



PARTNERSHIP ON AI

WHITE
PAPER

Making AI Inclusive

Four Guiding Principles for Ethical Engagement

Dr. Tina M. Park

JULY 20, 2022



Contents

Introduction	3
Guiding Principles for Ethical Participatory Engagement	5
1. All Participation Is a Form of Labor That Should Be Recognized	5
2. Stakeholder Engagement Must Address Inherent Power Asymmetries	6
3. Inclusion and Participation Can Be Integrated Across All Stages of the Development Lifecycle	7
4. Inclusion and Participation Must Be Integrated to the Application of Other Responsible AI Principles	7
Recommendations for Ethical Engagement in Practice	8
1. Allocate Time and Resources to Promote Inclusive Development	8
2. Adopt Inclusive Development Strategies Before Development Begins	8
3. Train Towards an Integrated Understanding of Ethics	8
Conclusion	12
Acknowledgments	13
Notes	14

Introduction

While the concept of “human-centered design” is hardly new to the technology sector, recent years have seen growing efforts to build inclusive artificial intelligence (AI) and machine learning (ML) products. Broadly, **inclusive AI/ML refers to algorithmic systems which are created with the active engagement of and input from people who are not on AI/ML development teams.** This includes both end users of the systems and non-users who are impacted by the systems.* To collect this input, practitioners are increasingly turning to engagement practices like user experience (UX) research and participatory design.

Amid rising awareness of structural inequalities in our society, embracing inclusive research and design principles helps signal a commitment to equitable practices. As many proponents have pointed out, it also makes for good business: Understanding the needs of a more diverse set of people expands the market for a given product or service. Once engaged, these people can then further improve an AI/ML product, identifying issues like bias in algorithmic systems.

Despite these benefits, however, **there remain significant challenges to greater adoption of inclusive development in the AI/ML field. There are also important opportunities.** For AI practitioners, AI ethics researchers, and others interested in learning more about responsible AI, this Partnership on AI (PAI) white paper provides guidance to help better understand and overcome the challenges related to engaging stakeholders in AI/ML development.

Ambiguities around the meaning and goals of “inclusion” present one of the central challenges to AI/ML inclusion efforts. To make the changes needed for a more inclusive AI that centers equity, the field must first find agreement on foundational premises regarding inclusion. Recognizing this, **this white paper provides four guiding principles for ethical engagement grounded in best practices:**

1. All participation is a form of labor that should be recognized
2. Stakeholder engagement must address inherent power asymmetries
3. Inclusion and participation can be integrated across all stages of the development lifecycle
4. Inclusion and participation must be integrated to the application of other responsible AI principles

To realize ethical participatory engagement in practice, **this white paper also offers three recommendations aligned with these principles for building inclusive AI:**

1. Allocate time and resources to promote inclusive development
2. Adopt inclusive strategies before development begins
3. Train towards an integrated understanding of ethics

* Impacted non-users are people who are impacted by the deployment of an AI/ML system, but are not the direct user or customer of that system. For example, in the case of students in the United Kingdom in 2020 whose A-level grades were determined by an algorithm, the “user” of the algorithmic system is Ofqual, the official exam regulator in England, and the students are “impacted non-users.”

This white paper’s insights are derived from the forthcoming research study “Towards An Inclusive AI: Challenges and Opportunities for Public Engagement in AI Development.” That study drew upon discussions with industry experts, a multidisciplinary review of existing research on stakeholder and public engagement, and nearly 70 interviews with AI practitioners and researchers, as well as data scientists, UX researchers, and technologists working on AI and ML projects, over a third of whom were based in areas outside of the US, EU, UK, or Canada. Supplemental interviews with social equity and Diversity, Equity, and Inclusion (DEI) advocates contributed to the development of recommendations for individual practitioners, business team leaders, and the field of AI and ML more broadly.

This white paper does not provide a step-by-step guide for implementing specific participatory practices. It is intended to renew discussions on how to integrate a wider range of insights and experiences into AI/ML technologies, including those of both users and the people impacted (either directly or indirectly) by these technologies. Such conversations – between individuals, inside teams, and within organizations – must be had to spur the changes needed to develop truly inclusive AI.

Guiding Principles for Ethical Participatory Engagement

There is strong consensus^{1,2,3} that the inclusion of a diverse body of end-users and other stakeholders in the creation of new technology is good for both improving the quality and usability of a product or service and mitigating possible emergent harms. However, even when AI/ML practitioners make efforts to increase inclusion or participation, unclear definitions of these concepts can create a mismatch between the desires of inclusion advocates and the outcomes of these efforts.

