

Identifying Harm in Personalized, Generative AI Systems Requires User-Centered Auditing at the Interaction Level

Hannah Cha

Microsoft Research
Stanford University
v-hannahcha@microsoft.com

Abstract

Personalized, generative AI systems increasingly adapt their behavior to individual users over time, fundamentally changing model behavior. While existing auditing approaches have been effective at surfacing harms in non-personalized contexts, they often rely on static, simulated evaluations and definitions of harm that aggregate across broad, group categories. In this position paper, we argue that such approaches can fail to capture emergent harms in personalized generative AI systems, where harms surface through interpretations of ongoing interaction and evolve with user history. We identify three presuppositions underlying many harm auditing paradigms: that harms can be (1) specified outside real-world interaction, (2) defined non-pluralistically within groups, and (3) treated as static. One might argue that personalized systems could simply learn definitions of what constitutes harm to individual users through repeated interactions. However, we argue that attempts to surface user harms through deeper personalization risk imposing asymmetric burdens of labor and privacy on marginalized users. Consequently, we propose reframing understandings of harm as adaptive, user- and community-centered processes, and outline design directions that shift auditing from retrospective evaluation toward infrastructures that support ongoing articulation of harm in interaction. Our work highlights the need for auditing and design practices that better reflect the pluralistic and evolving nature of harm understanding in personalized generative AI systems.

Introduction

Generative AI systems are being increasingly integrated into everyday applications, such as chatbots (Neumann et al. 2024; Dam et al. 2024) and writing assistants (Mysore et al. 2024; Guo and Li 2024). These systems have increasingly grown to personalize, or adapt outputs to individual users (Eapen and Adhithyan 2023), based on past behavior (Ning et al. 2025; Xu et al. 2025) and user preferences (Kang et al. 2023; Gao et al. 2024). These methods promise more relevant and engaging interactions for users (Kim et al. 2025; Ha et al. 2024). While personalization is not new in computational systems (Lü et al. 2012; Joachims 2002; Narayanan et al. 2004), the nature of personalization in generative AI is fundamentally different. Un-

like selection-based personalization, which chooses from a fixed set of outputs, generative AI reshapes the distribution of possible outputs, producing indeterminate, context-dependent, and evolving behaviors (Ning et al. 2025; Kowsheer et al. 2024). In fact, personalization can fundamentally alter the nature of model behavior itself (Wang, Ho, and Koyejo 2025). As personalization continues to shape the ways that systems interact with users, it is crucial to be able to document exactly when and how harms may emerge.

Existing paradigms to audit harms in generative AI systems often occur as non-personalized, simulated evaluations (Sandvig et al. 2014; Hofmann et al. 2024; Salinas, Haim, and Nyarko 2024), where the model does not retain memory of any prior interactions and the simulated context can be abstracted away from real world use. While these approaches have systematically revealed biases in baseline models (Hofmann et al. 2024; Salinas, Haim, and Nyarko 2024; Mittermaier, Raza, and Kvedar 2023; Lewis et al. 2025), they may fail to capture true user harms arising as a result of personalized interaction. Harm cannot be determined solely by what the output constitutes, but from the way individual users experience and interpret these outputs (Wang et al. 2025a; Brancalone et al. 2025; Mosley and Biernat 2025). As AI systems become personalized to individual users, merely auditing for harms in non-personalized contexts risks missing emergent harms that impact real users. Additionally, current systems that users interact with provide few mechanisms for users to articulate harm, contest system assumptions, or directly shape how personalization unfolds over time. In the absence of such mechanisms, personalized AI systems default to making inferences about users (Purificato, Boratto, and Luca 2024). Ultimately, this risks treating user preferences, such as what constitutes harm, as something that can be inferred from proxies (Wang, Ho, and Koyejo 2025) rather than directly articulated by users.

