants of femicide in Kenya (e.g., child or elderly victims, or those connected to "suspicions of witchcraft"). When asked about the current and future usefulness of the Swahili tools, the A3 describes:

*I will find the Swahili tools most useful when they are adopted beyond Kenya and utilized by other Swahili-speaking African countries to document and analyze femicide systematically. Their real value lies in enabling cross-country learning, comparability, and collective accountability across the region. (A3)*

These ideas highlight how the FC-K's team involvement in technology development enabled them to exercise agency over the project's future. In line with the DF principle to *Consider Context*, the proposed future directions were situated in the Kenyan context and FC-K's objectives.

### **Celebrating Achievements and Communicating Impact**

After deploying the tools, we deliberately sought to share our work beyond academic settings. First, A1 and A2 presented at a broad community event held by DCF to a global audience of activists engaged in similar counterdata collection work. Second, A2 and A3 presented the project as a keynote talk at the Gendering AI Conference in Nairobi, Kenya. We found that the latter gathering in particular supported FC-K in gaining exposure, articulating the project on its own terms, and building new connections. These relationships laid the groundwork for new partnerships and opportunities for regional coalition-building. Notably, since communicating about the project regionally, FC-K has secured additional funding to continue its research mission. Finally, in addition to our academic conference article, we plan to publish for a general audience. Led by the community partner, this article will focus explicitly on the tool's impact on advancing the partner's objectives and vision.

Reflecting on the project's impact on her organization, A3 writes:

*The project has clarified the importance of academic-movement partnerships in strengthening regional collaboration. While I have long been interested in working with universities in Kenya, this work*

*made visible both the gaps in access and the possibilities that emerge when feminist, community-led research is taken seriously as knowledge production. Engaging in this project has opened new pathways to think about collaboration beyond informal networks and towards more structured partnerships with academic institutions. It has also highlighted the potential for [FC-K] to act as a pioneer in regional feminist data production on femicide, particularly within Swahili-speaking contexts. [...] Looking ahead, [FC-K] is intentionally strengthening its research and writing capacity and preparing to share its work more widely. In 2026, we will be collaborating with a technical university in Kenya, engaging students in feminist research practices and knowledge translation. These steps position the project as a bridge between community knowledge, academic inquiry, and regional feminist collaboration. (A3)*

## **6 Discussion**

As our project wrapped up, we found ourselves asking: *"Was the project successful?"*, a question many similar efforts likely face. The mainstream, outcome-based approach to quality assessment would lead us to answer "yes", since both tools were successfully deployed and performing well. To evaluate the project against its justice-oriented goals, however, requires going beyond standard performance metrics, and results in a more nuanced answer. We explored reflexive practice as a method for evaluating not only the project's outcomes but also its *process*. Collaborative autoethnography foregrounded our experiences and relationships, creating space to reflect on enduring challenges and tensions that are typically left out of published results. Attending to these dimensions helped us understand *how* the project advanced justice practices and *where* it fell short: the ways in which it did and did not share ownership; the ways in which it built relationships and trust; and the ways in which it supported the community partner's ability to pursue their mission.

We now turn to the broader lessons and strategic interventions to navigate structural barriers that can benefit other AI-driven and justice-oriented projects. We argue that more systematically embedding reflexive practices into evaluation toolkits can advance the field of participatory AI by: (1) elevating processes and relationships as central to project success; (2) making visible the structural barriers to justice-oriented work and the strategies for addressing them; and (3) fostering accountability by interrogating how participatory commitments were enacted and sustained throughout a project.