For many, the purpose of engaging users and other non-technical audiences isn't to use AI/ML technology to emancipate and provide restitution to oppressed communities. In most cases, the intent of incorporating participatory practices is more modest: to expand the circle of people who can use and benefit from a product or service and to avoid some of the more obvious harms related to algorithmic bias. Nevertheless, if we are to genuinely make space for those without explicit expertise in machine learning and AI to contribute to the overall success of AI, we must reconsider our foundational assumptions.

To make changes needed for a more inclusive AI, the field must first agree on some foundational premises regarding inclusion. Below are four guiding principles that practitioners should adopt as their operating assumptions to align their approach to engagement with ethical best practices. These principles build upon the work of many thought leaders in the fields of Indigenous AI,⁴ feminist HCI,⁵ crip technoscience,* data justice,⁷ and critical race theory^{8,9} who have far more substantial publications discussing the importance of these dimensions.

- 1 All Participation Is a Form of Labor That Should Be Recognized**
- 2 Stakeholder Engagement Must Address Inherent Power Asymmetries**
- 3 Inclusion and Participation Can Be Integrated Across All Stages of the Development Lifecycle**
- 4 Inclusion and Participation Must Be Integrated to the Application of Other Responsible AI Principles**

* Crip technoscience is the field of practice and scholarship outlined by Hamraie and Fritsch⁶ which centers disabled people as "experts and designers of everyday life" and harnesses technoscience for political action. "Crip" serves as a reclamation of a derogatory term used against people with disabilities to describe the "non-compliant, anti-assimilationist position that disability is a desirable part of the world." "Technoscience" refers to the "co-production of science, technology, and political life."

All Participation Is a Form of Labor That Should Be Recognized

To ensure that the material benefits of AI/ML systems are experienced by all, we must first recognize that the value of these technologies is dependent on users and public participation.

Given the need for large amounts of data to build algorithmic systems, AI/ML technology already relies heavily on the participation of the public. As such, it is necessary to define “participation” as any direct or indirect contribution to the creation, development, deployment, and sustainment of an AI/ML system.

By recognizing all participation as work, active consent (where users are asked for consent, given the ability to opt out, and offered compensation if they choose to participate)¹⁰ becomes necessary. This also empowers participants to withdraw from projects or AI/ML systems they might find harmful or otherwise unappealing.¹¹ Differences between the passivity, purpose, and expertise required for this participation can help determine what is suitable recognition and compensation.

Stakeholder Engagement Must Address Inherent Power Asymmetries

Many members of the public, especially those who are members of marginalized and historically exploited communities, are wary of contributing to participatory efforts led by companies or other entities.^{12,13,14} Historic and contemporary experiences of giving valuable data to others who are able to profit from them creates an environment of mistrust. Historically oppressed communities, such as the Black community in the US, are often very familiar with the use of their bodies and labor for the financial profit of others, especially the dominant class of White Americans.

The relationship between developer and user is often presented as neutral, but several factors often position users and the public in a subordinate position than AI/ML developers and researchers.^{15,16} Users and members of the public rarely have the power to access proprietary information or drive major decisions at AI/ML organizations. Access to information and decision-making authority is mediated by developers (who may not have the final say in granting that access themselves), with limited recourse to obtain it.

Furthermore, structural power asymmetries exist in these relationships due to histories of colonization, discrimination, and other forms of social, economic, and political exclusion. These structural inequalities, which privilege and empower some over others, contribute to a sense of apprehension to engage in participatory processes. This gives community members reasons to doubt their opinions will be taken seriously and their participation will result in meaningful impact. Even the best practitioners will have to grapple with the repercussions of bad faith, actors engaged in extractive practices, and the structural dynamics of inequality.

Thus, even if interpersonal relations are established on more equitable footing, societal dynamics such as anti-Black racism or misogyny may (and likely will) affect the ongoing relationship. Since structural inequality is intersectional, it is also not enough to find parity across one specific dimension of difference, such as race or gender.

Recognizing these dynamics — and putting policies and practices into place to mitigate differences — is integral to establishing respectful and mutually beneficial relationships between developers and community members.¹⁷ It paves the way for the shift from AI/ML practitioners “building for” to “building with” users and the public.¹⁸ It is also necessary as a means to avoid deepening existing harm and inequality through participatory engagement.

