Accountability Case Labs

Founder & Lead: Borhane Blili-Hamelin (he/him)
Contact: borhane.blilihamelin@gmail.com; https://borhane.xyz/
(**NEW**) ACL Website! **https://accountabilitycaselabs.xyz**

Roadmap & Milestones

This document details the project roadmap, goals, and milestones for Accountability Case Labs. During our first year, we were part of the Mozfest Trustworthy AI Working Group. We are a collaborative, open community for algorithmic accountability practitioners. We are currently volunteer run. We aim to release all of our outputs under CC BY 4.0. This document is a work in progress. We warmly welcome input (use the comment function!)

Folks interested in contributing as participants in the project can fill this interest form, as well as reach out to Bo by email, on Linkedin or on *Twitter*.

Join us

- Interest form (stay in touch or contribute)

Project Updates!

- **September 2022**
  - **We have an official website! https://accountabilitycaselabs.xyz**
  - For now it's quite barebones, but it's setup to grow as we write new content. Over time, we hope to turn it website into a resource for people interested in remixing and adapting our insights!
  - Stay tuned for new ACL workshops in Fall 2022!
- **July 2022**
  - We are now in the process of gathering input from project contributors, & of making explicit our governance structure moving forward.
  - On **July 5**, we were invited to speak to Mozilla's Movement Building From Home community call. Bo talked about how to

design community calls that involve accessible, purposeful, and impactful **collaboration**, drawing on ACL's activities from Nov 2021 to June 2022. Slides are available here:
    - https://app.pitch.com/app/presentation/f8ce2616-fb74-4ba9-b0ee-f555313cab4c/b64cdd0f-9e0e-442c-b181-037f0c7fc7fd

- **June 2022:**
    - We delivered a workshop at **RightsCon 2022 (June 10)**
        - https://rightscon.summit.tc/t/2022/events/taking-algorithms-to-court-empowering-communities-to-enact-legal-accountability-7goN7yPKzzv3EXYkjJFRnz
    - Our session, co-facilitated Ranjit Singh, Gina Helfrich, Jillian Powers, and Borhane Blili-Hamelin, invites participants to explore the place of the courts in accountability for algorithmic harms.
    - Workshop agenda and collaborative document: 2022_06_10 Rightscon Workshop (Accountability Case Labs)
    - Presentation slides: https://app.pitch.com/app/presentation/f8ce2616-fb74-4ba9-b0ee-f555313cab4c/6f14c519-242c-4760-948d-a3a6016d0c07
- **May 2022:**
    - Over the coming months, ACL will be collaborating with Mozilla's **'Transparency** for AI Builders Best Practices Framework', aimed at building a resource to help builders navigate questions of transparency for AI systems.
- **April 2022:**
    - Our proposal for RightsCon about the place of the courts in algorithmic accountability was accepted! See you at Rightscon in June!
    - April 7, noon ET. We're holding a collaborative meetup to explore the results of our Winter 2022 survey. Join us!
        - https://www.eventbrite.com/e/accountability-case-labs-winter-2022-survey-insights-tickets-288692064737
    - See our survey results here: ACL Feb 2022 survey results

- **March 2022:**
    - We have 2 workshops at Mozfest! Join us!
    - Insights from the Accountability Case Labs Project on March 11th
        - Themes: Who are the agents of algorithmic accountability? (Who should hold whom accountable?) What is algorithmic accountability? What are the hard problems around algorithmic accountability? How can you build workshops that address the needs of your own communities on these topics?
        - We explore these questions through fun collaborative activities that draw on our qualitative research!
        - Shout out to Mrin Bhattacharya and Jillian Powers for co-facilitating!
        - Session info: https://schedule.mozillafestival.org/session/LEMZRH-1
    - Accountability CaseLabs: ShotSpotter (March 8th)
        - The main theme is taking algorithmic evidence to court. (Fun secondary theme: pseudo-science, and the misuses of metrics.)
        - Our case study is a Frye motion asking a US court to scrutinize the reliability of ShotSpotter evidence
            - Scoping the 'action' to feature for this case was tricky, so we did it in house. Here are the slides for our presentation!
            - 2022_03_08 Mozfest ShotSpotter Workshop...
        - Shout out to Tina Lassiter and Gina Helfrich for co-facilitating!
        - Session info: https://schedule.mozillafestival.org/session/TDXFMB-1
- **February 2022**
    - We released our winter 2022 survey.
    - We piloted activities that became the basis for our Mozfest insights workshop during our February community call.