### **6.1 Relationality as a Foundation**

A throughline in our project was relationality: honoring "interdependence, relationships, and connectedness" (Birhane 2021). The notion of relationality is central to many Indigenous and Black feminist epistemologies (Collins 2022; Tynan 2021), which highlight how it supports mutuality and the richness of diverse lived experiences. Relationality formed

the foundation for our project, shaping how responsibility, authority, and decision-making were distributed, and helping to advance DF and ADE principles, including *Challenging Power*, *Elevating Emotion and Embodiment*, and *Grounding Design in Community-Engaged Consensus*. For example, we found that bringing local technical talent into the project was a key way to support FC-K's ownership of the project: A2 added a situated technical perspective alongside FC-K's input and became a potential future technical collaborator with FC-K on its own projects. These impacts stemmed from trusting relationships and team members' mutual accountability. Relationality extended beyond the immediate team as well: presenting at regional venues and organizing public gatherings enabled FC-K to build visibility and new relationships grounded in their local community, seeding future work towards their mission that extends beyond the project's formal scope. Our findings highlight how relationality as a foundational commitment can shift the underlying structure of a partnership, redistributing authority and responsibility in ways that challenge extractive models of participation.

## 6.2 The Many Forms of (and Barriers to) Shared Ownership

From its initiation, our project aimed to share "ownership" across the team, though we found our reflection surfaced diverse ideas of what ownership means and how it manifests in practice. Shared ownership over *decision-making* emerged from a foundation of trusted relationships, taking place in meetings, communication channels, and concrete procedural steps: for example, DCF and FC-K team members coled the developer interviews, and only candidates that FC-K felt were aligned with their mission continued to the technical assessment round. Ownership over the *narrative* of the project was also important, with FC-K team members presenting the work on their own or with DCF team members (including as co-authors on this paper).

Still, we found that structural barriers limited the extent to which we could share ownership over *material resources and technical infrastructure*, such as servers hosting data and tools. FC-K, for instance, does not currently have the capacity or technical staff to maintain its own server, while DCF does—and it is not clear that developing this capacity is of core interest to FC-K's mission. Given this dynamic, we found that transparency (e.g., technical and organizational documentation, open access to the system) served as a key way to still share *technical* ownership with the local developer (A2), and we raise the importance of maintaining documentation that makes tacit technical knowledge explicit.

Ultimately, while DCF owning technical infrastructure was convenient and reduced the burden of full ownership and maintenance for FC-K, we recognize the colonial dynamic that remains. The very nature of the USA-Kenya partnership makes it challenging—if not impossible—to fully overcome this dynamic, and to fully realize the ADE principles to *Assert Data Self-Determination* and *Invest in Local Data Institutions & Infrastructures*. As the project wrapped up, though, FC-K formed a new partnership and acquired funding with a technical university in Kenya that stemmed

from relationships built while sharing the Swahili tools. Perhaps, then, the role of the current partnership is as an imperfect bridge, providing the catalyst and foundation for FC-K to independently build relationships and local technical capacity.

## 6.3 Assessment and Accountability

As we reflect on assessment, we also raise concerns about project sustainability and long-term community benefit. While researchers typically recognize the short-term benefits of research and technology projects (e.g., research insights, publications, novel datasets), communities often rely on delayed benefits (e.g., improved community outcomes over time). How, then, can we ensure that the projects yield sustained benefits for the community? Beyond critical reflexivity, we raise a provocation for future work to define measures of mutual project benefit and mechanisms for monitoring them. Much like Machine Learning Operations (MLOps) provides infrastructure for continuous monitoring of machine learning systems, could a "Community Benefit Operations" framework provide infrastructure for measuring community value, equity, and social impact? Future work could explore what measures of success might be (e.g., the extent to which community partners use a developed system, establish new relationships, or build local technical capacity), how to measure them (e.g., through community feedback mechanisms or audits), and what infrastructure is needed to ensure that feedback is traceable and actionable. More generally, what would it look like for assessments of community benefit to be conducted with as much rigor and care as those of AI systems themselves? Of course, measurement alone should not be treated as a replacement for critical reflexivity, a way to flatten community experiences, or a means to reinforce technosolutionism. Rather, we identify the need for a wide-ranging evaluation toolkit to support the research community to better "question and challenge the consequences of the work we put into the world" (Keyes, Hoy, and Drouhard 2019).

