

Participation in the Age of Foundation Models

Harini Suresh*
Cornell Tech
USA

Emily Tseng*
Cornell University
USA

Meg Young*
Data & Society Research Institute
USA

Mary L. Gray
Microsoft Research
USA

Emma Pierson
Cornell Tech
USA

Karen Levy
Cornell University
USA

ABSTRACT

Growing interest and investment in the capabilities of *foundation models* has positioned such systems to impact a wide array of services, from banking to healthcare. Alongside these opportunities is the risk that these systems reify existing power imbalances and cause disproportionate harm to historically marginalized groups. The larger scale and domain-agnostic manner in which these models operate further heightens the stakes: any errors or harms are liable to reoccur across use cases. In AI & ML more broadly, *participatory* approaches hold promise to lend agency and decision-making power to marginalized stakeholders, leading to systems that better benefit justice through equitable and distributed governance. But existing approaches in participatory AI/ML are typically grounded in a specific application and set of relevant stakeholders, and it is not straightforward how to apply these lessons to the context of foundation models. Our paper aims to fill this gap.

First, we examine existing attempts at incorporating participation into foundation models. We highlight the tension between participation and scale, demonstrating that it is intractable for impacted communities to meaningfully shape a foundation model that is intended to be universally applicable. In response, we develop a blueprint for participatory foundation models that identifies more local, application-oriented opportunities for meaningful participation. In addition to the “foundation” layer, our framework proposes the “subfloor” layer, in which stakeholders develop shared technical infrastructure, norms and governance for a grounded domain such as clinical care, journalism, or finance, and the “surface” (or application) layer, in which affected communities shape the use of a foundation model for a specific downstream task. The intermediate “subfloor” layer scopes the range of potential harms to consider, and affords communities more concrete avenues for deliberation and intervention. At the same time, it avoids duplicative effort by scaling input across relevant use cases. Through three case studies in clinical care, financial services, and journalism, we illustrate how this multi-layer model can create more meaningful opportunities for participation than solely intervening at the foundation layer.

*These authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution International 4.0 License.

FACCT '24, June 03–06, 2024, Rio de Janeiro, Brazil
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0450-5/24/06
<https://doi.org/10.1145/3630106.3658992>

CCS CONCEPTS

• **Computing methodologies** → **Machine learning; Artificial intelligence; Natural language generation; Computer vision;**
• **Human-centered computing** → **Human computer interaction (HCI); HCI theory, concepts and models;** • **Social and professional topics** → **Computing / technology policy.**

KEYWORDS

Foundation models, public participation, governance, communities, stakeholders

ACM Reference Format:

Harini Suresh, Emily Tseng, Meg Young, Mary L. Gray, Emma Pierson, and Karen Levy. 2024. Participation in the Age of Foundation Models. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FACCT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3630106.3658992>

1 INTRODUCTION

Recent years have seen a notable rise in interest and investment into *foundation models* [14], exemplified by systems such as GPT-4 or CLIP. Foundation models are unique in their generalizability, with the ability to adapt to a range of tasks not explicitly introduced during training. While many of the methods underpinning foundation models are not new (e.g., pre-training via self-supervised learning on unlabeled data), they are now being developed and deployed at an unprecedented scope and scale. These systems have spurred broad interest across industries including medicine [21, 62], software engineering [104], and education [38]. Alongside these opportunities, new and heightened risks have emerged, including environmental [66, 99] and economic impacts [18], extractive data labor [89], legal concerns [63], data inscrutability [9], and the homogenization of discriminatory behavior across applications [14].

Of particular concern is that these risks and opportunities will disproportionately impact different groups—further advantaging those who already benefit from existing power structures, while historically marginalized communities bear the brunt of the resulting harms [9]. This concern—where by default, technological systems operate within and reflect back systems of structural oppression and power—has long been discussed in the context of technology and ML systems more broadly [10, 32, 79]. To mitigate these power imbalances, there have been increasing calls for more *participation* in ML, i.e., for a broader range of people and communities to be involved in shaping whether and how systems are built and deployed [30, 61, 82, 107]. In recent years, participatory ML efforts have gained traction, spanning applications such as machine

translation for low-resourced languages [22, 50, 78], matching algorithms for food donation services [64], and news classification systems to support activists monitoring gender-related violence [100]. Alongside these endeavors, scholars have also described how participatory efforts can fall short of meaningfully shifting power to the marginalized, such as through *participation-washing* or co-optation by powerful interests [12, 26, 30, 96].

Given the scale at which foundation models are already affecting society and capturing imagination, it is critical for technologists to actively mitigate the power imbalances they produce, as well as the disproportionate harms. But **how can the benefits of participatory ML be realized within the unique affordances of foundation models?** In a classic ML setup, the downstream task and likely users are known upfront, and participatory ML approaches might look like community control over problem formulation, data collection and storage, or the design of evaluations that best reflect real-world use. However, this is much more challenging in the foundation model paradigm, where downstream use cases and stakeholders are disconnected from the model—both conceptually (the model is intended to perform well on an unbounded set of potential use cases) and practically (development primarily happens within large tech companies that are not accountable to specific communities). If valuing local expertise and context are core to participatory methods, can foundation models be meaningfully participatory?

In this paper, we investigate the proposition of participatory foundation models by first examining existing efforts that attempt to incorporate stakeholder input into foundation model development (Section 3). By analyzing them through the lens of participatory scholarship, we find consistent limitations in the ability of these mechanisms to meaningfully shift power, highlighting the tension between participation and scale. Based on these findings, we conceptualize the *participatory ceiling*: an inherent limit on the ability for impacted communities to meaningfully shape a foundation model that is intended to be almost universally applicable.

Then, we develop a blueprint for participation in foundation models that identifies more local, application-oriented opportunities to lend stakeholders meaningful agency and decision-making power (Section 4). The framework extends the carpentry metaphor of a *foundation*—an unfinished base that can support many different kinds of structures. Built on top of the foundation, a *subfloor* layer provides a level and structurally-sound base to support the top-level *surface* layer. In our framework, the subfloor layer encompasses technical infrastructure, norms, and/or governance for a grounded domain (e.g., reproductive health). It helps scope the range of potential uses to consider, lend clarity to who should be involved, and ground considerations of harm and equity in the sociohistorical context of a domain. The surface layer, then, builds on the subfloor, and encompasses specific downstream use cases (e.g., a chatbot to help provide patients with accurate information on fertility concerns). The surface layer provides opportunities for task- and locale-specific participatory engagement—while also benefiting from the domain-specific infrastructure, norms and governance developed at the subfloor layer. We walk through three case studies in clinical care, financial services, and journalism, to illustrate how this multi-layer framework can support public power and decision-making over ML in the foundation model paradigm.

2 RELATED WORK

2.1 Participatory machine learning

Participatory traditions have a long history in and outside of technology design [5, 42, 46, 69, 74, 75]. Specific motivations for participation vary, and include redistributing decision-making power to those who have less [75], learning from participants’ expertise or preferences [39, 87], or meeting epistemic goals such as procedural fairness [45]. In this paper, we focus on the promise of participatory approaches to shift power to those with less, following a range of participatory scholarship centering issues of power, agency and accountability [5, 25, 27, 47, 86, 96].

In AI and ML, recent work has applied participatory methods at various points throughout the ML lifecycle, including problem formulation, data collection and annotation, model development, evaluation, and governance [22, 50, 55, 64, 78, 92, 100]. Alongside these examples, research has also pointed out pitfalls of “participation-washing.” Sloane et al. [96] distinguish *participation as work* (e.g., user data used to train models without consent or compensation) and *participation as consultation* (e.g., one-off focus groups to elicit user preferences) from *participation as justice* (sustained and mutually beneficial relationships with communities, who co-determine if and how a model should be built). This gradation aligns with historical perspectives on the co-optation of participation, such as Arnstein’s “Ladder of Citizen Participation,” which describes the degree to which political and economic processes redistribute power to citizens [5].

Several recent papers have used additional frameworks and heuristics to understand the range of participatory ML work. Corbett et al. [26] use Arnstein’s ladder to compare and contrast the extent to which different participatory approaches redistribute power, highlighting eight case studies of each rung on the ladder. Delgado et al. [30] synthesize the ladder with other participatory scholarship to develop the “Parameters of Participation,” a framework they then use to analyze the goals and scope of 80 research papers in participatory AI. Birhane et al. [12] describe further axes of interest (reflexivity, empowerment, reciprocity, duration), using these to analyze three case studies [12]. Focusing on commercial AI labs, Groves et al. [48] interview industry practitioners, illustrating barriers that arise when attempting to use participatory approaches in industrial practice (e.g., a lack of resources or misaligned incentives).

Thus far, the participatory ML literature has focused on small-scale, application-focused case studies. This is expected, as scholarship on participatory approaches encourages context-specificity. However, it leaves us with a gap between how participatory ML is being conceptualized, and the increasing prevalence of foundation models that operate on much larger scales. To address this gap, our paper builds on the analysis and heuristics developed in these prior works to characterize the landscape of participatory efforts in the foundation model ecosystem, and present a framework for more participatory alternatives.