        - 2022_02_03 February Meetup (Accountability Case ...
- **January 2022**
    - **We delivered our first MVP: a case study based workshop about Twitter's Bias Bounty, featuring Kyra Yee and Irene Font Peradejordi**
        - 2022_01_26 Bias Bounty Workshop (Accountability ...
        - Blog post from our January presenters, Irene Font Peradejordi & Kyra Yee, about the story of Twitter's bias bounty program
    - January Blog post about the team work behind our January Milestone of a workshop on the accountability ramifications of bounties for algorithmic harms, and our next steps!
    - We finished our first round of exploratory research interviews.
- ADD NOV & DECEMBER

Resources

- 2022 survey results
    - ACL Feb 2022 survey results
- Workshop agendas
    - 2022_01_26 Bias Bounty Workshop (Accountability Case La...
    - 2022_03_08 Mozfest ShotSpotter Workshop (Accountabilit...
    - 2022_03_11 Mozfest Insights Workshop (Accountability Ca...
    - 2022_04_07 Winter 2022 Survey Insights (ACL)
    - 2022_06_10 Rightscon Workshop (Accountability Case Labs)
- Monthly meetup notes (see what we're up to!)
    - 2022_02_03 February Meetup (Accountability Case Labs)
    - 2022_01_06 January Meetup (Accountability Case Labs)
    - 2021_12_02 December Meetup (Accountability Case Labs)
    - 2021_11_04 Meetup (Accountability Case Labs)
- Sign up for our Google Calendar

Document Contents

Note

The roadmap below describes our activities during our seeding phase — November 2021 to April 2022 — as part of the first cohort of the Mozfest CSAWG working group.

We are still going! Stay tuned for a more up to date description of our current activities!

Mission

Our mission is to connect and build common strategic insights across the full range of technical and social experts, researchers, builders, and advocates who care about AI accountability

Why Build This?

Decisions made around AI systems need to face adequate scrutiny. However creating an ecosystem of accountability for AI systems is a complex and disorienting puzzle. Our project looks at one of the roots of this complexity, the dizzying range of expertise involved:

- Disciplines as diverse as policy-making, public administration, law, ethics and social science to design, data science, engineering
- Subject matter areas as different as data stewardship, biomedical R&D, financial risk modeling, content moderation, labor, human rights law,

criminal justice, surveillance, third party auditing, cybersecurity, ML ops, etc.

Our solution involves case studies about specific real world AI accountability problems and solutions, but with a twist: we will build the case studies as collaborative events, in which we invite a cross-disciplinary group of participants to co-create strategic insights. Our hope is to build bridges across those expert groups through the experience of targeted, purposeful collaboration.

1. Open Canvas[1]

Problems

Top 1-3 problems we want to solve

1. **Strategy Gaps:** The complexity of the problem of bringing about a world with adequate AI accountability makes it difficult to helpfully connect the specific decisions and opportunities available to different actors in this space to broader strategic goals.
2. **Expert Silos:** The sheer complexity and massively cross-disciplinary nature of the problem of building accountability at every layer of the governance of AI systems leads to **silos** among the expert and stakeholder groups involved: **barriers** to sharing knowledge, and to coordination across groups.
3. **Translation:** Common language gaps across expert groups create barriers to specificity and depth in cross-disciplinary AI accountability conversations.

[1] Follows the Mozilla Open Canvas model described here: https://mozilla.github.io/open-leadership-training-series/articles/opening-your-project/develop-an-open-project-strategy-with-open-canvas/

Solution

Outline our proposed solution for each problem

1. **Case studies:** Case studies are an effective tool for filling strategy gaps — for connecting broad, complex strategic problems to the specifics of the real world situations within which opportunities for action arise. They are also extremely flexible, allowing us to tailor them to the specific questions and needs we have in building them.

2. **Insights about expert and actor groups:** Combine qualitative research w. multi-stakeholder consultation approach to insights that identify the specific *needs* + *pain points* + *opportunities* of different expert and actor groups.

3. **Experiential learning approach to bridging silos:** **Collaborative** cross-disciplinary and multi-stakeholder **events,** with adequate scaffolding to navigate knowledge and common language gaps, allow participants to take a *learning by doing* approach to dissolving the boundaries that keep expert and stakeholder groups apart.

