Chapter 6

The AI Impact Engagement Report Kit

Tamara Kneese, Meg Young, Ranjit Singh, Briana Vecchione, Emnet Tafesse and Jacob Metcalf

Background and Context

This method is the result of eighteen months of piloting and experimentation at the Data & Society Research Institute's Algorithmic Impact Methods Lab, which we launched in May 2023 to better understand how impacted communities might drive the terms of impact assessment. Impact assessment, broadly, is a structured process that evaluates both the potential and actual harms of algorithmic systems. Our method, called "Impact Engagement," is an approach to algorithmic impact assessment that elicits the diverse perspectives and expertise of impacted communities. Impacted communities are the group(s) of people whose well-being will be affected by an Artificial Intelligence (AI) system: people who are local residents in the same area where an AI system is deployed, people who might use a system, or (more importantly) those whose well-being or material livelihoods will be changed by a system without being end-users themselves.

Eliciting community input can flag developers' flawed assumptions, better understand their needs on the ground, anticipate what could go wrong, or redirect investment toward more pressing problems and opportunities. An impact assessment process that is missing community engagement can partially address these needs but is missing a fundamental factor for meaningful accountability: relationships between the people developing or deploying the system and those whose well-being is at stake. Accountability relies on these relationships: that is, it requires that those impacted by a system can make demands of the organization deploying it. Meaningful accountability also requires an organization to produce an "account" or assurance to justify the outcome of its work. Whereas domains like food and drug safety have established oversight authorities and judicial means of appealing harm, AI systems operate under still-incipient accountability regimes that are not yet well established in law.

Our method—the AI Impact Engagement Report (AIER)—centers on relationship building with the intention of creating stronger lines of

DOI: 10.4324/9781003511588-7

accountability. AIER describes which stakeholder groups were considered, which organizations were understood to represent them, how each group and individual was invited to participate, the frequency and scope of their feedback, and if and how that feedback shaped the system's design and deployment. It includes a set of heuristics for whom to invite to engage, guidelines for the scope and agenda of the meeting(s), and instructions for how to translate community feedback into actionable insights for system design.

AIERs reflect an ongoing conversation. This conversation can be broadly divided into two kinds of ongoing debates: first, what is "impact" and how to evaluate it; and second, who should be involved in "engagement" and how to organize it? Starting with the first debate, impact is a way to evaluate potential consequences of algorithmic systems in an effort to make them legible to assessment practices and regulatory control (Metcalf et al. 2021; Selbst 2021). Impact is often anticipatory by design, but must account for the potential of unintended consequences that emerge from concrete harms experienced by individuals and groups as they contend with the operation of algorithmic systems (Tufekci 2015; Shelby et al. 2023; Data & Society and Center for Democracy & Technology 2024). Given that safety of algorithmic systems can never be assumed or guaranteed and there are fundamental limits of knowledge in anticipating the potential ways in which such systems might fail, articulating a clear relationship between impact and harm is the central concern that animates this debate (Moss et al. 2021).

An effective strategy to frame this relationship between impact and harm is articulating who is impacted by these systems and engaging with these communities, which in turn shapes the contours of the second debate. This debate around "engagement" often focuses on who counts as a participant in algorithmic systems as defined, and how this relates to the histories and political implications of participation as an experience (Kelty 2020), and the unevenly distributed challenges of claiming autonomy, power, and agency among data subjects (Kennedy 2018; Couldry and Mejias 2018; Kasy and Abebe 2021)—people who are "both resources and targets" (Ziewitz and Singh 2021, 2) for algorithmic systems. The uneven distribution of these challenges can be mapped around well-recognized intersections of race, gender, caste, class, ability, occupation, and power; engagement on such challenges requires valuing voices and standpoints (Haraway 1988) of impacted communities who inhabit these intersections (Birhane, Isaac, et al. 2022; Delgado et al. 2023; Young et al. 2024). Those who are excluded from social and political privilege have a different standpoint grounded in their own experiences of how algorithmic systems do not work for them (Birhane, Ruane, et al. 2022). What they see is not just

dysfunction; they also see the ordinary ways in which these systems and the infrastructures they support become daily objects of concern (Anand, Gupta, and Appel 2018). While impacted communities may not be able to account for algorithms in terms of statistics and code, they still develop an idea of how these systems work and impact their lives (Bucher 2017; Shen et al. 2021). The moral case for engaging with them and valuing their standpoint is that they have authoritative and unique knowledge of the workings of algorithmic systems and thus, they must be afforded substantial power in determining how these systems are designed and operated.