2.2 Foundation models

Coined in 2021, the term *foundation model* [14] refers to a paradigm of ML in which a base model (the *foundation*) is trained via self-supervised learning to compactly represent the statistical

distribution of a vast dataset (*pretraining*). Pretrained representations (also called *embeddings*) can then be adapted to downstream tasks (*fine-tuning* for specific *applications*). Pretraining-finetuning predates Bommasani et al. [14]; we adopt the foundation model terminology here to expand on its metaphor, and speak to the current era of large and centralized models, including both closed-source (e.g., OpenAI’s GPT) and open-source (e.g., DBRX, Llama) variants.

A foundation model approach offers application developers tremendous advantages—and offers foundation model developers tremendous power. Instead of training their models with as large of a dataset as they can find, store, and compute, an application developer may take an available foundation model (e.g., one trained on vast amounts of Internet text to represent the English language) and adapt it for their task, thus inheriting the knowledge representations learned during pretraining. Achieving the largest and most performant foundation model has thus become an arms race within the machine learning community. Institutions like tech companies or well-resourced universities look to amass vast stores of data—e.g., LLMs require datasets in the range of billions of words—that is typically only available at the scale of Internet-wide scrapes [9]. Training and deploying these models similarly requires massive amounts of compute, which comes with growing environmental costs [66, 99]: foundation models today can feature trillions of parameters. Nevertheless, the race accelerates. Since the initial set of foundation models cited in Bommasani et al. [14], including BERT, GPT, and CLIP, foundation models have seen rapid uptake in the tech industry, and a new wave of startups has emerged, all aiming to commercialize the best foundation models, and capture the widest share of downstream applications.

The foundation model paradigm poses fundamental challenges for the participatory ML approaches built for the task-specific era. Foundation models are inherently context-less; they are trained to simply represent a dataset, and provide baseline performance for an infinite horizon of future tasks. Moreover, due to their universal aim and the resources needed to develop them, foundation models are primarily built and controlled by large tech companies that do not have a particular investment in or accountability to any given community or domain. These departures from the task-specific paradigm pose challenges for participation at each step of ML lifecycle — from problem definition to data collection to deployment. For example, dataset auditing may have previously been informed by a domain-specific understanding of the data generation and curation processes. For a foundation model where the data is intended to inform a universally-applicable knowledge representation, specifying and finding dataset harms becomes a more nebulous problem. And this is before considering the lack of transparency into those datasets by foundation model providers [15]. Indeed, numerous harms of foundation model datasets have emerged, including LLMs that memorize and/or leak private information [17, 76], or image generation models that rely on datasets that contain hateful or even illegal content (such as child sexual abuse material) [13].

There are clear challenges to building meaningful participation into the foundation model paradigm. While fine-tuning may address some of these challenges, a fine-tuned model inherits the limitations of a foundation model, including its biases, and it is not yet clear how a community would participate in ensuring a fine-tuned model

was fit-for-purpose. Our work interrogates the possibility of participatory approaches to foundation model development, towards mitigating harm in this new scientific and industrial paradigm.

3 THE PARTICIPATORY CEILING IN FOUNDATION MODELS

We begin by reviewing proposed mechanisms for participation in foundation model development. We analyze these attempts through the lens of the Parameters of Participation, a conceptual framework developed by Delgado et al. [30] to characterize participatory initiatives in AI and ML. Based on this analysis, we argue there is a *participatory ceiling* that limits the extent to which participatory approaches can meaningfully redistribute decision-making power when directly intervening on a foundation model.

3.1 The Parameters of Participation

We use Delgado et al. [30]’s Parameters of Participation as a conceptual framework because it is focused on participatory AI efforts in particular, in contrast to more general theories or frameworks of participation. The framework articulates key *dimensions* along which participatory approaches differ. The dimensions, also framed as questions, include the **goal** (why is participation needed?), the **stakes** (what is on the table?), the **scope** (who is involved?), and the **form** (what form does participation take?). The answers to these questions are structured along a spectrum of four *modes* of participation, which span **consultation** (e.g., eliciting user preferences), **inclusion** (e.g., deliberation around specific design choices), **collaboration** (e.g., co-creation of design possibilities), and **ownership** (e.g., stakeholders shape the entire design process). This spectrum of modes reflects a long tradition in participatory scholarship that encourages engagements that cede a greater degree of decision-making power to those most directly affected by the outcome (i.e., more “meaningful participation”) [5, 42, 52].

3.2 How participatory are existing participatory foundation model efforts?

We used purposive sampling [88] to review academic and gray literature on approaches to using participatory methods or promoting public input for foundation models. Relevant examples were identified from websites of foundation model providers (e.g., OpenAI blog posts), AI/ML conference proceedings, and arXiv. We inductively clustered examples into several broad categories, including RLHF, methods that develop rulesets, guidelines or policies, and red teaming. We then applied Delgado et al.’s Parameters of Participation as an analytical lens to characterize these approaches by the degree to which they afford meaningful participation, or how they might deepen along that axis [30]. A summary of these findings is illustrated in Figure 1.

There are a few forms of human input we did not include in our analysis. First, while all development hinges on human decisions and norms, we consider participation to involve some kind of public external to the model development team, and excluded approaches that were limited to developer feedback. In addition, we do not include the generic use of human-generated data as a mechanism for human input. This data is certainly important for foundation models, which, like many ML models, rely on “human

infrastructures” of annotators, data workers, and content created by people [68]. We defer to existing literature that has showed how this mode of input is nominal and extractive [68, 71, 72, 96], and instead focus on approaches that intend to give participants more agency (e.g., red teaming) or that are unique to foundation models (e.g., RLHF).

3.2.1 Reinforcement learning with human feedback (RLHF). A primary mechanism for improving the quality of foundation model outputs is reinforcement learning with human feedback (RLHF) [6–8, 20, 56, 60, 85, 93, 95, 98, 109, 112]. As currently practiced, participants in RLHF are typically crowd-workers or contractors; most documentation does not provide further detail on recruiting strategies or demographics. The form of feedback is typically through online questionnaires where participants assign a comparative rating to two model outputs. This feedback is used to train a “preference model,” which serves as a reward function in a subsequent reinforcement learning training procedure. As it is most commonly instantiated, RLHF is limited to the *consult* mode of participation, in all dimensions. Its primary goal is to improve the quality of the model according to desiderata set by the development team. The scope of participants is limited to those who can give a prescribed quantity of discrete feedback, rather than those who can give specific kinds of expertise. The form of feedback is a single lever, typically ranking pairs of outputs. And the stakes are limited to adjusting a pre-trained model.

3.2.2 Rulesets and policies. A variety of approaches aim to synthesize principles, rules, or policies about foundation model behavior based on human input or deliberation. For example, Anthropic’s Collective Constitutional AI (CCAI) uses a public polling process to determine a ruleset (also referred to as a “constitution”) [3], which takes the form of a set of instructions for the model that reflects specific values or principles (e.g., “choose the response that is most respectful”). Participants (in this case, “a representative sample of 1,000 U.S. adults across age, gender, income, and geography”) vote on candidate principles, or submit their own; those with high consensus make it into the ruleset. Similarly, many of the projects funded by OpenAI’s “Democratic Inputs to AI” grant program develop methodologies for collecting and consolidating public input to produce a set of representative beliefs or statements [37, 41, 58, 59, 70, 94, 101, 105]. Participants are typically identified and recruited by the research team, and range from representative samples of the public to more demographically- or geographically-focused groups. The forms of participation in these approaches include voting or contributing free-form thoughts on online platforms [35, 41, 58, 59, 94], deliberation in online chat rooms [105], and focus groups or interviews [11, 70, 101]. In many cases, foundation models are further used as a tool in producing the final output—e.g., LLMs facilitate online discussions [70] or summarize participant input [94]. The outputs of these methods are used in different ways. Sometimes, as in the case of CCAI, the resulting ruleset is used to steer the model via a reinforcement learning procedure; in many other cases, there is not explicit guidance on how the resulting documentation or findings should be used.

The goal of rulesets is typically to adjust the model to reflect stakeholder preferences and values. The scope of participants is determined by the project team according to a particular value (e.g.,

representativeness). Along these dimensions, synthesizing public input into rulesets or guidelines falls under the *include* mode of participation. At the same time, the stakes are low. Participants can vote on constitutional principles, or contribute to discussions around model behavior. But how people’s inputs end up impacting the model is neither guaranteed nor transparent, mediated by complicated RL processes or LLM-facilitated analyses. And the *form* of participation—providing preferences or input during a discrete time window—reflects the fact that ultimately, participants have little say regarding the model’s impact in the world: whether it is developed, what other data it is trained on, what it may be used for, or if and how it should be deployed. Some methods aim to produce guidelines that touch on these broader questions, but to have impact, they require additional mechanisms of downstream control that are not currently in place. In short, with respect to the stakes and form of participation, these approaches remain in the *consult* mode of participation.