Key Metrics

How will we measure success?

- Successfully prototype and run multiple Case Labs event
- Successful cross-functional collaboration across multiple expert groups among project collaborators, and Case Labs event attendees.
- Case Labs process gets successfully forked by other communities
- Case Labs (case study) insights get used by educators, actors, and other experts

Resources Required

What do we need to build an MVP?

- Time + skills to build a prototype case lab

- (Potentially, especially if we want to build new case studies, rather than repurpose existing ones) Partnership with people interested in having their accountability problem + solution featured and discussed in a case lab
- Participants with the interest and passion to make 25-35 people collaborative events thrive, and representing a broad range of AI accountability expert groups.

Skills Required

Example skills we hope contributors will collectively bring to the table (no one is expected to have them all)

- **Subject matter expertise:** in *any field or community involved with AI accountability*, governance, and transparency. Our project is about the needs, pain points, and opportunities of the different communities involved in AI accountability, and about building bridges among them. We vitally need a diverse range of folks involved with AI accountability at the table throughout.
- **Qualitative research:** our project needs to generate **insights** to thrive — including into the specific needs + pain points + opportunities of the different expert and stakeholder groups interested in AI accountability. The skills of qualitative researchers are especially suited to that task.
- **Design and process:** our project needs to come up with its own approach to **collaborative, generative events**. Skills in designing collaborative, generative processes — e.g. multi-stakeholder consultation, participatory design, pedagogy — will be key to our success.
- **Writing/Research/Communication:** our project will need to share its results with the world. We want help building insightful **documentation** and/or **artefacts** about specific *case studies*, about our *process*, and/or about *other insights* we gather along the way.
- **Connectors:** our project needs ties with diverse expert communities to thrive, we need help cultivating those ties

Contributor Profiles

Contributor types and ideal contributors

- **Case Labs MVP builders:** contributors who help build our MVP, and figure out how to build an approach to collaborative events that achieves our goals
- **Insights builders:** contributors who want to help shape what we learn through our events, and help make those insights useful
- **Community builders:** contributors who want to help reach + build ties with folks who can make our case labs events tick, and to foster collaboration around and beyond our project
- **Documentation/artefact builders:** contributors who want to help build resources that enable other communities to fork our process, and to learn from our insights

User Profiles

Target audience and early adopters

- People who want more expansive bridges across the expert groups who work on AI accountability
- People who want insights that fill strategy gaps about building an ecosystem of AI accountability
- People who want insights into the specific needs + pain points + opportunities of different AI accountability expert groups.

Unique Value Proposition

A clear message that states what you offer and why you are different from other solutions that are available

- Experiential learning approach to bridges across expert communities: bridging silos across communities requires changing habits and

perspectives; we believe that purposeful collaborative events can bring about such changes through “learning-by-doing”.

- New use of case studies
- New approach to insights into the problem of cross-disciplinarity in the AI accountability space

2. Scoping Map

Need-to-have

- Prototype and run multiple Case Labs event
- Share our process and insights in a Mozfest session
- Insights into the specific needs + pain points + opportunities of different groups working on AI accountability
- Insights into the strategic problem of bringing about an ecosystem of AI accountability
- Breadth of AI accountability cases: we need our cases, over time, to cover a diverse range of technical to social real world AI accountability problems and solutions.

Nice-to-have

- Documentation/artefacts sharing our process + how to fork it
- Documentation/artefacts sharing our insights
- Novel case studies of our own.
    - One way to lower the barrier to prototyping Case Labs events would be to repurpose existing case studies on AI accountability; where we would build an event around a well documented case study, but design the event to arrive at new insights. It major ‘nice to have’, though, would be for us to cover fresh case study territory with our events — either focusing on cases that are known but poorly documented, or by focusing on emergent, poorly known real world

AI accountability cases, say, in partnership with people currently attempting to solve a real world AI accountability problem.

Out of Scope

- Academic publications
- Novel accountability solution of our own (we build new, cross-disciplinary insights from examining existing real world AI accountability problems and solutions.)