In conceptualizing impact engagements, we address both the debates by centering the experiences of impacted communities as the central perspective that binds impacts to harms of algorithmic systems. In documenting such engagements in AIERs, we offer ethnographically-informed methodological strategies (Seaver 2017; Christin 2020) to manage the challenges of engaging with lived experiences of impacted communities who may not be fully aware of their position and who are themselves figuring out their position vis-à-vis algorithmic systems (Ziewitz and Singh 2021; Kear 2022). Moss et al. (2021) provide a history of impact assessment in other settings. Whereas prior impact assessment approaches, such as privacy impact assessment, produce static documentation of system risks and how these have been mitigated (Moss et al. 2021), AIER is ongoing and updated with each respective exchange with communities affected. One intended use is within state and local government, where an agency (such as a Department of Transportation) is considering adopting a new AI system and is able to document its outreach to residents and the feedback received. Another application is for AI products or services created by companies or organizations, as a set of guidelines to follow when performing stakeholder engagement as part of community feedback sessions or user testing focus groups.

The AI Impact Engagement Report Kit Components

- **Partnership Assessment Template (Component 6.1)**
  A systematic way of determining whether to partner with a developer or deployer organization to conduct impact engagements, or whether any community feedback is unlikely to have a substantive role in shaping the product or deployer organization.
- **Email Outreach Template (Component 6.2) & Informed Consent Template (Component 6.3)**
  A set of "turn-key" email templates to send to organizations to invite them to work with you and informed consent forms for when you sit down to meet.

- **Workshop Run-of-Show Guide (Component 6.4)**
  Suggested "run-of-show" document for timing and structuring the discussion.
- **AI Engagement Report Template (Component 6.5)**
  A template for sharing the results of your engagement
- **Additional Materials Needed**
  - Zoom account and access (for online workshops)
  - Transcription services
  - Presentation slides for workshop run-of-show
  - Snacks or a meal (for in-person workshops)
  - Laptop or pen and paper for taking notes
  - Honoraria/gift cards

The AI Impact Engagement Report Kit Guide

Step 1: Identify Potential AI Systems and Partners

The first step is to identify potential systems to evaluate. Ideally, research partnerships and evaluations of AI systems would happen pre-deployment, when the technology is still in the design phase. But that is unlikely given the nature of technology development, which happens much faster than the pace of ethnographic research. This requires establishing trust with developers (such as a small startup company or a large, enterprise-scale corporation) or deployers (such as city governments or healthcare providers) within different kinds of organizations. Because algorithmic impact assessments are not yet legally mandated, this means that developers and deployers must voluntarily invite you into their systems and internal processes. While you are working with some members of a company, government, or other organization, you are also treating their workplace as a research site. This means that you need to have at least one or more allies on the inside who are willing to work with you and internally advocate for your presence.

Step 2: Use the Partnership Assessment Template

There are many risks attached to partnering with a developer or deployer, and it is vital that you, as a researcher and advocate, understand the full weight of your choice to align yourself with the developers or deployers of an AI system. The Partnership Assessment Template (Component 6.1) is designed to help you and your larger research team, if applicable, determine whether or not your values and goals are compatible with the partner. Some questions you and/or your team should take the time to answer include: Do we have the funding/resources to take on this project?