## 7 Conclusion

We describe a justice-oriented research-community partnership between Data Against Feminicide and Femicide Count - Kenya, a Kenyan-based activist organization focused on monitoring and combating femicide. Our work involved adapting two AI-based tools to the Kenyan/Swahili-language context, embedding situated technical capacity-building within the project by hiring a local developer to lead the effort. Using collaborative autoethnography, we reflect on the project, examining the barriers we faced, how we navigated them as a team, and how ownership and control were meaningfully shared. We contribute to a growing body of work on justice-oriented AI, distilling key strategies, such as the more systematic integration of reflexive practice into researchers' evaluation toolkits, to advance these efforts in practice.

## Acknowledgments

We sincerely thank Dr. Kathomi Gatwiri, co-founder of Femicide Count-Kenya, for her much-needed advice and guidance during the project's early stages, Fiona Wanyoni for her help with Swahili translation, and Angela Langat for her assistance in identifying Swahili media sources. We are grateful to Olivia Fiol for her crucial help with project coordination, and to Ray Fong, Nikko Stevens, Wonyong So, Walla Rahama, and Rahul Bhargava for their work and support. We also thank Yujia Gao for facilitating our focus group. Finally, we extend our gratitude to everyone at FC-K and Data Against Femicide, and every supporter of the project whom we have not named individually, but without whose contributions this project would not have been possible.

## References

- Adelani, D.; Abbott, J.; Neubig, G.; D'Souza, D.; Kreutzer, J.; Lignos, C.; Palen-Michel, C.; Buzaaba, H.; Rijhwani, S.; Ruder, S.; Mayhew, S.; Azime, I. A.; Muhammad, S.; Emezue, C. C.; Nakatumba-Nabende, J.; Ogayo, P.; Aremu, A.; Gitau, C.; Mbaye, D.; Alabi, J.; Yimam, S. M.; Gwadabe, T. R.; Ezeani, I.; Niyongabo, R. A.; Mukiibi, J.; Otiende, V.; Orife, I.; David, D.; Ngom, S.; Adewumi, T. P.; Rayson, P.; Adeyemi, M.; Muriuki, G.; Anebi, E.; Chukwuneke, C.; Odu, N.; Wairagala, E. P.; Oyerinde, S.; Siro, C.; Bateesa, T. S.; Oloyede, T.; Wambui, Y.; Akinode, V.; Nabagereka, D.; Katusiime, M.; Awokoya, A.; Mboup, M.; Gebreyohannes, D.; Tilaye, H.; Nwaike, K.; Wolde, D.; Faye, A.; Sibanda, B.; Ahia, O.; Dossou, B. F. P.; Ogueji, K.; Diop, T. I.; Diallo, A.; Akinfaderin, A.; Marengereke, T.; and Osei, S. 2021. MasakhaNER: Named Entity Recognition for African Languages. *ArXiv*, abs/2103.11811.
- Agid, S. 2016. "...it's your project, but it's not necessarily your work...": infrastructuring, situatedness, and designing relational practice. In *Proceedings of the 14th Participatory Design Conference: Full papers - Volume 1*, PDC '16, 81–90. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-4046-5.
- Arnstein, S. R. 1969. A ladder of citizen participation. *Journal of the American Institute of planners*, 35(4): 216–224.
- Barrett, T.; Okolo, C. T.; Biira, B.; Sherif, E.; Zhang, A.; and Battle, L. 2025. African Data Ethics: A Discursive Framework for Black Decolonial AI. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, 334–349. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-1482-5.