In this position paper, we specifically focus on *experiential harms*, which consist of harms that users can self-identify and report experiencing negative affect from (Wang et al. 2025a). We argue that personalized generative AI systems create experiential harms during the course of *interaction* that are often difficult to surface through traditional auditing mechanisms because interactions can create emergent behavior. More specifically, we argue that many, general-

purpose auditing approaches presuppose that:

1. **Harms can be specified outside of real-world interaction:** Auditing often happens in a memoryless, non-personalized contexts, abstracted away from how systems may truly harm users in the real world.
2. **Harms can be defined non-pluralistically within groups:** Auditing often defines harm at broad, demographic categories (e.g. race, gender), which can obscure the pluralistic, intersectional ways members of the same group can experience harm.
3. **Harms are static:** Many auditing methods lack continual or evolving evaluation over time, failing to capture how personalization introduces new harmful behavior over time, or how user definitions of harm may evolve.

We further argue that attempting to personalize definitions of harm to individual users with current system paradigms can create asymmetries of *labor* and *privacy*, disproportionately burdening users in minoritized groups. To address these challenges and bridge the gap in understanding true user harms, we propose directions for auditing that acknowledge defining harm as an *adaptive, user- and community-centered process*, emphasizing evolving, pluralistic understandings of harm without causing disproportionate burden on users in minority groups. Ultimately, this work motivates the need for better harm auditing mechanisms that allow all users, especially those from minoritized backgrounds, to reap the benefits of personalized, generative technologies.

Related Work

Personalized, Generative AI Systems

Personalization is not new in computational systems. Recommendation systems (Lü et al. 2012), search engines (Joachims 2002), and adaptive interfaces (Narayanan et al. 2004) have long tailored outputs to users. In these traditional systems, personalization primarily affects selection: which item is recommended (Ricci, Rokach, and Shapira 2010), which result is ranked higher (Joachims 2002), or which predefined label is assigned (Barocas and Selbst 2016). Although users may interpret these outputs differently, the system’s behavior itself often remains stable and enumerable. Thus, algorithmic decision rules can be probed to surface discrimination or harms because outputs are repeatable and enumerable across test inputs (Sandvig et al. 2014).

By contrast, generative AI systems produce outputs through stochastic generation that may be conditioned on user behavior and preferences (Ning et al. 2025; Xu et al. 2025), prior conversational history (Kowsher et al. 2024), or even inferred attributes (Salinas, Haim, and Nyarko 2024). Thus, personalization in generative settings does not merely select among existing outputs, but reshapes the distribution over possible behaviors. The system may adopt different tones, frames, assumptions, and normative stances depending on the user it is interacting with (Wang et al. 2025b). In fact, personalization fundamentally changes model behavior, even modifying behavior on standard benchmarks (Wang, Ho, and Koyejo 2025). Consequently,

the model does not exhibit a single, stable behavior, but rather a family of behaviors that continually evolves in a personalization context that requires new ways of evaluation (Hopkins, Boggust, and Suresh 2025).

Recent work has attempted to better understand and evaluate how personalized interaction can affect behaviors for generative AI systems, such as chatbots (Geng et al. 2025; Shen et al. 2024; Wang et al. 2025b). We extend this need for new evaluation paradigms for personalized, generative AI systems specifically for the harms that users may experience.

Surfacing Harms Via Algorithm Auditing

Algorithm auditing describes the process of surfacing harms by directly analyzing the outputs of such systems (Birhane et al. 2024; Lacmanovic and Skare 2025; Metaxa et al. 2021). Such harms include stereotypes (Ghosh et al. 2024), problematic outputs (Birhane et al. 2024; Lacmanovic and Skare 2025), or surfacing demographic-based disparities in outcomes in domains such as healthcare (Obermeyer et al. 2019; Mittermaier, Raza, and Kvedar 2023; Seyyed-Kalantari et al. 2021), employment (Chen et al. 2018; Wilson and Caliskan 2024; Wen et al. 2025), and housing (Asplund et al. 2020; Zou and Khern-am nuai 2023; Juhn et al. 2022). Many of these audit-based methods, especially in the context of generative AI systems, often consist of researchers simulating a use case of an LLM through prompts and iteratively evaluating a series of AI outputs. Potential harms from these systems are often surfaced by providing proxies for demographic groups, such as name (Salinas, Haim, and Nyarko 2024) or dialect (Hofmann et al. 2024), and measuring discrepancies across groups. Such simulated studies have measured these discrepancies in the natural language used to describe groups (Hofmann et al. 2024; Lim and Pérez-Ortiz 2024; Cheng, Durmus, and Jurafsky 2023), outcomes groups achieve (Salinas, Haim, and Nyarko 2024; Bouguettaya, Stuart, and Aboujaoude 2025), or the way marginalized groups are portrayed in generative imagery (Wang, Morgenstern, and Dickerson 2025; Yang 2025).