Partnership Assessment This framework helps to explicitly articulate the nature of the partnership, its values and mission alignment, what is at stake, and the scope of what the collaboration entails.		
Step 1: In what ways does this potential partner share our values and mission?		
Factors	*Consideration*	*Assessment*
Independence: Are researchers independent from the collaborating organization in terms of funding and publications on the research?		(✓ 0)(… 0)(× 0)
Integrity: Do researchers feel confident in standing by the research process/methods and results?		(✓ 0)(… 0)(× 0)
Equity: How is decision-making in the research process equitable?		(✓ 0)(… 0)(× 0)
Mutuality: How are elements of this research organized to ensure that the exchange of ideas and resources is reciprocal?		(✓ 0)(… 0)(× 0)
Creativity: Does this collaboration achieve a result that would not have been possible without it?		(✓ 0)(… 0)(× 0)
Mission fit: In what ways is this partner currently working to advance just and equitable design, use, and governance of technology?		(✓ 0)(… 0)(× 0)
Step 2: What is at stake?		
Factors	*Consideration*	*Assessment*
Does working with this potential partner carry any reputational risks or present risks to existing relationships?		(✓ 0)(… 0)(× 0)

Component 6.1 Partnership Assessment Template.

Do the potential benefits outweigh the costs? If so, how might we mitigate the risks?		(✓ 0)(… 0)(× 0)
Does the potential partner organization have clearly stated goals in taking on this collaboration? And do they have a sense of what they are contributing?		(✓ 0)(… 0)(× 0)
Does the organization have clearly stated goals in taking on this collaboration? And are we clear on what we are contributing?		(✓ 0)(… 0)(× 0)
Are these goals aligned or complementary to each other? And what is the added value of carrying out this work through a partnership?		(✓ 0)(… 0)(× 0)
Step 3: What does the scope of this collaboration entail?		
Factors	*Consideration*	*Assessment*
Do we have the technical and human resources needed to achieve the goals of this collaboration?		(✓ 0)(… 0)(× 0)
Do we have the funding/resources/ capacity to take on this partnership? And what are their resources?		(✓ 0)(… 0)(× 0)
Are financial resources being exchanged or joint funds being procured? How does it comply with our institution's gift acceptance policy?		(✓ 0)(… 0)(× 0)
Do we have pre-existing relationships with individuals based in this partner institution?		(✓ 0)(… 0)(× 0)
If there is a collaborative research project involved, is it attached to an IRB process?		(✓ 0)(… 0)(× 0)

Component 6.1 (Continued).

Do the terms and conditions of the research include agreement on academic research outputs and support for independence in writing?		(✓ 0)(… 0)(× 0)
Are there any other considerations that might pertain to the specific nature of this potential partnership?		(✓ 0)(… 0)(× 0)
Final Decision:		
(✓ 0) Go ahead | (… 0) More info needed | (× 0) Pass		

Component 6.1 (Continued).

What kind of budget and labor power is needed to complete it? Does working with this potential partner carry any reputational risks or present risks to existing relationships? Are we facilitating an extractive or abusive relationship between a developer and a community, or can we trust that the developer will take the community feedback seriously and revise or halt their proposed technology?

The following set of guiding questions should be used to assess the value, risks, and implications of a partnership in order to confidently go ahead, define whether more information is needed, or decide not to proceed. The template should be filled in by the research team leader, exploring the partnership in consultation with their team, strategy, and engagement team members, and any implicated colleagues across the organization to arrive at a final decision.

In practice, our team uses Google Docs voting chips to collaboratively assess each factor. These chips allow team members to register their judgment using a visual interface—such as ✓ (go ahead), … (more info needed), or × (pass)—and to see at a glance where there is agreement or divergence. While this interactive feature cannot be reflected in the printed version of the table, readers are encouraged to use checkmarks or initials in the assessment column to capture group consensus or divergent viewpoints. The aim is not to tally scores but to prompt structured discussion, surface tensions, and guide a deliberative decision about whether to proceed with the partnership.

Step 3: Determine Impacted Communities and Conduct Outreach

After assessing your potential partner and deciding to move forward with an impact engagement pilot project, you need to determine which communities will be impacted by the AI system. This means end users in some

cases, but they may also encompass members of a particular district or city if the AI system is being deployed in a city environment. The AI system may also have downstream impacts on non-users and the communities, caregivers, or loved ones of users, and can have labor impacts on particular groups of workers. The system may even impact people who are unaware they are affected, such as when an AI system is deployed in a city without sufficient public notice.