3. Milestone Mapping

Q3 2021

- Project Proposal (July 2021)
- Project Roadmap and Milestones (Sept 2021)
- Begin inviting and onboarding collaborators (Late Sept 2021)

Q4 2021

- Monthly Project Meetups (Beginning Oct 2021)
- Mozfest Event Proposals (Oct 2021)
- Design Case Labs process w. Project collaborators (Q3)
- Begin prototyping Case Labs Events (ideally running 2-3 between December 2021 and March 2022)
-
-

Q1 2022

- Complete Series of First 2-3 Case Labs events (March 2022)
    - Completed: January (Twitter's bias bounty program) March (
- Begin exploratory research on the pain points of algorithmic accountability practitioners
    - Completed exploratory research interviews (Dec-Jan)

    - Completed survey aimed at getting baseline insights on how practitioners in different roles think about where we are with algorithmic accountability, and where we need to go.
- Begin knowledge sharing (help other communities fork our insights)
    - Completed (in March, at Mozfest 2022)
    - Will be supplemented by meetup about our survey results on April 7
-

MOZFEST

- 1x Case Labs Project + Outputs Workshop: What did we build? What did we learn?
    - Done, March 11
- 1x Case Labs Demo Event: run 1x Case Labs event at Mozfest
    - Done, March 8
-

BEYOND TAI WG

- Build a sustainable home for our open community
- Run monthly meetups featuring insights from accountability practitioners
- Run more case study based workshops!
- Conduct more in depth research about the strategy gaps and expert silos algorithmic accountability practitioners face
- Document and share insights about case studies
- Document and share insights from our research
- Document and share insights about workshop design
-
-
-

4. Resources, Definitions & Bibliography

2022 ACL DEFINITION

Accountability is distinct from but closely related to the following notions: governance, transparency, auditing, impact assessment, and responsibility. **Accountability** means scrutinizing, evaluating, or judging one another's actions in a way that has real consequences. In this general sense, it is part of most human relationships. In many institutional and organizational settings, impactful scrutiny of actions is performed through standardized evaluation methods: audits. Algorithmic **auditing** refers to institutionalized, standardized, and impactful methods for scrutinizing algorithms. Algorithmic accountability is also often thought to require processes such as **documentation** that lower the barrier to meaningful scrutiny of the decision that get made throughout an algorithm's lifecycle.

It's helpful to realize that scholarship on accountability and auditing has a long history. The above characterization is a paraphrase of the account of Michael Power's 1997 book:

- Power, M. (1997). *The Audit Society: Rituals of Verification*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198296034.001.0001

See also

- Raji, I. D. (2022, February 2). It's Time to Develop the Tools We Need to Hold Algorithms Accountable. *Mozilla Foundation*. https://foundation.mozilla.org/en/blog/its-time-to-develop-the-tools-we-need-to-hold-algorithms-accountable/
- Raji, I. D., Xu, P., Honigsberg, C., & Ho, D. (2022). Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance. *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 557–571. https://doi.org/10.1145/3514094.3534181
- Rosenblat, A., Kneese, T., & boyd, danah. (2014). Workshop Primer: Algorithmic Accountability. SSRN Electronic Journal. https://doi.org/10/gpddzx

- Wieringa, M. (2020). What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 1–18. https://doi.org/10/ggjpch
- Kohli, N., Barreto, R., & Kroll, J. A. (2018). *Translation Tutorial: A Shared Lexicon for Research and Practice in Human-Centered Software Systems*. 7. https://facctconference.org/static/tutorials/fatconf18_lexicon_tutorial.pdf

2021 ACL Definitions

Accountability can be difficult to define. It is closely related, but distinct from each of the following notions: governance, transparency, auditing, impact assessment, responsibility. In a loose sense, it is sometimes equated with good governance. More narrowly, a world with more accountability for AI would be a world in which decisions made by those who control, develop, deploy, procure, maintain AI systems face scrutiny, especially in light of the harms and wrongs that those systems risk incurring, and in which such scrutiny leads to consequences.

The following definition of the *narrow sense* of accountability from public administration scholarship is helpful to have in mind:

- "a relationship between an actor and a forum, in which the actor has an obligation to explain and to justify his or her conduct, the forum can pose questions and pass judgement, and the actor may face consequences." (Bovens (2007), 450)[2]

N. Kohli, R. Barreto & J.A. Kroll offer a succinct characterization in their shared lexicon for the inaugural FAccT/FAT* conference on fairness, accountability, and transparency in socio-technical systems:

- "Accountability is fundamentally about the answerability of actors for outcomes." (N. Kohli, R. Barreto & J.A. Kroll, 3)

[2] See also Wieringa (2020), and Moss et al. (2021).