PRINCIPLE **3** **Inclusion and Participation Can Be Integrated Across All Stages of the Development Lifecycle**

While significant strides are being made,¹⁹ it remains all too common to see participatory practices implemented at the end of the AI/ML development lifecycle instead of being fully integrated throughout the process.²⁰

Often this is done by engaging UX researchers, who are far more likely to be trained to identify possible end-users and think about who might be directly and indirectly impacted by the deployment of the AI/ML system. UX researchers, however, may not have detailed knowledge of the specifics of the algorithmic model or be given the ability to improve it.

As many equity and inclusion advocates have pointed out, every stage of the development process has the potential to be shaped and directed by users and impacted community members. The deepest and longest-term inclusive participatory practices will establish relationships with community stakeholders and give them the space to direct the purpose and intention of an AI/ML project. While most practitioners, especially those working on commercial products and services, are less likely to engage in co-development practices such as these, practitioners can be mindful of when input is not being gathered and address the issues that might arise due to that absence.



It remains all too common to see participatory practices implemented at the end of the AI/ML development lifecycle instead of being fully integrated throughout the process.

PRINCIPLE **4**

Inclusion and Participation Must Be Integrated to the Application of Other Responsible AI Principles

Responsible AI frameworks often discuss foundational principles individually, rather than as integrated values which intersect and build upon each other. The capacity to access and benefit from the development and use of the technology should be considered in conjunction with other responsible AI principles.²¹

In addition to inclusion, these principles commonly include transparency, accountability, security/privacy, reliability, and fairness. Transparency of algorithmic models, for instance, cannot support responsible development if it is practiced through documentation that is incomprehensible for non-technical audiences or there are no mechanisms for holding developers accountable for harms done to different communities.

Recommendations for Ethical Engagement in Practice

To center inclusion as a guiding principle in the development of AI/ML technology, practitioners must overcome several challenges. These challenges include the need for organizational support throughout the development lifecycle, grappling with histories of stakeholder exclusion, and an incomplete understanding by many of inclusion's importance.

Too often, practitioners we interviewed who were working in organizations where user or public input was not as highly prioritized reported having to be creative with building in the opportunity for feedback outside of the project team. Some expressed frustration with the additional time spent trying to convince project managers and other decision-makers within their organizations about the value of public user or stakeholder input and AI ethics more broadly.

In line with the principles offered above, the three recommendations below speak to the challenges facing individual practitioners. To address deeply rooted and broadly distributed challenges, these recommendations require buy-in, not only from the practitioners themselves, but the organizations and teams they work within.

- 1** Allocate Time and Resources to Promote Inclusive Development
- 2** Adopt Inclusive Development Strategies Before Development Begins
- 3** Train Towards an Integrated Understanding of Ethics

Allocate Time and Resources to Promote Inclusive Development

The majority of interviewees in PAI's study of practitioners noted that their interest in inclusive participatory practices emerged from a personal commitment to social equity. However, these personal ambitions were only impactful when they were girded by material institutional support – something often lacking in for-profit companies.

Managing relationships with user and community-based contributors, strategizing different ways to engage diverse audiences, and synthesizing feedback across different stages of the development lifecycle requires a different set of skills, as well as time and resources to properly execute, than what is necessary for developing algorithmic models.

While technical team members should be part of inclusive participatory practices, teams and organizations should not displace the responsibility of inclusive participatory practices to individuals without broader organizational support. Beyond a stated commitment to responsible and inclusive development, committing organizational resources demonstrates a different level of dedication to pursuing inclusively and responsibly developed technology.

ACTIONS TO TAKE

- Build teams with explicit roles to support community-based relationships and focus on inclusive development, as well as other responsible AI practices.
- Draw from expertises outside of computer science or machine learning, such as anthropology, community organizing, disability studies, ethnic studies, gender studies, humanities, and sociology.
- Plan for sprint cycles that permit time for the collection of insights from users and/or impacted communities, as well as the synthesis of those findings for incorporation.

LIMITATIONS

If teams are unable to act on the input of community stakeholders, additional resources to conduct public engagement and more diverse, multidisciplinary staff members may not substantially shift how the work is completed or grow an organization's capacity to mitigate future harms. In the absence of having community members directly involved in decision-making and direction-setting, those on staff who serve as liaisons with the public should be empowered and have the authority to act in the interests of the community. Without this, there is a high likelihood any public engagement activities will be read as "participant washing,"* as public input will be perceived to have little to no impact on the final product or service.