- Becker, K.; Vögele, N.; Surmann, H.; Lubos, D.; Rogowski, T.; Yildiz, M. U.; Abdelazim, A.; Koruhama, A. K.; and Winschiers-Theophilus, H. 2025. Challenges in AI image generation of underrepresented groups: the case of the indigenous OvaHimba people. In *Proceedings of the 5th Biennial African Human Computer Interaction Conference*, AfriCHI '25, 338–347. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-1849-6.
- Birhane, A. 2021. Algorithmic injustice: a relational ethics approach. *Patterns*, 2(2).
- Birhane, A.; Isaac, W.; Prabhakaran, V.; Diaz, M.; Elish, M. C.; Gabriel, I.; and Mohamed, S. 2022a. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–8.
- Birhane, A.; Ruane, E.; Laurent, T.; S. Brown, M.; Flowers, J.; Ventresque, A.; and L. Dancy, C. 2022b. The Forgotten Margins of AI Ethics. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 948–958. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-9352-2.
- Cambo, S. A.; and Gergle, D. 2022. Model positionality and computational reflexivity: Promoting reflexivity in data science. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–19.
- Chang, H.; Ngunjiri, F.; and Hernandez, K.-A. C. 2016. *Collaborative autoethnography*. Routledge.
- Cioffi Felice, M.; Feldfeber, I.; Glasserman Apicella, C.; Quiroga, Y. B.; Ansaldo, J.; Lapenna, L.; Bezchinsky, S.; Barriga Rubio, R.; and García, M. 2025. Doing the Feminist Work in AI: Reflections from an AI Project in Latin America. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–18. Yokohama Japan: ACM. ISBN 979-8-4007-1394-1.
- Collins, P. H. 2022. *Black feminist thought: Knowledge, consciousness, and the politics of empowerment*. routledge.
- Cooper, N.; Horne, T.; Hayes, G. R.; Heldreth, C.; Lahav, M.; Holbrook, J.; and Wilcox, L. 2022. A Systematic Review and Thematic Analysis of Community-Collaborative Approaches to Computing Research. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, 1–18. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-9157-3.
- Cooper, N.; and Zafiroglu, A. 2024. From fitting participation to forging relationships: The art of participatory ML. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–9.
- Corbett, E.; Denton, R.; and Erete, S. 2023. Power and public participation in ai. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–13.
- Correa Lucero, H.; and Martens, C. 2025. Colonial structures in AI: a Latin American decolonial literature review of structural implications for marginalised communities in the Global South. *AI & SOCIETY*, 1–17.
- Costanza-Chock, S. 2020. *Design Justice: Community-Led Practices to Build the Worlds We Need*.
- Dantec, C. A. L.; and DiSalvo, C. 2013. Infrastructuring and the formation of publics in participatory design. *Social Studies of Science*, 43(2): 241–264.
- Delgado, F.; Yang, S.; Madaio, M.; and Yang, Q. 2023. The participatory turn in ai design: Theoretical foundations and the current state of practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–23.
- D'Ignazio, C. 2024. *Counting Femicide: Data Feminism in Action*. The MIT Press. ISBN 978-0-262-37801-7.