3.2.3 Red teaming. Red teaming, a practice developed in computer security, involves enlisting domain experts to adversarially test systems to uncover specific kinds of weaknesses. Recent efforts at foundation model governance have explored how red teaming can be applied to AI: e.g., in September 2023, OpenAI announced a call for a red teaming network, made up of external domain experts, “to help develop domain specific taxonomies of risk and evaluating possibly harmful capabilities in new systems” [83]. Here, participants are recruited by the project team for their domain expertise in predefined areas (typically, along with technical expertise), consistent with the *inclusion* mode of participation. Red teaming programs have involved individuals interacting with a model, group discussions, and shared documentation of findings from weeks of use [80, 81]. When red teaming involves deliberation and discussion with the project team, it becomes *inclusion*, which improves upon methods like online surveys (which remain at the *consultation* level), but does not yet achieve the *collaboration* mode, which would involve more durable relationships and decision-making power. Red teaming is also usually scoped to uncovering specific technical vulnerabilities. Therefore, the stakes are limited to *consultation*. Findings from a red teaming engagement might inform adjustments to the model or its documentation on harms, but these decisions are still ultimately up to the development team.

3.2.4 Domain-oriented efforts. Other approaches—while similar in many ways to creating rulesets with public input (Section 3.2.2) or red teaming (Section 3.2.3)—stem more from understanding risks that are relevant to a particular domain or group. For example, researchers have engaged stakeholders in maternal health fields [4], creative professions [53], or disability [43] and queer communities [40] about risks or opportunities for foundation models that are important to them. Most commonly, findings result in high-level guidelines for usage or harm measurement. For example, Antoniak et al. [4] propose guiding principles for using LLMs in maternal health applications. In other cases, findings shape evaluative benchmarks [40]. And in one exploratory example we encountered, participants (in this case, African artists) contributed to a collaborative licensable dataset with a payout structure for compensating creators [77].

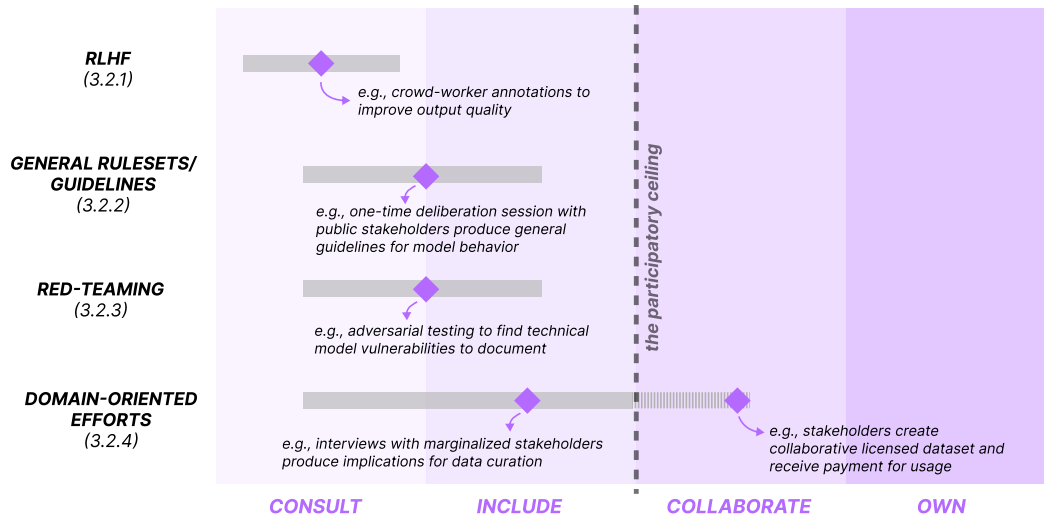


Figure 1: Different categories of participatory approaches, along with the modes of participation (as described in [30]) they currently cover. The bars are a qualitative depiction of the modes covered by each method – e.g., dimensions for red teaming approaches we reviewed fell equally under consult and include; so the range spans both modes equally. The purple diamonds are exemplars of each category. In the domain-oriented efforts category, we illustrate the possibility for a few approaches to pass the participatory ceiling with the dashed bar. We describe these exploratory approaches in Section 3.3.

For the most part, as in Sections 3.2.2 and 3.2.3, these efforts reflect the *inclusion* mode of participation along each dimension. But some aim to expand the stakes, raising issues related to data quality and acceptable use cases [4, 43], or designing interventions targeting data collection and curation [77]. These opportunities for more meaningful participation arise, in part, through the focus on a grounded domain. That said, these approaches are exploratory, and it remains challenging for domain-specific stakeholders to influence the foundation model with any decision-making power. We expand on this challenge in Section 3.3; then, in section Section 4, we propose a framework that envisions the potential for deeply domain-oriented efforts to be a crucial site for participation in the foundation model ecosystem.

3.3 The participatory ceiling

Looking across the approaches described above, we find the majority exhibit consistent limitations in the mode of participation they achieve. Why are participatory approaches thus far limited to those that lend little power to stakeholders? We argue that the intersection of primarily corporate control and context-agnostic models leads to a *participatory ceiling* on what is possible when attempting to directly intervene on a foundation model.

3.3.1 Foundation model developers currently lack incentives to share control with communities. In our analysis, the stakes of participation are most often limited to *consultation*. Participatory engagements primarily produce knowledge for model adjustments, rather than informing higher-level decisions around data sources or system purpose. Contributing to a ruleset or a red team does not guarantee one’s input winds up reshaping the final model (Section 3.2.2), and contributing to a domain-specific evaluation does not guarantee

others in one’s community will not be harmed by products the model is based on (Section 3.2.4).

We argue this lack of capacity to govern models among communities contributing to them occurs, in part, because the current ecosystem of foundation models is massively centralized, resting primarily in the hands of well-resourced technology companies that amass the data and compute to deploy models. In any setting, participatory processes rarely guarantee outcomes—e.g., in democratic voting, participants usually do not decide how and when to vote, and are not guaranteed that their position will emerge as the majority. But the primacy of *proprietary* models managed by corporate actors creates an additional layer of detachment separating public stakeholders from decision-making processes. If meaningful participation requires the distribution of *some* decision-making power, this shift is not easily managed by firms, which are primarily constituted to protect shareholder interests rather than open collaboration to societal stakeholders. In recent history, large technology firms have especially emphasized they are not liable for downstream individual uses of their products; when a firm does not have legal liability for harm, it has even less incentive to collaborate with impacted communities to identify and mitigate risks that may arise once technologies enter everyday lives. Moreover, most large organizations are risk-averse: partnerships with advocacy groups or community advisory boards create openings for the company to be publicly criticized, or exposed to federal regulation and enforcement. There are also practical challenges: contending with intellectual property claims and integrating diverse user feedback into engineering are both known challenges across responsible AI work [48].

More broadly, the methods we surveyed tend to reflect an approach to “participation” that views model builders as the arbiters

of representativeness, responsible for eliciting preferences from people and shaping the model *for* them. In CCAI, for example, Anthropic developers culled its list of crowdsourced constitutional principles by selecting only those which displayed high public agreement, and eliminating those which were more controversial (e.g., whether the model should prioritize individual or collective good) [3]. This is not to say that developers do not have expertise that should be respected; the issue is that they remain in control of deciding what counts as ‘foundational’ in a foundation model.

Participants’ lack of meaningful governance is exacerbated by corporate control of foundation models—but corporate dominance is not the sole issue here. As we argue in the next subsection and unpack further in Section 5.2, the fundamental premise of a foundation model approach assumes the need for *a* centralized entity—be it a corporation, an academic institution, a government, or an open-source community—to orchestrate model development. This core assumption separates model building from the contexts it aims to reflect, producing a disjointed supply chain of datasets in its wake.

3.3.2 Meaningful participation necessitates context-specificity, but foundation models aim for universality. Because foundation models are intended to be applicable across domains, geographies, and other contexts, most of the participatory efforts we found aimed to construct general and universally applicable guidelines. The participatory tradition, however, has historically been grounded in *context-specificity*. Prioritizing context in fields such as Participatory Action Research (PAR) stems from the acknowledgement that local stakeholders (i.e., people embedded in the day-to-day of a particular context) hold complex and valuable expertise about the needs, dynamics, and intricacies of their environment [46]. It also acknowledges that local knowledge systems are differentiated; values or norms in one context may not neatly transfer to another. The same holds for understanding how harm manifests — as described in intersectional feminist theory, the reality of marginalization is highly varied [23, 24, 28, 32]. Considering context lends concreteness to *what* harm actually looks like in a grounded domain and *who* may be disproportionately affected. Finally, from the perspective of human-AI interaction design, deliberations around abstract or unbounded capabilities pose unique challenges for user-centered design processes [110]. It becomes difficult for potential users to anticipate and reason about system behavior: e.g., empirical work has shown that the ways people react to and think about AI systems changes depending on whether they are engaged in an abstracted proxy task versus a more realistic, application-grounded task [19, 33].

When participatory efforts aim to produce universal outputs, they no longer prioritize local knowledge. Instead, they ask participants to imagine and then reason about hypothetical, distant, or abstract scenarios. Consider a high-school teacher prompted to reason about whether “being respectful” or “conveying clear intentions” is a more important value for an LLM to adhere to (both examples from CCAI [3]), versus considering the concrete trade-offs and risks of an LLM-based tutoring tool integrated into their classroom. While participants may be perfectly capable of abstraction and imagination, foundation models are subject to tremendous hype under the banner of artificial general intelligence, and their governance requires a willingness to speculate on risks that may

be quite far from participants’ lived realities—a unique challenge within human-AI interaction [110]. Such speculation might lead to greater technological awareness and literacy, but as Harrington et al. [51] argue, asking marginalized communities to engage in the “*blue-sky ideation*” of technology design risks ultimately frustrating underserved individuals. By focusing instead on the real harms and concerns people are presently experiencing, developers can limit the material and affective demands of participation [34], acknowledge participants as experts on local knowledge systems, and create outputs that more closely reflect actual downstream needs.