The next step is for you to decide, given the likely impacted communities, which community-based organizations should be approached to participate in pilot workshops. Prioritize groups that are least likely to already be involved in development of the technology or decision-making around it, or those who are most likely to be negatively impacted (or have concerns about harm) based on previous research about that class of technologies. You will likely have to choose relevant groups under time and resource constraints, and the most relevant groups may be unable or unwilling to participate. For example, advocacy groups may not have the capacity to join your workshops because they have more pressing matters to attend to. Be sure to take note of declined invitations to participate and to "learn from the no."

Use the Outreach Email Template (Component 6.2) to email the nonprofits or other grassroots/advocacy organizations that are relevant to your research site. It may take several attempts before you get a response, and you should be very clear about the consent process and terms of the research relationship. Use the Informed Consent Template (Component 6.3) to provide detailed information in plain language about your research aims. Secure consent forms from all participants.

Outreach Email Template

Dear [Organization Staff Member],

My name is [Name]. I work at [Institution] on a project to inform how [Deployer] uses new technologies and promotes residents' digital rights.

We deeply admire [Organizations]'s work to connect and support the local [Community Name] community; we are working with the [Deployer] to invite groups like yours to elicit community members' perspectives on what the [Deployer] should consider when evaluating a new AI-based technology for applications related to housing and people experiencing homelessness.

Component 6.2 Outreach Email Template.

If you were interested in collaborating with us on this project, we would like to:

- Collaborate on organizing one or more ninety-minute conversations with up to eight total community members and staff on an AI-based technology and community members' questions, concerns, and expectations.
- Ask for two hours of a staff member's time over the coming month or two for the purpose of setting up these sessions.

We are hoping to complete a session in [General Time Period] and could hold it as early as [Specific Date Options].

People attending the sessions would receive a [Detailed Description of Honorarium, e.g., $100 Visa gift card] honorarium for each session; we would also provide $[Amount] stipends payable to your organization to acknowledge your time and expertise.

Do you have thirty minutes in the coming week or so to talk about this project idea?

Warmly,
[Your Name]

Component 6.2 (Continued).

Informed Consent Agreement for Working Group Participants

Purpose of This Project

You are invited to participate in a working group convened by [Institution]. The goal of this project is to ask [Deployer] residents to share questions, concerns, and expectations about technologies that the [Deployer] is considering in its [Specific Program].

Today's working group is focused on your:

- Questions, expectations, and concerns about city-owned technologies.
- Thoughts about a draft document that tries to answer these questions.

Component 6.3 Informed Consent Template.

This project will have two outputs:

- A document called an Impact Engagement Report that will be delivered to the [Deployer] that summarizes community input about [System]; and
- Academic paper(s) about how community groups' questions and concerns about technology can lead to conversations about technology oversight.

What You Will Do in the Working Group

If you decide to participate in this project, we will discuss hopes and concerns about technologies the city uses, and questions for the [Deployer] about how these technologies are designed or used. Your participation is voluntary: please feel free to skip questions, topics, or to leave the room at any time.

Time Required

The working group session is ninety minutes long. There will be future working groups that you will be invited to participate in—but these are optional.

Honorarium Payment

You will be paid $100 for your time and expertise. If you decide to leave the working group at any time during today's session or later, you will still receive this payment.

Confidentiality

Your identity will be completely anonymous; we will refer to the organization by name. The working group session will be audio-recorded and transcribed. Only our team will have your information; we will anonymize them in any public writing.

How to Withdraw

We care about your comfort, and you are free to stop participating at any time. If you decide not to participate or withdraw later, it will

Component 6.3 (Continued).

not be a problem in any way; you can contact us or [Institution] to ask for your input not to be used.

Questions about This Project?

You can ask questions about this consent form or the project at any time. Before we go any further, do you have any questions regarding the project, your participation, or payment?

_____

(Optional) If You Are Willing to be Photographed, Please Select Below

I consent to having my photo taken for the purpose of public presentations and reports.

Acknowledgment

I understand I can withdraw my input before the project is published and that I would still receive my $100 honorarium payment for my time.

Consent

We are so grateful for your input on this project. Please sign below to confirm your willingness to participate.

"I consent to participating in this study. I understand that **my identity will be anonymous**, but using **my organization's name with my quotes is OK**."