* **Participant washing** refers to the way minimal public or user engagement is spun and exaggerated to present a company or organization as being more inclusive and civic-minded than they actually are.

Adopt Inclusive Development Strategies Before Development Begins

Much like AI ethics should not be treated as an afterthought,²² inclusive participation strategies should not be created after much of the AI/ML development lifecycle has passed. Having early discussions about inclusive development and participation goals for the project at the very beginning not only situates values of diversity, equity, and inclusion at the core of the project, but provides an opportunity to initiate important conversations that more broadly relate to the responsible development of AI.²³

While an organization may consider drafting more general guidelines to help each team identify their own approach, given the high degree of variability between projects, it may be more effective to develop project-specific inclusion strategies alongside project work plans. It is important to align the participatory objective with the appropriate mechanisms,²⁴ given the social context in which it is being developed.²⁵

ACTIONS TO TAKE

- **Identify marginalized stakeholders.** Who are the people who may use or be impacted by the use of the AI/ML product or service, but are not typically consulted?
- **Understand dynamics of power.** What are the power dynamics between the organization (developers) and members of the public (users / impacted communities), both specifically (interpersonal) and structurally (societal)?
- **Identify resources needed.** What is needed to build and sustain relationships with key (marginalized) stakeholders throughout / at different points of the AI/ML development lifecycle?
- **Identify integration points.** At what stage(s) of development should key stakeholders be engaged? Can stakeholders change these integration points?
- **Recognize contributions of participants.** What is the compensation policy for stakeholders who participate in the development of the AI/ML product or service? How are passive participants (e.g., people who contribute important data points for training dataset) compensated? How will they be credited? Are there opportunities to redistribute future success with participants (e.g., profit sharing)? Can participants withdraw their contributions or support (including any data taken)?
- **Build accountability mechanisms.** What processes or mechanisms exist for participants or future users/members of the public to hold the organization or company accountable for any harm experienced due to use of the algorithmic model?

LIMITATIONS

Drafting the best set of guidelines and policies for an organization or project will likely be an iterative process requiring resources that smaller organizations or leaner teams may not have. Currently, there are no “off-the-shelf” guidelines or best practices organizations and practitioners can draw upon to support their efforts and ensure that some guidance and direction is provided (rather than none at all). Additionally, having a plan does not mean that all team members will understand or implement it. Having incentives to encourage adoption is key.

Train Towards an Integrated Understanding of Ethics

The relevance and value of incorporating inclusive practices into AI/ML development may not be readily apparent to some practitioners. This is an opportunity to not only discuss the important role of users and impacted communities, but more generally about the need for the responsible and ethical development of AI.

When PAI put out an open call for interviewees who self-identified as incorporating participatory practices into AI/ML systems, individuals in a wide range of roles (from engineers to UX researchers to change management consultants) responded. Even among practitioners working at the same organization, there can be substantial differences in both their knowledge about how AI/ML systems were created or will be used and their ability to incorporate inclusive practices.

Creating a body of practitioners who are conversant on both equity and responsible AI issues will significantly help shift principles into practice by enabling more robust and thoughtful colleagues who share common definitions and understandings.

ACTIONS TO TAKE

- Develop and implement trainings and regular workshops on responsible AI principles and best practices, including inclusive practices for all staff members, that cover:
 - How “inclusion” works with and alongside other principles of responsible AI
 - The aims and implications of various participatory frameworks and approaches to understand that participation is not a “one-size fits-all” or singular concept

LIMITATIONS

As with any non-mandatory employee training, or supplemental professional development, those who are disinclined to engage in inclusive or responsible AI practices cannot be forced to learn and engage. Having leadership throughout the organization, including senior leadership, prioritize and highlight the value of these trainings (and the practices themselves) as being core to the success of the organization’s work can go a long way to improve adoption and commitment among employees. Also, many AI/ML practitioners operate outside of formal organizational spaces. Without structured professional development in place through a workplace, it is important to provide both free and paid learning opportunities for independent or start-up practitioners.

Conclusion

The growing diversity of actors and circumstances involved in AI/ML deployment makes establishing a set of ethical participatory practices especially difficult for this field. Even among practitioners working at the same organization, there can be substantial differences in both their knowledge about how AI/ML systems were created or will be used and their ability to incorporate inclusive practices. Additionally, the greater availability of AI development platforms, including “no-code” platforms, means that algorithmic models can be deployed without having deep expertise, expanding the number of circumstances automated systems are deployed in. This means many more instances where automated systems are being deployed without consideration of how the algorithms were developed, the provenance of the datasets and nature of the bias on which they were trained and tested, and the ethical implications of their development and deployment.