Pushing the participatory ceiling. Notably, in our review of the current state of play, a few examples out of those surveyed stood out in straining against this ceiling. For example, UbuntuAI [36, 77] proposes a system that responds to the expropriation by foundation models of African artists’ intellectual property by collaboratively creating a licensable dataset of their work. While the project is still exploratory, its approach touches on the *collaboration* mode of participation, with ongoing co-creation of the dataset and a compensation structure that lends itself to community control. Here, gains on the parameters of participation are realized, in part, through a focus on a specific domain. In the next section, we draw out the implications of examples like this into a broader framework for more meaningful participation into foundation models.

4 A BLUEPRINT FOR MORE PARTICIPATORY FOUNDATION MODELS

We propose a three-layer framework to enable more effective public participation in foundation models. As discussed in the prior section, we claim there is a fundamental tension between the scale and generalizability of foundation models, and the power that local communities can wield in shaping them. Our framework addresses this limitation by building in additional layers for participation at more local, application-oriented scales—the *subfloor* layer and the *surface* layer. The naming of our framework builds on the metaphor of a *foundation*: an unfinished base upon which many different kinds of structures might be built. The *subfloor* is a stable and level ground built on the foundation, which provides a structurally sound base for any number of top-level *surfaces*.

In this section, we first describe the three-level blueprint. We then delve into three case studies, illustrating opportunities for participation it affords when applied to foundation model usage in healthcare, banking, and journalism. Each case study, while hypothetical, is grounded in current uses of LLMs, real-world organizations, and discussions with a domain expert in each context. We intend the case studies to serve as starting points for further iteration and testing out in practice.

4.1 The Blueprint

4.1.1 The Foundation Layer. The foundation layer is where a base model—not yet intended to be domain- or application-specific—is created and maintained. Our analysis of current attempts at participation in foundation models (section Section 3) shows there are already numerous approaches used and stakeholders involved; still, there are certainly avenues to deepen this work with respect to participatory desiderata. For example, foundation model developers




Figure 2: An illustration of the three layers proposed in our blueprint.

might pursue forms of participation that are ongoing and sustainable, or widen the scope of interventions to those that impact other parts of the lifecycle beyond the steering of the model.

That said, we contend that participation at the foundation layer is *not, by itself, enough*, due to inherent limitations in what is possible at scale (Section 3.3). The subfloor and surface layers of our framework are instead grounded in context, and necessary for tractable participation. In this setup, the foundation layer takes on additional responsibilities: ensuring the base functionality required by the network of subfloors and surfaces, and remaining porous to issues or demands raised by them (e.g., compensating and attributing creative workers for the use of their work, as in Section 4.3).

4.1.2 The Subfloor Layer. The subfloor layer encompasses technical infrastructure, norms and governance for a grounded domain. This might include fine-tuned models, curated datasets, auditing process, mechanisms for recourse, standards-setting procedures, and shared governance structures (we expand on specific instantiations of the subfloor layer in Sections 4.2-4.4). Importantly, participation at the subfloor mitigates the pitfalls of generality (Section 3.3.2) by being grounded in a specific domain, and the barriers of corporate incentives (Section 3.3.1) via ownership by entities such as nonprofits, local governments, community advocacy organizations, or unions. Context-specificity makes more clear *who* should be involved: for instance, in the healthcare domain, participatory efforts might engage medical practitioners, health equity scholars, health policy makers, patient advocacy groups, and voices from communities we know to be marginalized within the current healthcare system [91]. It also becomes easier for people and communities to participate. Rather than coming up with abstract, universal desiderata (“the model should produce outputs that are good for humanity”), the scope of harms to consider is bounded by a particular domain. Stakeholders can thus contribute their concrete expertise and lived experience of the domain (“the model should not replicate known biases in pain assessment”) [e.g., 90, 111].

4.1.3 The Surface Layer. The surface layer corresponds to a specific downstream use case, built on top of the subfloor layer. This might encompass a specific tool or system, along with task-specific datasets(s), documentation, and mechanisms for accountability and refusal. Participation at the surface layer likely looks similar to existing case studies of participatory ML. Affected communities can shape problem formulation, and co-determine whether a foundation model-based solution is desirable in their context. For example, a

specific healthcare facility might want a tool to share information to patients on reproductive health. Local patients and care providers, along with reproductive justice experts, patient advocacy groups, and other stakeholders, might collectively decide on the appropriateness of an LLM-based tool. If the outcome of this deliberation is to move forward, it might involve participatory data collection and annotation that is specific to the task and locality; this data can be used for fine-tuning the subfloor model or for context-specific validation. Participation at this layer might further involve avenues for public accountability and recourse, determined and designed by the affected stakeholders.

Importantly, the existence of the subfloor layer means that the burden of domain-specific validation and specialization is not put entirely on stakeholders developing a particular use case. For example, our local healthcare facility can inherit the baseline equity assurances put in at the subfloor layer (“the model should not replicate known biases in pain assessment”), while building specialized constraints for their patient population (“the model should accurately answer questions that are most common among our patient population”). We elaborate on a related example in Section 4.2.

4.2 Case study 1: Clinical care

Our first case study illustrates the participatory opportunities provided by the nested subfloor-surface structure of our framework. We consider the development of LLM-based tools that aim to address growing administrative burdens clinicians face due to electronic health records (EHRs) [44]. Such tools may transcribe the doctor-patient interaction with a speech-to-text model; these transcriptions may then be used for clinical note summarization. This paradigm (also called “*ambient intelligence*”) is increasingly used in dominant EHRs like Epic [62].

The problem. A tool for transcribing medical interactions must robustly handle a diversity of languages, accents, and idioms, without inadvertently reproducing harmful stereotypes about specific minoritized groups, or compromising medical accuracy. There are also issues of trust and privacy—patients and providers must be able to verify a summary before it becomes part of the record, and to opt out if they do not want the tool listening in the exam room.

Participatory approaches could improve these systems with respect to equity and trust. However, these concerns are hard to address at the foundation layer: they require engaging with complex and domain-specific issues at the intersection of dialect, geography, race and ethnicity, service provision, and medical history. Still, if the responsibility of thoroughly auditing a model were to rest with individual clinics, they would each likely run into data limitations (e.g., for under-represented dialects), and may needlessly duplicate effort across different settings.

The opportunity and participatory methods. Here, the subfloor layer could encompass a constellation of patient advocacy organizations who pool data collected with communities to represent various dialects and accents, as well as identify key risks to anticipate and audit. For example, a patient advocacy organization in a predominantly Black community might collect data representing African American Language (AAL) [29]. They might further identify Black maternal health as a key equity risk, and elevate

the perspectives of Black midwives to shape the data collection process and identify specific harms to audit. Similarly, an organization representing an Asian-American immigrant community might collect data representing the different kinds of accents and idioms in that community, and flag Asian-American mental health as an important equity risk.

Data collected across communities could be used to fine-tune a shared speech-to-text model that performs well on speech from both of these communities, and to design ongoing audits of this model for both sets of identified risks. The pooled data could include data and deliberation from as many communities as is relevant to the subfloor's identified constituency. Additional expertise (e.g., on legal, privacy, or health equity issues) could also shape technical infrastructure and broader guidelines for use. Such a convening might build, for example, on the format proposed in Antoniak et al. [4], where healthcare workers and birthing people collaborated on guiding principles for the use of NLP in maternal health.

At the surface layer, then, resources made available at the subfloor (e.g., guidelines for use, audits across equity dimensions) can inform concrete deliberation at a specific care site around whether an LLM-based transcription tool is desired in that context. The surface layer thus provides an important site of refusal, for the opt-out and oversight mechanisms important to ensure trust. If the outcome of deliberation is to move forward, the surface layer provides an opportunity to incorporate context-specific logic and validation around, e.g., the resources, procedures, and patient populations at a given hospital, while building on the technical infrastructure and trust established at the subfloor.

4.3 Case Study 2: Journalism

In this case study, we examine opportunities for collective stakeholder power at the subfloor to shape changes at the foundation layer. We consider this opportunity within the context of foundation models' impact on professional writers and other creatives. We focus on those employed by newsrooms—organizations which themselves assert a copyright claim over their material.

The problem. In January 2024, the New York Times filed a lawsuit against OpenAI and Microsoft for the unauthorized use of the publication's copyrighted material to train AI systems [49]. Several other organizations and individuals have filed similar complaints objecting to the wide-scale collection and use of copyrighted material without compensation or credit [2, 97]. The foundation model-based systems in question further risk leading to labor displacement of the same people whose work was expropriated to create it. When trying to enact changes in the data used at the foundation layer (for example, by filing individual lawsuits), each individual creator lacks enough leverage to shift data collection norms. The prohibitive cost of challenging this practice in effect establishes a precedent that data misuse is permissible.