- Print Name
- E-mail address:
- Signature:
- Date:

If You Have Questions about This Study, Please Contact

- [Name]
- [Organization]
- [Email]
- [Phone]

Component 6.3 (Continued).

Step 4: Produce Workshops for Community Engagement

If financially possible and permissible according to local laws and regulations, you should offer an honorarium to the organization itself to compensate them for their time and labor, and you should also offer individual honoraria to attendees that make sense for your organization's budget. If the event is in person, you should provide food and do everything to ensure that caregivers, disabled people, and others who might experience barriers to participating are able to attend. If you are on Zoom, make sure that everyone participating has access to a Zoom account and is able to connect. Whether online or in-person, hire interpreters and captioners as needed.

Use the discussion guides to create a run-of-show for the community engagement workshops. Organizers can begin a Google doc or other editable, shareable document that allows you and the research partner to collaborate, which will be printed out for community members. Your research partner may have different goals for these workshops, so please consult with the developer/deployer on the workshop run-of-show and leave space for their needs. Also, leave space after workshops for participants to add feedback via email or other channels, because some of this informal feedback outside of the workshop itself can be useful while formulating the eventual impact engagement report.

Determine if there are particular kinds of background information that would be helpful to the community you are engaging with. For example, would community members like to learn about a specific aspect of an AI system? Do they have any questions about the technology and its applications? What kinds of technologies or social changes would they like to see instead?

Prepare to deliver a short presentation to do some basic table setting at the start of the workshop, providing some educational background on its applications and potential harms. Your job is to make community members feel welcome and supported so they can offer meaningful feedback. Ask questions and facilitate a dialog among community members, managing differences of opinion and any conflicts that arise. Record the session and take notes so you can revisit the feedback they provide. You will rely on these materials, along with your own ethnographic observations, to write your impact engagement report.

Ideally, you will have at least two different workshops with each impacted community. We have found that it is important to use the first session for table setting and introducing the technology, which makes people feel comfortable, and then deepen the discussion in the second session. It is especially fruitful when participants are able to directly engage with a system and observe its effects in real-time.

Workshop Run-of -Show Agenda

1:00–1:05 PM: Welcome and Introduction (five minutes)

- Facilitator Introduction:
  - Brief self-introductions and background
  - Introduction to [Organization(s) Involved]
- Workshop Objectives:
  - Overview of [The Technology] and its purpose(s)
  - Goals of the study
  - Importance of participant feedback
- Purpose of the Workshops:
  - Discuss the purpose of the workshops[1]
  - Explain how the conversation will contribute to [The Organizations, The Product, and The Broader Field].

1:05–1:10 PM: Participant Introductions (five minutes)

- Quick round of participant introductions
- Share reasons for participating in the study (optional)

1:05–1:15 PM: Overview of [The Technology] (ten minutes)

- Intended Functions
  - Main functions and features
  - Ongoing and future product developments
- Guidance for Use:
  - Guide participants to launch/use the product [if applicable]
  - Troubleshoot any setup issues

1:15–1:35 PM: Interactive Session with [The Technology] (fifteen to twenty minutes)[2]

- Breakout Group Activity Models:
  - *Role-Playing Scenarios:* Participants use the technology to simulate real-world interactions (e.g., using a chatbot for customer service training or support)—mainly to assess how it handles varied/unexpected inputs.**Exploratory Testing:** Participants engage with the technology in unconventional ways (e.g., creating

Component 6.4 Workshop Run-of-Show Guide.

edge cases or challenging prompts) to observe and document unintended behavior or biases.
- *Ethical Implications Review:* Participants brainstorm and write prompts that test the AI's responses to ethical dilemmas or sensitive issues, analyzing how the technology navigates morally complex situations.
- *User Journey Mapping:* Participants interact with the technology as different user personas (e.g., users with accessibility needs) to identify strengths and weaknesses in user experience, focusing on issues around inclusivity and accessibility.
- *Societal Impact Challenges:* Participants write prompts designed to explore how the technology affects particular demographics or user groups (e.g., marginalized communities), documenting the social and cultural implications of its responses.
- *Scalability Testing:* Participants push the technology's functionality by engaging in high-volume activities (e.g., sending multiple complex queries in quick succession) to investigate its capacity for handling extensive or simultaneous input. [When using prompts that relate to harming users]