- D'Ignazio, C.; and Klein, L. F. 2023. *Data feminism*. MIT press.
- Drahota, A.; Meza, R. D.; Brikho, B.; Naaf, M.; Estabillo, J. A.; Gomez, E. D.; Vejnoska, S. F.; Dufek, S.; Stahmer, A. C.; and Aarons, G. A. 2016. Community-Academic Partnerships: A Systematic Review of the State of the Literature and Recommendations for Future Research. *The Milbank Quarterly*, 94(1): 163–214. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0009.12184>.
- Edenberg, E.; and Wood, A. 2023. Disambiguating Algorithmic Bias: From Neutrality to Justice. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 691–704. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-0231-0.
- Ellis, C.; Adams, T. E.; and Bochner, A. P. 2011. Autoethnography: an overview. *Historical social research/Historische sozialforschung*, 273–290.
- Eslami, M.; Fox, S.; Shen, H.; Fan, B.; Lin, Y.-R.; Farzan, R.; and Schwanke, B. 2025. From Margins to the Table: Charting the Potential for Public Participatory Governance of Algorithmic Decision Making. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 2657–2670.
- Feffer, M.; Skirpan, M.; Lipton, Z.; and Heidari, H. 2023. From Preference Elicitation to Participatory ML: A Critical Survey & Guidelines for Future Research. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 38–48. Montréal QC Canada: ACM. ISBN 979-8-4007-0231-0.
- Garcia, P.; and Cifor, M. 2019. Expanding our reflexive toolbox: Collaborative possibilities for examining socio-technical systems using duoethnography. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW): 1–23.
- Haraway, D. 2013. Situated knowledges: The science question in feminism and the privilege of partial perspective 1. In *Women, science, and technology*, 455–472. Routledge.
- Harding, S. 2013. Rethinking standpoint epistemology: What is “strong objectivity”? In *Feminist epistemologies*, 49–82. Routledge.
- Howell, N.; Desjardins, A.; and Fox, S. 2021. Cracks in the success narrative: Rethinking failure in design research through a retrospective trioethnography. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 28(6): 1–31.
- Hutchinson, B.; Louro, C. R.; Collard, G.; and Cooper, N. 2025. Designing Speech Technologies for Australian Aboriginal English: Opportunities, Risks and Participation. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 108–124.
- Iyer, N.; Chair, C.; and Achieng, G. 2021. Afrofeminist Data Futures Report. Report, Pollicy. Accessed: 2026-01-13.
- Kaltenhauser, A.; Stefanidi, E.; and Schöning, J. 2024. Playing with perspectives and unveiling the autoethnographic kaleidoscope in HCI—A literature review of autoethnographies. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–20.
- Kamande, W.; Rheault, C.; Gatwiri, M.; Duchêne, N.; Gachambi, B.; and Jauffret, P. 2025. Capturing Linguistic Diversity in Data Annotation: Learnings From Reviewing AI Generated Content in Swahili and English. In *Proceedings of the 5th Biennial African Human Computer Interaction Conference*, AfriCHI '25, 454–458. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-1849-6.
- Karasti, H. 2014. Infrastructuring in participatory design. In *Proceedings of the 13th Participatory Design Conference: Research Papers - Volume 1*, PDC '14, 141–150. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-2256-0.
- Kasirzadeh, A. 2022. Algorithmic Fairness and Structural Injustice: Insights from Feminist Political Philosophy. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '22, 349–356. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-9247-1.
- Katell, M.; Young, M.; Herman, B.; Dailey, D.; Tam, A.; Guetler, V.; Binz, C.; Raz, D.; and Krafft, P. 2019. An algorithmic equity toolkit for technology audits by community advocates and activists. *arXiv preprint arXiv:1912.02943*.
- Keyes, O.; Hoy, J.; and Drouhard, M. 2019. Human-Computer Insurrection: Notes on an Anarchist HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, 1–13. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-5970-2.
- Klein, L.; and D'Ignazio, C. 2024. Data Feminism for AI. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, 100–112. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-0450-5.
- Leslie, D.; Katell, M.; Aitken, M.; Singh, J.; Briggs, M.; Powell, R.; Rincón, C.; Chengeta, T.; Birhane, A.; Perini, A.; Jayadeva, S.; and Mazumder, A. 2022. Advancing Data Justice Research and Practice: An Integrated Literature Review. Technical report. ArXiv:2204.03090 [cs].
- Maas, J.; and Inglés, A. M. 2024. Beyond Participatory AI. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 932–942.
- Magaña, M. I.; and Shilton, K. 2025. Frameworks, Methods and Shared Tasks: Connecting Participatory AI to Trustworthy AI Through a Systematic Review of Global Projects. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, 2166–2179. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-1482-5.
- Miceli, M.; Schuessler, M.; and Yang, T. 2020. Between subjectivity and imposition: Power dynamics in data annotation for computer vision. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2): 1–25.
- Miceli, M.; Yang, T.; Naudts, L.; Schuessler, M.; Serbanescu, D.; and Hanna, A. 2021. Documenting computer vision datasets: An invitation to reflexive data practices. In

- Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 161–172.
- Mohamed, S.; Png, M.-T.; and Isaac, W. 2020. Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33(4): 659–684.
- Nekoto, W.; Marivate, V.; Matsila, T.; Fasubaa, T.; Fagbohungebe, T.; Akinola, S. O.; Muhammad, S.; Kabenamualu, S. K.; Osei, S.; Sackey, F.; et al. 2020. Participatory research for low-resourced machine translation: A case study in african languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2144–2160.
- Ngaruiya, N.; Donner, J.; Baru, J. K.; and Chege, B. W. 2024. The domestication of AI by Kenyan digital creators. In *Proceedings of the 4th African Human Computer Interaction Conference, AfriCHI '23*, 71–75. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-0887-9.
- Patton, D. U.; Frey, W. R.; McGregor, K. A.; Lee, F.-T.; McKeown, K.; and Moss, E. 2020. Contextual analysis of social media: The promise and challenge of eliciting context in social media posts with natural language processing. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 337–342.
- Pierre, J.; Crooks, R.; Currie, M.; Paris, B.; and Pasquetto, I. 2021. Getting Ourselves Together: Data-centered participatory design research & epistemic burden. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, CHI '21*, 1–11. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-8096-6.
- Png, M.-T. 2022. At the tensions of south and north: Critical roles of global south stakeholders in AI governance. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1434–1445.
- Queerinaï, O. O.; Ovalle, A.; Subramonian, A.; Singh, A.; Voelcker, C.; Sutherland, D. J.; Locatelli, D.; Breznik, E.; Klubicka, F.; Yuan, H.; et al. 2023. Queer in AI: A case study in community-led participatory AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1882–1895.
- Rapp, A. 2018. Autoethnography in human-computer interaction: Theory and practice. In *New directions in third wave human-computer interaction: Volume 2-Methodologies*, 25–42. Springer.
- Ricaurte, P.; and Zasso, M. 2023. AI, ethics and coloniality: A feminist critique. *Cebral-Loureda, Manuel (Hg.): What AI can Do. Milton Park*, 53–72.
- Scheuerman, M. K.; and Brubaker, J. R. 2024. Products of positionality: How tech workers shape identity concepts in computer vision. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Schroeder, H.; Pareek, A.; and Barocas, S. 2025. Disclosure without Engagement: An Empirical Review of Positionality Statements at FAccT. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 1195–1210.
- Seger, E.; Ovadya, A.; Siddarth, D.; Garfinkel, B.; and Dafoe, A. 2023. Democratising AI: Multiple Meanings, Goals, and Methods. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society, AIES '23*, 715–722. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-0231-0.
- Septiandri, A. A.; Constantinides, M.; Tahaei, M.; and Quercia, D. 2023. WEIRD FAccTs: How western, educated, industrialized, rich, and democratic is FAccT? In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 160–171.
- Sloane, M.; Moss, E.; Awomolo, O.; and Forlano, L. 2022. Participation is not a design fix for machine learning. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–6.
- Suresh, H.; Movva, R.; Dogan, A. L.; Bhargava, R.; Cruxen, I.; Cuba, A. M.; Taurino, G.; So, W.; and D'Ignazio, C. 2022. Towards Intersectional Feminist and Participatory ML: A Case Study in Supporting Femicide Counterdata Collection. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, 667–678. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-9352-2.
- Suresh, H.; Tseng, E.; Young, M.; Gray, M.; Pierson, E.; and Levy, K. 2024. Participation in the age of foundation models. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1609–1621.
- Tseng, E.; Young, M.; Le Quéré, M. A.; Rinehart, A.; and Suresh, H. 2025. "Ownership, Not Just Happy Talk": Co-Designing a Participatory Large Language Model for Journalism. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 3119–3130.
- Tynan, L. 2021. What is relationality? Indigenous knowledges, practices and responsibilities with kin. *Cultural geographies*, 28(4): 597–610.
- Walker, R. D.; Cruse, B.; Cora, A.; Rogers, K.; D'Ignazio, C.; Brion-Meisels, G.; and Breazeal, C. 2025. Archaeology of Self: Reflexivity in Data Activism to Address Systemic Injustices. In *Proceedings of the 5th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 28–41. Pittsburgh USA: ACM. ISBN 979-8-4007-2140-3.
- Young, M.; Ehsan, U.; Singh, R.; Tafesse, E.; Gilman, M.; Harrington, C.; and Metcalf, J. 2024. Participation versus scale: Tensions in the practical demands on participatory AI. *First Monday*.
- Young, M.; Katell, M.; and Krafft, P. 2022. Confronting power and corporate capture at the FAccT Conference. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1375–1386.