The opportunity and participatory methods. The subfloor layer presents an opportunity for entities with copyright concerns to partner with each other, collectively asserting copyright claims over their data and establishing a license agreement for compensation over time. Such an effort might be housed within an existing collective action organization, such as the NewsGuild, a primary

union for media workers. Inspired by Nayebare et al. [77], the subfloor could support the creation of a curated dataset of license-ready contributions, alongside organizational infrastructure that would keep it updated with successively published material. Collectively, publishers and creators could demand the expungement or filtering of their data from existing datasets used to train foundation models. They could then establish an ongoing data use agreement and payment structure such that the data cannot be used without compensation. Such a setup might be supported by technical approaches such as SILO [73], which could enable creators to opt in/out of a datastore used only during inference and be attributed for the use of their work. In this case, the subfloor leverages collective action to influence changes at the foundation layer that reflect domain-specific needs and requirements. In its absence, individual lawsuits may both duplicate efforts and lack the power needed to demand a compensation agreement.

4.4 Case Study 3: Financial services

Finally, we consider how the subfloor layer could provide avenues for participation in contexts such as finance, where major financial firms are exploring the use of LLMs for improving fraud risk scoring algorithms [65, 67, 103]. While machine learning techniques are already used for assigning risk scores to transactions, LLMs are beginning to be used to identify more subtle patterns in fraudulent transactions, and may introduce new and unforeseen risks.

The problem. As a task, fraud detection carries high-stakes fairness and equity risks. To avoid reputational damage, credit card issuing banks and payment networks aim to avoid declining transactions frivolously or inequitably. They are bound by legal requirements to protect equity as well: in the U.S., the Federal Trade Commission and the Equal Credit Opportunity Act hold credit card companies accountable for credit discrimination.

Robust reporting mechanisms can help to catch new risks and unfair patterns of credit decline in deployment; however, competitiveness concerns incline firms to conduct system evaluation strictly behind closed doors. As a result, potential harms like false positives in fraud detection and inequitable credit service provision remain under-examined—particularly by the communities who would be most affected and the organizations that represent them. Because these concerns are deeply domain-specific, addressing them at the foundation layer proves difficult. If ensuring equity is left to individual providers, however, we risk leaving affected communities to navigate consequences downstream of a complex chain of systems that each implement fraud detection technology differently.

The opportunity and participatory methods. The subfloor in this case could focus on robust on-the-ground reporting and recourse mechanisms to catch new or unexpected patterns of discrimination. For example, participatory reporting mechanisms might include: (1) *in vivo* field testing “spot-checks” of the payment system with end-users from high risk groups; (2) reporting hotlines by which the broader public can escalate their experiences and concerns; and (3) formal reports of concerns and common experiences from various constituencies.

Working groups of experiential experts and advocates can then use this input to inform foundation model governance in financial services—identifying larger patterns in model behavior, documenting harms and vulnerabilities, and producing updated guidance. These working groups can be organized around an institutional body already common to financial services providers, such as the Payment Card Industry (PCI) Security Standards Council, which convenes and issues guidance to financial services firms on information security and privacy protections. Its remit is also inclusive of best practices for new tools and technologies, such as generative AI. Guidance produced by working groups can inform PCI member firms’ in-house counsel, privacy and risk management personnel, and regulatory teams. These audiences can then translate findings for their own firm’s product teams.

Together, these reporting and organizational infrastructures would create avenues for more robust, real-world, and variegated external input to shape surface-layer systems, in spite of a highly closed and rigid overall environment. It would also serve as an “early warning system” when cardholder communities or merchants face an uptick in declined transactions because of spurious or discriminatory correlations made by an LLM-backed product.

5 DISCUSSION

We have presented a conceptual contribution for how meaningful participation can shape the foundation model lifecycle. Our blueprint supports shifting power within the foundation model ecosystem through a focus on context. We hope our work helps organize the FAcCT community’s thinking around this new and rapidly growing paradigm. But participation is not meant as a panacea, and our framework inherits its familiar risks and limitations: diffusion of accountability (Section 5.1), power asymmetry and co-optation (Section 5.2), and labor for already disempowered stakeholders (Section 5.3). In this section, we unpack each of these limitations, and outline future work for the FAcCT community addressing these risks.

5.1 Accountability through the subfloor layer

Our case studies highlight a wide range of participatory mechanisms afforded by the subfloor, including collaborative data collection and model auditing (Section 4.2), collective action for data ownership and refusal (Section 4.3), and avenues for on-the-ground reporting and recourse (Section 4.4). In each case, the subfloor provides a route for meaningful participation by creating more tractable accountability relationships between the foundation model provider and downstream surface layers. Rather than engagement that is atomized and discrete (as we see with efforts at the foundation layer), because it is much closer to—if not directly controlled by—affected communities, the subfloor enables participation that is *deliberative* and *longitudinal*. It breaks the participatory ceiling through being grounded in context and primarily owned by entities such as nonprofits, local governments, community advocacy organizations, or unions. Our framework shifts the locus of responsibility from a central model creator to an ecosystem of subfloors, enabling a domain-specific actor to become accountable for a system. At the same time, collective action by subfloors, who together can represent a substantial stake in an industry, becomes harder for a

foundation layer to brush off. The impact of this engagement scales across surface-layer use cases, which provide sites for additional, task-specific participatory processes while inheriting trust built at the subfloor.

Still, a layered framework diffuses accountability for harm in ways that scholars have yet to resolve. The relationship between the foundation layer and the subfloor layer may resemble that of an app store provider to its app developers. Apple’s Terms of Service define guidelines for apps to satisfy to be hosted in their marketplace; a failure to meet these requirements results in removal from its App Store, and thus a limited channel for distribution to iOS users. But how does a foundation model become responsible (or not) for harms it enables at the subfloor and surface layers? And what responsibilities do subfloors have for harms they may inherit from foundations, or host and pass on to surfaces on their own? More participatory and community-controlled infrastructure may enable bad actors to fine-tune foundation models for abusive ends [84]. The stakes of this debate are rising: as a stark example, it remains unclear who ought to be responsible for an LLM-powered chatbot encouraging a person to harm themselves [108]. Understanding where liability for harm can and should rest will be key to enabling participatory infrastructure while mitigating its drawbacks.

5.2 Centralization and the limits of transparency

Our analysis identified the inherent centralization in foundation model development as a key bulwark of the participatory ceiling (Section 3.3). Whether a venture-backed startup, a technology giant, or a well-resourced academic or government institution, foundation model developers are an obligatory passage point for influencing the operation of a given foundation model. Crucially, this is a feature of the foundation model paradigm that remains consistent across open-source and closed-source approaches. Many researchers have advocated for transparency and openness in foundation models: whether model weights and data are available for inspection, or whether a corporate or nonprofit entity hosts a given model [16]. We contend, however, the issue is more in the fundamental premise that models can be disconnected from meaningful social governance, even as they aim to represent the complexity of language, moral reasoning, and other human social interactions. While the precedents around radical transparency set by free and open-source software communities are an important first step towards meaningful participation in *development*, Widder et al. [106] have described how in practice, the sheer scale of resources needed to *deploy* a foundation model mean only a handful of well-resourced institutions are positioned to engage.

It remains to be seen to what extent this inherent power asymmetry can be allayed by a subfloor. A subfloor with a large remit (e.g., to represent the English language) could recapitulate this centralization of power. Moreover, like all participatory mechanisms, the subfloor risks participation-washing by powerful actors in service of their own aims. Ahmed [1] warns that participation can be a means not of advancing collective well-being, but of powerful actors creating a heavily circumscribed mechanism for co-optation that impedes real change. At best, our blueprint raises the stakes

for participation-washing at scale, and also creates a more effective target for advocacy and organizing.

To realize this best-case, we see ample future work understanding how to scope a subfloor or a surface so it remains accountable to communities' needs. In our clinical care case study (Section 4.2), for example, a subfloor must take on the responsibility of coordinating across different communities, each with their own goals and capacities, to ensure equitable performance and risk mitigation in a common speech-to-text model. What constitutes a subfloor's domain, and how should the boundaries between subfloors be determined? How can we collectively place constituencies between subfloors, and ensure common deliberative processes among them? Technical advancements in securely pooling data and training shared models are also important, and could build on existing approaches like federated learning and open science.¹ Progress on these fronts will help create the healthy ecosystem of subfloors and surfaces we need to establish decentralization and accountability in foundation models' future.

5.3 Ownership, refusal, and the burden of participation

Finally, there is the issue of how to make meaningful participation manageable for stakeholders—especially those who may already face marginalization and disempowerment. Better opportunities for participation still require time and energy of people who may not want to donate their time to governance efforts; particularly since so much of foundation model governance involves the labor of collecting and managing large-scale datasets. For participation in foundation models to achieve lofty aims of redistributing power to the marginalized, we need further work updating mechanisms for data stewardship and consent for the foundation model era.

These mechanisms are well-suited to subfloor and surface layer interventions. We are encouraged to see growing interest in how individuals and communities can exercise agency over the data collection and modeling processes underlying machine learning, via principles ranging from data refusal [113] and participatory data stewardship [54, 102] to AI contestability [31, 57]. The need for individuals to have control over the data that flow through foundation model-based systems has also begun to motivate new model architectures: e.g., techniques like SILO are explicitly motivated by the need to isolate sensitive data from a model, so it is used in inference but not training, and so individuals have the agency to remove their data from what the model can use [73]. In our framework, we envision that surface layer interventions could offer a space for individual data refusal; and subfloor layer interventions could offer a space for collective contestation around data usage.