1:35–1:55 PM: Group Discussion and Feedback

- Feedback and Discussion Models:
  - *Round-Robin Sharing:* Each group shares their main findings and reflections (note: for smaller groups only).
  - *Facilitated Group Discussion:* The facilitator leads an open discussion with questions like, "What surprised you the most?" or "What challenges did you face?"
  - *Strengths and Weaknesses Review:* Participants use examples from their activities to highlight the most effective aspects of the technology and areas that need improvement.
  - *Thematic Summarization:* Groups summarize their discussions around key themes (e.g., ethical concerns) and share them with the larger group.
  - *Interactive Polling:* The facilitators enact live polling or voting tools to collect participants' opinions on specific observations and/or themes.
  - *Written Reports/Artifact Sharing:* Groups submit brief written summaries or screenshots of their interactions with the technology.

Component 6.4 (Continued).

1:55–2:00 PM: Conclusion and Next Steps (five minutes)

- Summary:
  - Recap key points and takeaways
- Next Workshop:
  - Provide details about the next session (if applicable)
  - Encourage participants to continue experimenting with the technology (if applicable)
  - Give participants contact information to follow up with organizers or provide necessary artifacts (e.g., screenshots, written experiences, etc.).
- Thank You:
  - Thank participants for contributing to the study

Component 6.4 (Continued).

Step 5: Analyze Findings and Write Engagement Report

After you have completed your workshop sessions, review any transcripts, notes, or recordings to discern emergent themes. Write up your findings using the AI Engagement Report Template (Component 6.5). Share your findings with your research partner and determine how to share any relevant findings with participants.

Algorithmic Impact Engagement Report V1

1. Technology to be considered and proposed use cases
2. Current status/timeline
3. Reason(s) this technology is being considered
4. Which digital rights are relevant to this technology? (Privacy, accountability, equity, fairness)
5. What communities in your city are especially likely to have digital rights concerns? Please primarily list those who are not well-represented in policymaking to date.
6. For each community or stakeholder group listed above, what are some prominent local community-based organizations or advocacy groups that serve them? Invite several of these organizations to meet with your team, including staff and residents,

Component 6.5 AI Engagement Report Template.

to discuss the technology under consideration. Sample meeting agendas are attached below.

   a. [While maintaining anonymity, describe who didn't want to meet and why]

7. Provide the date, point of contact, and a summary of the conversation with each community organization, including general reactions, specific hopes and expectations, or concerns provided for this technology.
8. Re-state key community input in the form of questions to answer about this technology:

   a.
   b.
   c.

9. Answer the questions in (8) above. Where applicable, include them in the evaluation framework or other public-facing documentation about this technology. Where there is currently not enough information to answer a question, indicate how it will be addressed in the future.
10. Summarize all anticipated impacts that community members identified from this technology (both positive and negative).
11. Were there any overlapping concerns from different communities, or did any community members disagree with each other or with the concerns of other community organizations?
12. Did community members express any other concerns or unmet needs during the engagement?
13. What is the next anticipated touchpoint with each of the community partner organizations? How often will you check in with community partners?

    a. Follow-up meeting 1
    b. Follow-up meeting 2
    c. Follow-up meeting 3

On completion of this engagement report, please send it to your staff point of contact at the community partner organizations and ask for clarifying edits as needed.

Component 6.5 (Continued).

Step 6: Relationship Maintenance

Follow up with participating organizations and with individual participants about honoraria. Provide your research partner with an overview of your findings and arrange any necessary follow-up discussions. You may decide that you need to conduct more research or conduct further impact engagement workshops. Ensure that your research partner is able to maintain contact with the community-based organizations in case this relationship is beneficial to both parties.

Lessons Learned

Our team used this method to pilot impact assessment processes on three specific systems: a computer vision pilot by a municipal government, an LLM-based (Large Language Model) chatbot created by a nonprofit, and an LLM-based chatbot created by a private company. Impact engagements are not an exact science. Some observations and questions that emerged from doing this work include broader themes connected to the ethics of applied research and our proximity to power as researchers.