In philosophy scholarship on the concept, accountability is intimately connected to **moral responsibility and blame**. (See e.g. Shoemaker (2011) and Smith (2012) for a debate about whether accountability should be understood as one of many sub-species of moral responsibility, or whether it can be adequately captured by a unified account of responsibility.)[3]

Other definition of accountability from the **AI model development community**:

- "We use accountability to mean the state of being responsible or answerable for a system, its behavior and its potential impacts [38]. Although algorithms themselves cannot be held accountable as they are not moral or legal agents [7], the organizations designing and deploying algorithms can through governance structures. Proposed standard ISO 37000 defines this structure as "the system by which the whole organization is directed, controlled and held accountable to achieve its core purpose over the long term." 1 If the responsible development of artificial intelligence is a core purpose of organizations creating AI, then a governance system by which the whole organization is held accountable should be established.
- In environmental studies, Lynch and Veland [45] introduced the concept of urgent governance, distinguishing between auditing for system reliability vs societal harm. For example, a power plant can be consistently productive while causing harm to the environment through pollution [42]. Similarly, an AI system can be found technically reliable and functional through a traditional engineering quality assurance pipeline without meeting declared ethical expectations. A separate governance structure is necessary for the evaluation of these systems for ethical compliance. This evaluation can be embedded in the established quality assurance workflow but serves a different purpose, evaluating and optimizing for a different goal centered on social benefits and values rather than typical

[3] Shoemaker, D. (2011). Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility. Ethics, 121(3), 602–632. https://doi.org/10/d9tch4
Smith, A. M. (2012). Attributability, Answerability, and Accountability: In Defense of a Unified Account. Ethics, 122(3), 575–589. https://doi.org/10/gfkzm3

> performance metrics such as accuracy or profit [39]. Although concerns about reliability are related, and although practices for testing production AI systems are established for industry practitioners [4], issues involving social impact, downstream effects in critical domains, and ethics and fairness concerns are not typically covered by concepts such as technical debt and reliability engineering." (Raji & Smart et al.)

Examples of Accountability Mechanisms

Institutional review boards (IRB) for research involving human subjects (including biomedical research) are arguably an example of accountability in the narrow sense: researchers are required to explain and justify their research projects in light of their board's ethical standards, the board gets to examine and assess the proposals, and determines whether or not they are allowed to move forward.

Internal auditing process proposed by the ML community, see especially Raji & Smart et al.

- "External audits are designed to identify these risks from outside the system and serve as accountability measures for these deployed models. However, such audits tend to be conducted after model deployment, when the system has already negatively impacted users [26, 51]. In this paper, we present internal algorithmic audits as a mechanism to check that the engineering processes involved in AI system creation and deployment meet declared ethical expectations and standards, such as organizational AI principles." (Raji & Smart et al.)"

External audit as found in other industries, where model auditing is done by independent, certified third party auditors. (See e.g. https://forhumanity.center/independent-audit-of-ai-systems/)

Impact assessment as found in other industries (environmental impact assessment, financial impact assessment, human rights impact assessment, etc.)

but also in the AI space (privacy impact assessment, human rights impact assessment, etc.) (See Moss et. al. (2021))

For a comprehensive overview of examples of algorithmic auditing practices in the public sector, see (Ada Lovelace Institute et al., 2021).

Bibliography

Selection of highly recommended readings for an overview of the topic

- Ada Lovelace Institute, AI Now Institute, & Open Government Partnership. (2021). *Algorithmic Accountability for the Public Sector*. https://www.opengovpartnership.org/documents/algorithmic-accountability-public-sector/
- Bovens, M. (2007). Analysing and Assessing Accountability: A Conceptual Framework. *European Law Journal*, *13*(4), 447–468. https://doi.org/10/b4hmbf
- Hutchinson, B., Smart, A., Hanna, A., Denton, E., Greer, C., Kjartansson, O., Barnes, P., & Mitchell, M. (2021). Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure. *ArXiv:2010.13561 [Cs]*. http://arxiv.org/abs/2010.13561
- Kohli, N., Barreto, R., & Kroll, J. A. (2018). *Translation Tutorial: A Shared Lexicon for Research and Practice in Human-Centered Software Systems*. 7.
- Moss, E., Watkins, E. A., Singh, R., Elish, M. C., & Metcalf, J. (2021). *Assembling Accountability: Algorithmic Impact Assessment for the Public Interest*. Data & Society. https://datasociety.net/library/assembling-accountability-algorithmic-impact-assessment-for-the-public-interest/
- Nissenbaum, H. (1996). Accountability in a computerized society. *Science and Engineering Ethics*, *2*(1), 25–42. https://doi.org/10/dn6p6h
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal

Algorithmic Auditing. *ArXiv:2001.00973 [Cs]*. http://arxiv.org/abs/2001.00973

- Rosenblat, A., Kneese, T., & boyd, danah. (2014). Workshop Primer: Algorithmic Accountability. SSRN Electronic Journal. https://doi.org/10/gpddzx
- Wieringa, M. (2020). What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 1–18. https://doi.org/10/ggjpch

Resources about workshop design

- Borhane Blili-Hamelin, Beth M. Duckles, and Marie-Eve Monette (Open Post Academics), *Toolkit for Cross-Disciplinary Workshops*. February 2022
    - https://doi.org/10.5281/zenodo.6026972
    - Blog post about the Toolkit: https://openpostac.org/blog/toolkit-for-cross-disciplinary-workshops/
- Timelines (World-Building Activity)
    - Wong, R. (2021). *Timelines: Design Activities for Surfacing Values and Ethics in Technologies - CLTC UC Berkeley Center for Long-Term Cybersecurity*. https://cltc.berkeley.edu/2021/07/01/timelines/
    - Wong, R. Y., & Nguyen, T. (2021b). *Timelines (Facilitation Guide)*. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, 1–15. https://doi.org/10/gj2fkm
    - Wong, R. Y., & Nguyen, T. (2021a). *Timelines: A World-Building Activity for Values Advocacy*. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, 1–15. https://doi.org/10.1145/3411764.3445447
- (CSCCE) Woodley, L., Pratt, C., Ainsworth, R., Amsen, E., Bakker, A., Butland, S., O'Donnell, S., Penfold, N., Pope, A., Quigley, T., & Tsang, E.

(2020). *A guide to using virtual events to facilitate community building: Event formats*. https://doi.org/10.5281/zenodo.3934385

5. Our Contributors

We strive to credit everyone who contributes to this project in any way, unless they request not to.

All contributors

Here is the current list of everyone who contributed to this project in any way.

- Vasundhra Dahiya
- Vanja Skoric
- Tina Lassiter
- Sundar Narayanan
- Sophia Katrenko
- Shlomi Hod
- Shea Brown
- Ranjit Singh
- Pamela Jasper
- Mrin Bhattacharya
- Luke Richards
- Kyle Smith
- Jonathan Poritz
- John Hurst
- Jillian Powers
- Jacqui Ayling
- Itzel Amieva
- Gina Helfrich
- Divij Joshi
- Debra Erickson
- Borhane Blili-Hamelin (Project Lead)

- Beth M. Duckles
- Bernease Herman
- Bernd Durrwachter

Project team contributors

Some of our work is organized into teams. People can contribute to as many teams as they want. Here are the contributors for each. The names are in reverse alphabetical order.

Cases Studies Team

Scout and select topics (case studies) for our events. Build partnerships with folks who know the case.

- Sophia Katrenko
- Shea Brown
- Ranjit Singh
- Kyle Smith
- Jillian Powers
- Jacqui Ayling
- Debra Erickson
- Borhane Blili-Hamelin

Design Team

Design events that helps a cross-disciplinary group of participants think through a specific real world step towards increasing AI accountability

- Vanja Skoric
- Tina Lassiter
- Shlomi Hod
- Ranjit Singh
- Mrin Bhattacharya
- Kyle Smith
- Jillian Powers
- Gina Helfrich

- Borhane Blili-Hamelin
- Beth Duckles

Qual Team

Exploratory research (interviews and surveys) to learn about the pain points of folks working to increase accountability around AI systems, including around cross-disciplinarity and expert silos.

- Tina Lassiter
- Mrin Bhattacharya
- Kyle Smith
- Jillian Powers
- Ian Lowrie
- Borhane Blili-Hamelin
- Beth Duckles

6. Credits

Photo (cropped) by USGS on Unsplash

This document was written by Borhane Blili-Hamelin, with helpful feedback from Beth M. Duckles, Bernease Herman, John Hurst, and the wonderful team running the 2022 Mozfest Trustworthy AI working group, Temi Popo and Chad Sansing!