Importantly, we hope that our framework leads to more participatory foundation models—but also lends communities greater agency in refusing the use foundation models at all. We envision an important part of participation at the surface layer, for example, involving deliberations about if a foundation model is the right or desired approach to a given problem. Domain-specific guidelines or audits at the subfloor layer can help ground this deliberation,

but do not guarantee that the outcome will be to move forward. In other words, the blueprint is not an endorsement or guarantee of foundation model usage. Rather, it creates more participatory alternatives, but holds that control around downstream usage should always rest with those closest to a specific context.

6 CONCLUSION

In this paper, we examined the intersection between participatory approaches and foundation models. Because foundation models are developed for use across a wide array of settings, any individual group seeking to provide input into the foundation model or mitigate harms on-the-ground must vie for influence amid a vast number of similarly impacted parties for the time and resources of a single developer firm. This convergence of scale and power asymmetry exacerbates the challenges of shaping foundation models with public input. To better support meaningful participation over foundation models, we define an organizational-level intervention that could address this challenge: a “subfloor” layer supports more local, application-oriented opportunities for meaningful participation in the form of shared technical infrastructure, norms, and governance for a grounded domain such as journalism or financial services. The subfloor layer scopes the range of potential harms to consider and affords individual participants more concrete avenues for intervention, scaling that engagement to downstream surface layer use cases.

ACKNOWLEDGMENTS

We would like to thank our colleagues, peer reviewers, and subject matter experts who have been so generous with their feedback and insight. In particular, we are grateful to Robert Lewis, Seth Lazar, Jon Kleinberg, Mark Sendak, Merill Fabry, Albert Gehami, the students and faculty of Australia National University's School of Cybernetics, and presenters at the Stanford Sociotechnical AI Safety workshop. This work is also supported in part by a Microsoft Research PhD Fellowship (ET) and the Principles G.I. Coffee House in Gowanus, Brooklyn.

REFERENCES

- [1] Alex Ahmed. 2020. We Will Not Be Pacified Through Participation. *Tech Otherwise* 15, 10 (2020), 2020.
- [2] Alexandra Alter and Elizabeth A. Harris. 2023. Franzen, Grisham and Other Prominent Authors Sue OpenAI. <https://www.nytimes.com/2023/09/20/books/authors-openai-lawsuit-chatgpt-copyright.html>
- [3] Anthropic. 2023. *Collective Constitutional AI: Aligning a Language Model with Public Input*. Technical Report. <https://www.anthropic.com/index/collective-constitutional-ai-aligning-a-language-model-with-public-input>
- [4] Maria Antoniak, Aakanksha Naik, Carla S Alvarado, Lucy Lu Wang, and Irene Y Chen. 2023. Designing Guiding Principles for NLP for Healthcare: A Case Study of Maternal Health. *arXiv preprint arXiv:2312.11803* (2023).
- [5] Sherry R Arnstein. 1969. A ladder of citizen participation. *Journal of the American Institute of planners* 35, 4 (1969), 216–224.
- [6] Amanda Askill, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Jackson Kernion, Kamal Ndousse, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, and Jared Kaplan. 2021. A General Language Assistant as a Laboratory for Alignment. *arXiv:2112.00861* [cs.CL]
- [7] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askill, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared

¹e.g., in clinical care, Nightingale (<https://www.ngsci.org/>), Observational Health Data Sciences and Informatics (OHDSI) (<https://www.ohdsi.org/>), and the Patient-Led Research Collaborative (<https://patientresearchcovid19.com/>).

- Kaplan. 2022. Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback. *arXiv:2204.05862* [cs.CL]
- [8] Michiel Bakker, Martin Chadwick, Hannah Sheahan, Michael Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matt Botvinick, et al. 2022. Fine-tuning language models to find agreement among humans with diverse preferences. *Advances in Neural Information Processing Systems* 35 (2022), 38176–38189.
- [9] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 610–623.
- [10] Ruha Benjamin. 2019. *Race after technology*. Polity.
- [11] Stevie Bergman, Nahema Marchal, John Mellor, Shakir Mohamed, Iason Gabriel, and William Isaac. 2024. STELA: a community-centred approach to norm elicitation for AI alignment. *Scientific Reports* 14, 1 (2024), 6616.
- [12] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to the people? opportunities and challenges for participatory AI. *Equity and Access in Algorithms, Mechanisms, and Optimization* (2022), 1–8.
- [13] Abeba Birhane, Vinay Prabh, Sang Han, Vishnu Naresh Boddeti, and Alexandra Sasha Luccioni. 2023. Into the LAIONs Den: Investigating Hate in Multimodal Datasets. *arXiv preprint arXiv:2311.03449* (2023).
- [14] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
- [15] Rishi Bommasani, Kevin Klyman, Shayne Longpre, Sayash Kapoor, Nestor Maslej, Betty Xiong, Daniel Zhang, and Percy Liang. 2023. The foundation model transparency index. *arXiv preprint arXiv:2310.12941* (2023).
- [16] Rishi Bommasani, Kevin Klyman, Shayne Longpre, Sayash Kapoor, Nestor Maslej, Betty Xiong, Daniel Zhang, and Percy Liang. 2023. The foundation model transparency index. *arXiv preprint arXiv:2310.12941* (2023).
- [17] Hannah Brown, Katherine Lee, Fatemehsadat Mirehghallah, Reza Shokri, and Florian Tramèr. 2022. What does it mean for a language model to preserve privacy?. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 2280–2292.
- [18] Erik Brynjolfsson, Danielle Li, and Lindsey R Raymond. 2023. *Generative AI at work*. Technical Report. National Bureau of Economic Research.
- [19] Zana Bucinca, Phoebe Lin, Krzysztof Z Gajos, and Elena L Glassman. 2020. Proxy tasks and subjective measures can be misleading in evaluating explainable ai systems. In *Proceedings of the 25th international conference on intelligent user interfaces*. 454–464.
- [20] Paul Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2023. Deep reinforcement learning from human preferences. *arXiv:1706.03741* [stat.ML]
- [21] Jan Clusmann, Fiona R Kolbinger, Hannah Sophie Muti, Zunamys I Carrero, Jan-Niklas Eckardt, Narmin Ghaffari Laleh, Chiara Maria Lavinia Löffler, Sophie-Caroline Schwarzkopf, Michaela Unger, Gregory P Veldhuizen, et al. 2023. The future landscape of large language models in medicine. *Communications Medicine* 3, 1 (2023), 141.
- [22] Donavyn Cofey. 2021. Maori are trying to save their language from Big Tech. *Wired* (2021).
- [23] Combahee River Collective. 1977. The Combahee river collective statement. (1977).
- [24] Patricia Hill Collins. 2019. *Intersectionality as critical social theory*. Duke University Press.
- [25] Bill Cooke and Uma Kothari. 2001. *Participation: The new tyranny?* Zed books.
- [26] Eric Corbett, Emily Denton, and Sheena Erete. 2023. Power and Public Participation in AI. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization* (<conf-loc>, <city>Boston</city>, <state>MA</state>, <country>USA</country>, </conf-loc>) (EAAMO '23). Association for Computing Machinery, New York, NY, USA, Article 8, 13 pages. <https://doi.org/10.1145/3617694.3623228>
- [27] Sasha Costanza-Chock. 2020. *Design justice: Community-led practices to build the worlds we need*. The MIT Press.
- [28] Kimberlé Williams Crenshaw. 2013. Mapping the margins: Intersectionality, identity politics, and violence against women of color. In *The public nature of private violence*. Routledge, 93–118.
- [29] Nicholas Deas, Jessi Grieser, Shana Kleiner, Desmond Patton, Elsbeth Turcan, and Kathleen McKeown. 2023. Evaluation of African American Language Bias in Natural Language Generation. *arXiv preprint arXiv:2305.14291* (2023).
- [30] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The participatory turn in ai design: Theoretical foundations and the current state of practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–23.
- [31] Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, Hilary Nicole, and Morgan Klaus Scheuerman. 2020. Bringing the people back in: Contesting benchmark machine learning datasets. *arXiv preprint arXiv:2007.07399* (2020).
- [32] Catherine D'ignazio and Lauren F Klein. 2023. *Data feminism*. MIT press.
- [33] Finale Doshi-Velez and Been Kim. 2018. Considerations for evaluation and generalization in interpretable machine learning. *Explainable and interpretable models in computer vision and machine learning* (2018), 3–17.
- [34] Paul Dourish, Christopher Lawrence, Tuck Wah Leong, and Greg Wadley. 2020. On being iterated: The affective demands of design participation. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [35] Joe Edelman, Oliver Klingefjord, Ivan Vendrov, and Ellie Hain. 2023. Democratic Fine-tuning. *Meaning Alignment Institute* (2023).
- [36] Ron Eglash, Kwame P Robinson, Audrey Bennett, Lionel Robert, and Mathew Garvin. 2024. Computational reparations as generative justice: Decolonial transitions to unalienated circular value flow. *Big Data & Society* 11, 1 (2024), 20539517231221732.
- [37] Tyna Eloundou and Teddy Lee. 2024. Democratic inputs to AI grant program: lessons learned and implementation plans. <https://openai.com/blog/democratic-inputs-to-ai-grant-program-update>
- [38] Andy Extance. 2023. ChatGPT has entered the classroom: how LLMs could transform education. *Nature* 623, 7987 (2023), 474–477.
- [39] Michael Feffer, Michael Skirpan, Zachary Lipton, and Hoda Heidari. 2023. From preference elicitation to participatory ml: A critical survey & guidelines for future research. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 38–48.
- [40] Virginia Felkner, Ho-Chun Herbert Chang, Eugene Jang, and Jonathan May. 2023. WinoQueer: A Community-in-the-Loop Benchmark for Anti-LGBTQ+ Bias in Large Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 9126–9140.
- [41] Sara Fish, Paul Gölz, David C Parkes, Ariel D Procaccia, Gili Ruskak, Itai Shapira, and Manuel Wüthrich. 2023. Generative Social Choice. *arXiv preprint arXiv:2309.01291* (2023).
- [42] Archon Fung. 2006. Varieties of participation in complex governance. *Public administration review* 66 (2006), 66–75.
- [43] Vinita Gadiraju, Shaun Kane, Sunipa Dev, Alex Taylor, Ding Wang, Emily Denton, and Robin Brewer. 2023. "I wouldn't say offensive but...": Disability-Centered Perspectives on Large Language Models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 205–216.
- [44] Atul Gawande. 2018. Why doctors hate their computers. *The New Yorker* 12 (2018).
- [45] Jerald Greenberg and Robert Folger. 1983. Procedural justice, participation, and the fair process effect in groups and organizations. In *Basic group processes*. Springer, 235–256.
- [46] Davydd J Greenwood and Morten Levin. 2006. *Introduction to action research: Social research for social change*. SAGE publications.
- [47] Judith Gregory. 2003. Scandinavian approaches to participatory design. *International Journal of Engineering Education* 19, 1 (2003), 62–74.
- [48] Lara Groves, Aidan Peppin, Andrew Strait, and Jenny Brennan. 2023. Going public: the role of public participation approaches in commercial AI labs. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1162–1173.
- [49] Michael M. Grynbaum and Ryan Mac. 2023. The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work. <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>
- [50] Karen Hao. 2022. A new vision of artificial intelligence for the people. *MIT Technology Review* (2022).
- [51] Christina Harrington, Sheena Erete, and Anne Marie Piper. 2019. Deconstructing community-based collaborative design: Towards more equitable participatory design engagements. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–25.
- [52] IAP2. 2024. IAP2 Spectrum of Public Participation. https://cdn.ymaws.com/www.iap2.org/resource/resmgr/pillars/Spectrum_8.5x11_Print.pdf
- [53] Nanna Inie, Jeanette Falk, and Steve Tanimoto. 2023. Designing Participatory AI: Creative Professionals' Worries and Expectations about Generative AI. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [54] Ada Lovelace Institute. 2021. Participatory data stewardship. *Ada Lovelace Institute* (2021). <https://www.adalovelaceinstitute.org/report/participatory-data-stewardship/>
- [55] Michael Katell, Meg Young, Dharma Dailey, Bernease Herman, Vivian Guetler, Aaron Tam, Corinne Bintz, Daniella Raz, and PM Krafft. 2020. Toward situated interventions for algorithmic equity: lessons from the field. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 45–55.
- [56] Hannah Rose Kirk, Andrew M Bean, Bertie Vidgen, Paul Röttger, and Scott A Hale. 2023. The past, present and better future of feedback learning in large language models for subjective human preferences and values. *arXiv preprint arXiv:2310.07629* (2023).
- [57] Daniel N Kluttz, Nitin Kohli, and Deirdre K Mulligan. 2022. Shaping our tools: Contestability as a means to promote responsible algorithmic decision making in the professions. In *Ethics of Data and Analytics*. Auerbach Publications, 420–428.