- Impact engagements are a matter of logistics; they are shaped by budgets, managerial priorities, time constraints, staff turnover, and precarity. How do you build meaningful relationships and trust over time? A lack of resources within organizations means that there is also a lack of continuity or institutional knowledge and a general inability to see something through various stages of engagement and development.
- How do you get buy-in from leadership? The impact engagement process depends not only on fostering trust with community-based organizations but also with the key internal stakeholders involved with the product development or procurement/deployment process. Working with internal allies, we developed a value proposition that showed how our work would improve relationships between developers/deployers and communities and mitigate risk. While our real goal was to empower communities, framing the business case of algorithmic impact assessments and demonstrating the value of our work was an essential part of getting access to systems and internal organizational structures.
- How do you know which communities to engage with? We focused on the most marginalized groups—asking, who is most likely to be harmed by the technology? We found that mitigating risks for the most vulnerable will help everyone. Communities are not monolithic, and different communities might have opposing views.
- It is useful for participants to have as much exposure to a system as possible before the scheduled workshops. Providing them with access ahead of the workshop itself and encouraging them to experiment with

the system will result in a degree of comfort, which will make the workshop run more smoothly. It can be a struggle to get developers on the same page as researchers, in that researchers would rather provide background information to participants before initiating the workshop process, whereas developers are more comfortable with having users test products without any experience with the system. In one case, a product dramatically shifted twenty minutes before our first participatory workshop, leaving us scrambling to arrange the parameters of the engagement.

- Balancing positive relationships with developers/deployers and providing accurate (and often critical) feedback is a craft. We need to build trust with the internal staff at particular research sites and are sometimes in positions where we need to tell them things they don't want to hear. We also need to build trust with community-based organizations and other stakeholders whom we are bringing into the process. Sociotechnical research takes a long time and can be emotionally taxing.
- In a number of case studies, participants had concerns about ensuring safety and privacy in their interactions with algorithmic systems, including chatbot interactions with sensitive information or with computer vision technologies that might surveil them in their everyday activities. There is also a gap between what a technology is imagined to do and its actual affordances. Sometimes communities might wonder why AI is being used in the first place—is an automated system the best use of resources?
- As a team, we are navigating the challenge of observing internal processes and facilitating community engagements without having direct influence and often without having a final say in outcomes. One of our goals is to strengthen connections between developers or deployers and impacted communities to help create more space for accountability, contestation, and refusal. If we hold one or two workshops with community-based organizations, even if the community's recommendations are taken up by the technical team, that should not be the end of the story. We want to ensure that communities are part of the process in the long term and are able to make decisions about technologies that will affect their lives. In short, the nature of the contribution that communities can make is directly related to the processes employed to solicit their contribution, which brings us back to our first observation that impact engagements are a matter of logistics.

Notes

1 If the workshops follow a sequence with different goals that build on each other, make that distinction upfront so participants understand their roles and responsibilities for each session.
2 If participants have already engaged with the product, this step may be skipped, and the workshop can start with discussion and feedback.

References

Anand, Nikhil, Akhil Gupta, and Hannah Appel. 2018. *The Promise of Infrastructure*. Duke University Press.

Birhane, Abeba, William Isaac, Vinodkumar Prabhakaran, et al. 2022. "Power to the People? Opportunities and Challenges for Participatory AI." In *EAAMO '22: Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. Association for Computing Machinery. https://doi.org/10.1145/3551624.3555290.

Birhane, Abeba, Elayne Ruane, Thomas Laurent, et al. 2022. "The Forgotten Margins of AI Ethics." In *FAccT '22: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery. https://doi.org/10.1145/3531146.3533157.

Bucher, Taina. 2017. "The Algorithmic Imaginary: Exploring the Ordinary Affects of Facebook Algorithms." *Information, Communication & Society* 20 (1): 30–44. https://doi.org/10.1080/1369118X.2016.1154086.

Christin, Angèle. 2020. "The Ethnographer and the Algorithm: Beyond the Black Box." *Theory and Society* 49 (5–6): 897–918. https://doi.org/10.1007/s11186-020-09411-3.