When AI/ML projects have small development teams, quickly approaching deadlines, and limited budgets, it can be easy to deprioritize the inclusion of users and other stakeholders. End-users, impacted non-users, and the public, however, are integral to the ethical development of AI/ML systems. Engaging them can identify issues which can be better resolved through algorithmic systems, lead to products and services that are useful and accessible by many, and formulate policies for creating and deploying AI.

On their own, technical understandings of (and solutions to) ethical issues related to AI/ML-enabled systems are insufficient. To both meaningfully and ethically implement stakeholder engagement practices, it is necessary to draw together understandings of structures of power and social inequality and apply them to the development and deployment of digital technology. Ignoring asymmetries of power can result in greater harm between those who develop AI/ML technology and those who are impacted by it. Given the long history of marginalized communities being asked to freely give their time and labor, extractive stakeholder engagements have the potential to deepen the social inequality many ethically oriented practitioners are trying to mitigate.

While AI/ML practitioners should not be expected to be both proficient developers and experts in social inequality, they do need to have shared language and concepts with those who are. To develop truly inclusive AI-ML technology, practitioners need additional resources, including training to build expertise, funding to support community engagement, and time to incorporate stakeholders feedback. In addition to an alternative framework to understand how to develop technology more responsibly and inclusively, support structures are needed to advance these efforts. Individual advocates cannot be expected to instigate the change necessary in the field: organizations, and the field more broadly, must integrate deep changes to how work is conducted if we are to address the social inequality within our AI/ML products and systems.

Acknowledgments

Many thanks are owed to the past and current team members of PAI, but especially to Drs. Jeffrey Brown and Christine Custis for their support and helpful conversations throughout the “Methods for Inclusion” project and to Jingying Yang for getting the project off the ground.

This report continues the work and ongoing advocacy of ethical AI researchers committed to radically transforming both the field and the technology itself, including the researchers at Ada Lovelace Institute, PAI Partners Data & Society and the Alan Turing Institute, and other organizations.

Finally, we thank the participants in the PAI study for generously offering their time and insights. These participants were enthusiastic about AI ethics and sought more information on how to more effectively engage different audiences to ensure AI/ML-driven technology was fair, equitable, and free of bias. We owe them a great debt of gratitude.