- [58] Andrew Konya, Lisa Schirch, Colin Irwin, and Aviv Ovadya. 2023. Democratic Policy Development using Collective Dialogues and AI. *arXiv preprint arXiv:2311.02242* (2023).
- [59] Alex Krasodomski-Jones, Carl Miller, Flynn Devine, Jia-Wei (Peter) Cui, and Shu Yang Lin. 2023. vTaiwan and Chatham House: Bridging the Recursive Public. (2023).
- [60] Julia Kreutzer, Stefan Riezler, and Carolin Lawrence. 2021. Offline Reinforcement Learning from Human Feedback in Real-World Sequence-to-Sequence Tasks. In *Proceedings of the 5th Workshop on Structured Prediction for NLP (SPNLP 2021)*, Zornitsa Kozareva, Sujith Ravi, Andreas Vlachos, Priyanka Agrawal, and André Martins (Eds.). Association for Computational Linguistics, Online, 37–43. <https://doi.org/10.18653/v1/2021.spnlp-1.4>
- [61] Bogdan Kulynych, David Madras, Smitha Milli, Inioluwa Deborah Raji, Angela Zhou, and Richard Zemel. 2020. Participatory approaches to machine learning. In *International Conference on Machine Learning Workshop*, Vol. 7.
- [62] Heather Landi. 2023. Epic is going all in on generative AI in health-care. Here's why health systems are eager to test-drive it. (may 2023). <https://www.fiercehealthcare.com/health-tech/epic-moves-forward-bring-generative-ai-healthcare-heres-why-handful-health-systems-are>
- [63] Katherine Lee, A Feder Cooper, and James Grimmelmann. 2023. Talkin' Bout AI Generation: Copyright and the Generative-AI Supply Chain. *arXiv preprint arXiv:2309.08133* (2023).
- [64] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction 3*, CSCW (2019), 1–35.
- [65] David Lapeska. [n. d.]. AI-DRIVEN FRAUD IS BANKING'S NEXT GREAT RISK. Bank Director. <https://www.bankdirector.com/article/ai-driven-fraud-is-bankings-next-great-risk/>.
- [66] Alexandra Sasha Luccioni, Yacine Jernite, and Emma Strubell. 2023. Power Hungry Processing: Watts Driving the Cost of AI Deployment? *arXiv preprint arXiv:2311.16863* (2023).
- [67] Mastercard. [n. d.]. Detect: Decision Intelligence. Online. <https://www.mastercard.com/globalrisk/en/resources/all-resources/detect.html>.
- [68] Alexandra Mateescu and Madeleine Elish. 2019. AI in context: the labor of integrating new technologies. *Data & Society Research Institute* (2019).
- [69] Alice McIntyre. 2007. *Participatory action research*. Sage publications.
- [70] Gemma B. Mendoza, Gilian Uy, Don Kevin Hapal, Ogoy San Juan, and Maria Ressa. 2023. Making AI Transparent and Accountable by Rappler. (2023).
- [71] Milagros Miceli and Julian Posada. 2022. The Data-Production Dispositif. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 460 (nov 2022), 37 pages. <https://doi.org/10.1145/3555561>
- [72] Milagros Miceli, Martin Schuessler, and Tianling Yang. 2020. Between subjectivity and imposition: Power dynamics in data annotation for computer vision. *Proceedings of the ACM on Human-Computer Interaction 4*, CSCW2 (2020), 1–25.
- [73] Sewon Min, Suchin Gururangan, Eric Wallace, Hannaneh Hajishirzi, Noah A Smith, and Luke Zettlemoyer. 2023. Silo language models: Isolating legal risk in a nonparametric datastore. *arXiv preprint arXiv:2308.04430* (2023).
- [74] Meredith Minkler and Nina Wallerstein. 2011. *Community-based participatory research for health: From process to outcomes*. John Wiley & Sons.
- [75] Michael J Muller and Sarah Kuhn. 1993. Participatory design. *Commun. ACM* 36, 6 (1993), 24–28.
- [76] Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A Feder Cooper, Daphne Ippolito, Christopher A Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. 2023. Scalable extraction of training data from (production) language models. *arXiv preprint arXiv:2311.17035* (2023).
- [77] Michael Nayebara, Ron Eglash, Ussen Kimanuka, Rehema Baguma, J. Mounsey, and C. Maina. 2023. Interim Report for Ubuntu-AI: A Bottom-up Approach to More Democratic and Equitable Training and Outcomes for Machine Learning. OpenAI Foundation Democratic Inputs for AI.
- [78] Wilhelmina Nekoto, Vukosi Marivate, Tshinondiwa Matsila, Timi Fasubaa, Taiwo Fagbohunbe, Solomon Oluwole Akinola, Shamsuddeen Muhammad, Salomon Kabongo Kabenamualu, Salomey Osei, Freshia Sackey, Rubungo Andre Niyongabo, Ricky Macharm, Perez Ogayo, Orevaoghene Ahia, Musie Meressa Berhe, Mofetoluwa Adeyemi, Idris Abdulkadir Dangana, Herman Kamper, Hady Elshahar, Goodness Duru, Ghollah Kioko, Murhabazi Espoir, Elan van Biljon, Daniel Whitenack, Christopher Onyefuluchi, Chris Chinenye Emezue, Bonaventure F. P. Dossou, Blessing Sibanda, Blessing Bassey, Ayodele Olabiyi, Arshath Ramkilwan, Alp Öktem, Adewale Akinfaderin, and Abdallah Bashir. 2020. Participatory Research for Low-resourced Machine Translation: A Case Study in African Languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 2144–2160. <https://doi.org/10.18653/v1/2020.findings-emnlp.195>
- [79] Safiya Umoja Noble. 2018. *Algorithms of oppression*. New York university press.
- [80] OpenAI. 2022. DALL-E 2 Preview - Risks and Limitations. <https://github.com/openai/dalle-2-preview/blob/main/system-card.md>
- [81] OpenAI. 2023. GPT-4 System Card. <https://cdn.openai.com/papers/gpt-4-system-card.pdf>
- [82] OpenAI. 2023. How should AI systems behave, and who should decide? <https://openai.com/blog/how-should-ai-systems-behave>
- [83] OpenAI. 2023. OpenAI Red Teaming Network. <https://openai.com/blog/red-teaming-network>
- [84] Will Oremus. 2023. Elon Musk promised an anti-'woke' chatbot. It's not going as planned. <https://www.washingtonpost.com/technology/2023/12/23/grok-ai-elon-musk-x-woke-bias/>
- [85] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. *arXiv:2203.02155* [cs.CL]
- [86] Victoria Palacin, Matti Nelimarkka, Pedro Reynolds-Cuèllar, and Christoph Becker. 2020. The design of pseudo-participation. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 2*. 40–44.
- [87] Dev Patnaik and Robert Becker. 1999. Needfinding: the why and how of uncovering people's needs. *Design Management Journal (Former Series)* 10, 2 (1999), 37–43.
- [88] Michael Quinn Patton. 2014. *Qualitative research & evaluation methods: Integrating theory and practice*. Sage publications.
- [89] Billy Perrigo. 2023. Exclusive: OpenAI Used Kenyan Workers on Less Than 2 Dollars Per Hour to Make ChatGPT Less Toxic. <https://time.com/6247678/openai-chatgpt-kenya-workers/>
- [90] Stephen R. Pfohl, Heather Cole-Lewis, Rory Sayres, Darlene Neal, Mercy Asiedu, Awa Dieng, Nenad Tomasev, Qazi Mamunur Rashid, Shekoofeh Azizi, Negar Rostamzadeh, Liam G. McCoy, Leo Anthony Celi, Yun Liu, Mike Schaekermann, Alanna Walton, Alicia Parrish, Chirag Nagpal, Preeti Singh, Akeiyah Dewitt, Philip Mansfield, Sushant Prakash, Katherine Heller, Alan Karthikesalingam, Christopher Semturs, Joelle Barral, Greg Corrado, Yossi Matias, Jamila Smith-Loud, Ivor Horn, and Karan Singhal. 2024. A Toolbox for Surfacing Health Equity Harms and Biases in Large Language Models. *arXiv:2403.12025* [cs.CY]
- [91] Emma Pierson, Divya Shanmugam, Rajiv Movva, Jon Kleinberg, Monica Agrawal, Mark Dredze, Kadija Ferryman, Judy Wawira Gichoya, Dan Jurafsky, Pang Wei Koh, et al. 2023. Use large language models to promote equity. *arXiv preprint arXiv:2312.14804* (2023).
- [92] Organizers Of Queerinaï, Anaelia Ovalle, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J. Sutherland, Davide Locatelli, Eva Breznik, Filip Klubicka, Hang Yuan, Hetvi J. Huan Zhang, Jaidev Shiram, Kruno Lehman, Luca Soldaini, Maarten Sap, Marc Peter Deisenroth, Maria Leonor Pacheco, Maria Ryskina, Martin Mundt, Milind Agarwal, Nyx Mclean, Pan Xu, A. Pranav, Raj Korpan, Ruchira Ray, Sarah Mathew, Sarthak Arora, St John, Tanvi Anand, Vishakha Agrawal, William Agnew, Yanan Long, Zijie J. Wang, Zeerak Talat, Avijit Ghosh, Nathaniel Dennler, Michael Noseworthy, Sharvani Jha, Emi Baylor, Aditya Joshi, Natalia Y. Bilenko, Andrew Mcnamara, Raphael Gontijo-Lopes, Alex Markham, Evyn Dong, Jackie Kay, Manu Saraswat, Nikhil Vyta, and Luke Stark. 2023. Queer In AI: A Case Study in Community-Led Participatory AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAcCT '23). Association for Computing Machinery, New York, NY, USA, 1882–1895. <https://doi.org/10.1145/3593013.3594134>
- [93] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290* (2023).
- [94] Ethan Shaooran, Ido Pesok, and Sam Jones. 2023. Aligned: Platform-based Alignment. *Energize AI* (2023).
- [95] Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndotsse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2023. Towards Understanding Sycophancy in Language Models. *arXiv:2310.13548* [cs.CL]
- [96] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2022. Participation Is Not a Design Fix for Machine Learning. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization* (<conf-loc>, <city>Arlington</city>, <state>VA</state>, <country>USA</country>, </conf-loc>). (EAAMO '22). Association for Computing Machinery, New York, NY, USA, Article 1, 6 pages. <https://doi.org/10.1145/3551624.3555285>
- [97] Zachary Small. 2023. Sarah Silverman Sues OpenAI and Meta Over Copyright Infringement. <https://www.nytimes.com/2023/07/10/arts/sarah-silverman-lawsuit-openai-meta.html>