Couldry, Nick, and Ulises A Mejias. 2018. "Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject." *Television & New Media* 20 (4): 336–349. https://doi.org/10.1177/1527476418796632.

Data & Society, and Center for Democracy & Technology. 2024. "Ensuring 'AI Safety' Begins with Addressing Algorithmic Harms Now." *Data & Society*. March 18, 2024. https://datasociety.net/announcements/2024/03/18/ensuring-ai-safety-begins-with-addressing-algorithmic-harms-now.

Delgado, Fernando, Stephen Yang, Michael Madaio, and Qian Yang. 2023. "The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice." In *EAAMO '23: Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. Association for Computing Machinery. https://doi.org/10.1145/3617694.3623261.

Haraway, Donna. 1988. "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective." *Feminist Studies* 14 (3): 575–599. https://doi.org/10.2307/3178066.

Kasy, Maximilian, and Rediet Abebe. 2021. "Fairness, Equality, and Power in Algorithmic Decision-Making." In *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery. https://doi.org/10.1145/3442188.3445919.

Kear, Mark. 2022. "The Moral Economy of the Algorithmic Crowd: Possessive Collectivisim and Techno-Economic Rentiership." *Competition & Change* 26 (3–4): 467–486. https://doi.org/10.1177/1024529421990496.

Kelty, Christopher M. 2020. *The Participant: A Century of Participation in Four Stories*. 1st Edition. University of Chicago Press.

Kennedy, Helen. 2018. "Living with Data: Aligning Data Studies and Data Activism through a Focus on Everyday Experiences of Datafication." *Krisis: Journal for Contemporary Philosophy* 38 (1): 18–30. https://doi.org/10.21827/krisis.38.1.37184.

Metcalf, Jacob, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, and Madeleine Clare Elish. 2021. "Algorithmic Impact Assessments and Accountability: The Co-Construction of Impacts." In *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery. https://doi.org/10.1145/3442188.3445935.

Moss, Emanuel, Elizabeth Anne Watkins, Ranjit Singh, Madeleine Clare Elish, and Jacob Metcalf. 2021. "Assembling Accountability: Algorithmic Impact Assessment for the Public Interest." *Data & Society*. June 29, 2021. https://datasociety.net/library/assembling-accountability-algorithmic-impact-assessment-for-the-public-interest.

Seaver, Nick. 2017. "Algorithms as Culture: Some Tactics for the Ethnography of Algorithmic Systems." *Big Data & Society* 4 (2): 1–12. https://doi.org/10.1177/2053951717738104.

Selbst, Andrew D. 2021. "An Institutional View of Algorithmic Impact Assessments." *Harvard Journal of Law & Technology* 35 (1): 117–191. https://jolt.law.harvard.edu/volumes/volume-35.

Shelby, Renee, Shalaleh Rismani, Kathryn Henne, et al. 2023. "Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction." In *AIES '23: Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery. https://doi.org/10.1145/3600211.3604673.

Shen, Hong, Alicia DeVos, Motahhare Eslami, and Kenneth Holstein. 2021. "Everyday Algorithm Auditing: Understanding the Power of Everyday Users in Surfacing Harmful Algorithmic Behaviors." *Proceedings of the ACM on Human-Computer Interaction* 5 (CSCW2): 433:1–433:29. https://doi.org/10.1145/3479577.

Tufekci, Zeynep. 2015. "Algorithmic Harms beyond Facebook and Google: Emergent Challenges of Computational Agency." *Colorado Technology Law Journal* 13 (2): 203–218. https://ctlj.colorado.edu/wp-content/uploads/2015/08/Tufekci-final.pdf.

Young, Meg, Upol Ehsan, Ranjit Singh, et al. 2024. "Participation versus Scale: Tensions in the Practical Demands on Participatory AI." *First Monday* 29 (4). https://doi.org/10.5210/fm.v29i4.13642.

Ziewitz, Malte, and Ranjit Singh. 2021. "Critical Companionship: Some Sensibilities for Studying the Lived Experience of Data Subjects." *Big Data & Society* 8 (2): 1–13. https://doi.org/10.1177/20539517211061122.

Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com