Notes

- 1 Jean-Baptiste, A. (2020). *Building for Everyone: Expand Your Market with Design Practices from Google's Product Inclusion Team*. John Wiley & Sons, Inc.
- 2 Romao, M. (2019, June 27). "A vision for AI: Innovative, Trusted and Inclusive." Policy@Intel. <https://community.intel.com/t5/Blogs/Intel/Policy-Intel/A-vision-for-AI-Innovative-Trusted-and-Inclusive/post/1333103>
- 3 Zhou, A., Madras, D., Raji, D., Milli, S., Kulynych, B. & Zemle, R. (2020, July 17). "Participatory Approaches to Machine Learning." [Workshop]. International Conference on Machine Learning 2020.
- 4 Lewis, J. E., Abdilla, A., Arista, N., Baker, K., Benesiinaabandan, S., Brown, M., ... & Whaanga, H. (2020). *Indigenous protocol and artificial intelligence position paper*. Indigenous AI. <https://www.indigenous-ai.net/position-paper>
- 5 Costanza-Chock, S. (2020). *Design justice: Community-led practices to build the worlds we need*. The MIT Press.
- 6 Hamraie, A., & Fritsch, K. (2019). "Crip technoscience manifesto." *Catalyst: Feminism, Theory, Technoscience*, 5(1), 1-33. <https://catalystjournal.org/index.php/catalyst/article/view/29607>
- 7 Taylor, L. (2017). "What is data justice? The case for connecting digital rights and freedoms globally." *Big Data & Society*, 4(2), 2053951717736335. <https://doi.org/10.1177/2053951717736335>
- 8 Benjamin, Ruha. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity. <https://www.ruhabenjamin.com/race-after-technology>
- 9 Hanna, A., Denton, E., Smart, A., & Smith-Loud, J. (2020). "Towards a critical race methodology in algorithmic fairness." In *Proceedings of the 2020 conference on fairness, accountability, and transparency* (pp. 501-512). <https://arxiv.org/abs/1912.03593>
- 10 Sloane, M., Moss, E., Awomolo, O. & Forlano, L. (2020). "Participation is not a design fix for machine learning." arXiv. <https://arxiv.org/abs/2007.02423>
- 11 Cifor, M., Garcia, P., Cowan, T.L., Rault, J., Sutherland, T., Chan, A., Rode, J., Hoffmann, A.L., Salehi, N. & Nakamura, L. (2019). "Feminist Data Manifest-No." Feminist Data Manifest-No. Retrieved October 1, 2020 from <https://www.manifestno.com/home>
- 12 Harrington, C., Erete, S. & Piper, A.M. (2019). "Deconstructing Community-Based Collaborative Design: Towards More Equitable Participatory Design Engagements." *Proceedings of the ACM on Human-Computer Interaction* 3(CSCW):1-25. <https://doi.org/10.1145/3359318>
- 13 Freimuth V.S., Quinn, S.C., Thomas, S.B., Cole, G., Zook, E & Duncan, T. (2001). "African Americans' Views on Research and the Tuskegee Syphilis Study." *Social Science & Medicine* 52(5):797-808. [https://doi.org/10.1016/S0277-9536\(00\)00178-7](https://doi.org/10.1016/S0277-9536(00)00178-7)
- 14 George, S., Duran, N. & Norris, K. (2014). "A Systematic Review of Barriers and Facilitators to Minority Research Participation Among African Americans, Latinos, Asian Americans, and Pacific Islanders." *American Journal of Public Health* 104(2):e16-31. <https://doi.org/10.2105/AJPH.2013.301706>
- 15 Barabas, C., Doyle, C., Rubinovitz, J.B., & Dinakar, K. (2020). "Studying Up: Reorienting the Study of Algorithmic Fairness around Issues of Power." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 167-176).
- 16 Harrington, C., Erete, S. & Piper, A.M.. (2019). "Deconstructing Community-Based Collaborative Design: Towards More Equitable Participatory Design Engagements." *Proceedings of the ACM on Human-Computer Interaction* 3(CSCW):1-25. <https://doi.org/10.1145/3359318>.
- 17 Chan, A., Okolo, C. T., Terner, Z., & Wang, A. (2021). "The Limits of Global Inclusion in AI Development." arXiv. <https://arxiv.org/abs/2102.01265>
- 18 Sanders, E. B. N. (2002). "From user-centered to participatory design approaches." In *Design and the social sciences* (pp. 18-25). CRC Press. <https://www.taylorfrancis.com/chapters/edit/10.1201/9780203301302-8/user-centered-participatory-design-approaches-elizabeth-sanders>
- 19 Leslie, D., Katell, M., Aitken, M., Singh, J., Briggs, M., Powell, R., ... & Burr, C. (2022). "Data Justice in Practice: A Guide for Developers." arXiv. <https://arxiv.org/ftp/arxiv/papers/2205/2205.01037.pdf>
- 20 Zdanowska, S., & Taylor, A. S. (2022). "A study of UX practitioners roles in designing real-world, enterprise ML systems." In *CHI Conference on Human Factors in Computing Systems* (pp. 1-15). <https://dl.acm.org/doi/abs/10.1145/3491102.3517607>
- 21 Leslie, D., Katell, M., Aitken, M., Singh, J., Briggs, M., Powell, R., ... & Burr, C. (2022). "Data Justice in Practice: A Guide for Developers." arXiv. <https://arxiv.org/ftp/arxiv/papers/2205/2205.01037.pdf>
- 22 Saulnier, L., Karamcheti, S., Laurençon, H., Tronchon, L., Wang, T., Sanh, V., Singh, A., Pistilli, G., Luccioni, S., Jernite, Y., Mitchell, M. & Kiela, D. (2022). "Putting Ethical Principles at the Core of the Research Lifecycle." Hugging Face Blog. Retrieved from <https://huggingface.co/blog/ethical-charter-multimodal>
- 23 Ada Lovelace Institute. (2021). "Participatory data stewardship: A framework for involving people in the use of data." Ada Lovelace Institute. <https://www.adalovelaceinstitute.org/report/participatory-data-stewardship/>
- 24 Delgado, F., Yang, S., Madaio, M., & Yang, Q. (2021). "Stakeholder Participation in AI: Beyond 'Add Diverse Stakeholders and Stir.'" arXiv. <https://arxiv.org/pdf/2111.01122.pdf>
- 25 Sloane, M., Moss, E., Awomolo, O. & Forlano, L. (2020). "Participation is not a design fix for machine learning." arXiv. <https://arxiv.org/abs/2007.02423>