- [98] Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2022. Learning to summarize from human feedback. *arXiv:2009.01325* [cs.CL]
- [99] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in NLP. *arXiv preprint arXiv:1906.02243* (2019).
- [100] Harini Suresh, Rajiv Movva, Amelia Lee Dogan, Rahul Bhargava, Isadora Cruxên, Ángeles Martínez Cuba, Guilia Taurino, Wonyoung So, and Catherine D'Ignazio. 2022. Towards intersectional feminist and participatory ML: A case study in supporting Femicide Counterdata Collection. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 667–678.
- [101] Jorim Theuns, Evelien Nieuwenburg, Pepijn Verburg, Lei Nelissen, Brett Hennig, Rich Rippin, Ran Haase, Aldo de Moor, CeesJan Mol, Naomi Esther, Rolf Kleef, and Bram Delisse. 2023. Deliberation at Scale: Socially democratic inputs to AI. *Commonground* (2023).
- [102] Emily Tseng, Rosanna Bellini, Yeuk-Yu Lee, Alana Ramjit, Thomas Ristenpart, and Nicola Dell. 2024. Data Stewardship in Clinical Computer Security: Balancing Benefit and Burden in Participatory Systems. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 39 (apr 2024), 29 pages. <https://doi.org/10.1145/3637316>
- [103] Visa. [n.d.]. Visa Account Attack Intelligence. Online. <https://km.visamiddleeast.com/content/dam/VCOM/global/run-your-business/documents/visa-account-attack-intelligence-final.pdf>.
- [104] Junjie Wang, Yuchao Huang, Chunyang Chen, Zhe Liu, Song Wang, and Qing Wang. 2023. Software testing with large language model: Survey, landscape, and vision. *arXiv preprint arXiv:2307.07221* (2023).
- [105] Yang Wang, Yun Huang, Tanusree Sharma, Dawn Song, Sunny Liu, and Jeff Hancock. 2023. Inclusive.AI: Engaging Underserved Populations in Democratic Decision-Making on AI. (2023).
- [106] David Gray Widder, Sarah West, and Meredith Whittaker. 2023. Open (for business): Big tech, concentrated power, and the political economy of open AI. *Concentrated Power, and the Political Economy of Open AI (August 17, 2023)* (2023).
- [107] Christine T Wolf, Haiyi Zhu, Julia Bullard, Min Kyung Lee, and Jed R Brubaker. 2018. The changing contours of “participation” in data-driven, algorithmic ecosystems: Challenges, tactics, and an agenda. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 377–384.
- [108] Chloe Xiang. 2023. Man dies by suicide after talking with AI chatbot, widow says. *Motherboard by Vice News* (2023). <https://www.vice.com/en/article/pkadgm/man-dies-by-suicide-after-talking-with-ai-chatbot-widow-says>
- [109] Kai Yang, Jian Tao, Jiafei Lyu, Chunjiang Ge, Jiabin Chen, Qimai Li, Weihai Shen, Xiaolong Zhu, and Xiu Li. 2023. Using Human Feedback to Fine-tune Diffusion Models without Any Reward Model. *arXiv:2311.13231* [cs.LG]
- [110] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. In *Proceedings of the 2020 chi conference on human factors in computing systems*. 1–13.
- [111] Travis Zack, Eric Lehman, Mirac Suzgun, Jorge A Rodriguez, Leo Anthony Celi, Judy Gichoya, Dan Jurafsky, Peter Szolovits, David W Bates, Raja-Elie E Abdunour, et al. 2024. Assessing the potential of GPT-4 to perpetuate racial and gender biases in health care: a model evaluation study. *The Lancet Digital Health* 6, 1 (2024), e12–e22.
- [112] Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2020. Fine-Tuning Language Models from Human Preferences. *arXiv:1909.08593* [cs.CL]
- [113] Jonathan Zong and J Nathan Matias. 2024. Data Refusal From Below: A Framework for Understanding, Evaluating, and Envisioning Refusal as Design. *ACM Journal on Responsible Computing* 1, 1 (2024), 1–23.