However, one common critique of these audit-based methods is that these evaluations are created and run by either AI practitioners (Metaxa et al. 2021; Sandvig et al. 2014; Sweeney 2013) or AI researchers (Buolamwini and Gebru 2018; Cramer et al. 2018; Eslami et al. 2019, 2017; Raji et al. 2020), who may be disconnected from the communities actually experiencing harm (Blodgett et al. 2020). As a result, there has been a rise in end users and communities identifying emergent harms from AI systems (DeVos et al. 2022; Shen et al. 2021). Users can discover nuanced and community-specific harms that have previously been overlooked by researchers or practitioners (Mack et al. 2024; Mim et al. 2024; Zhang et al. 2024; Young, Magassa, and Friedman 2019; Qadri et al. 2025). For instance, Deng et al. (2025)’s WeAudit created an infrastructure allowing for users to audit amongst community members and report harms to AI practitioners in actionable ways. These user-centered auditing practices have been helpful in surfacing community-specific harms, such as stereotypical represen-

tations of non-Western cultures (Ghosh et al. 2024), inappropriate outputs in cultural contexts (Qadri et al. 2025) and the lack of disability (Mack et al. 2024) and gender (Ghosh, Lutz, and Caliskan 2024) representation in generative image models. Overall, these auditing methods focus on experiential harms, or harms that can be surfaced through self-reported experiences of negative affect (Wang et al. 2025a), from the users themselves.

Nonetheless, existing auditing approaches can miss harms emerging from the ways that users interact with personalized AI tools. Simulated, stateless audits happen outside of true user interaction, and many community-based methods may consult communities at the design phase rather than continuing to evaluate the impact of tools after they are deployed (Harrington, Erete, and Piper 2019). Furthermore, even end-user auditing often functions with the goal of providing technical recommendations and remediations for AI practitioners (Deng et al. 2025; Saleiro et al. 2019; Raji et al. 2020; Mökander 2023), rather than preventing harm or reshaping outputs for users as they occur. This can lead to an asymmetry of burden and benefit, where those who perform the work required by the system may not be the one who receive its benefits (Grudin 1988). Ultimately, our work highlights this gap in auditing paradigms by outlining the nature of harm in personalized, generative AI systems that current auditing methods may fail to capture.

Why Existing Auditing Approaches Are Insufficient for Personalized Generative AI Systems

In this section, we argue that personalized, generative AI systems create harms that current auditing approaches may be insufficient for surfacing. We outline how existing auditing approaches often presuppose harm (1) can be specified outside of real world interaction, (2) is non-pluralistic within groups, and (3) is static.

Auditing Approaches Presuppose Harm Can be Specified Outside Real-World Interaction

We argue that existing auditing approaches for defining and quantifying harm presuppose harms are specifiable outside interaction. Many auditing approaches have operationalized potential harms arising from AI systems through simulated studies. These simulated evaluations have quantified harms such as stereotyping (Hofmann et al. 2024; Melero Lázaro, García Ull et al. 2023; Lim and Pérez-Ortiz 2024; Cheng, Durmus, and Jurafsky 2023), unequal treatment (Salinas, Haim, and Nyarko 2024; Bouguettaya, Stuart, and Aboujaoude 2025), or erasure of marginalized groups (Wang, Morgenstern, and Dickerson 2025; Yang 2025). While these strategies have enabled important advances in understanding harms from generative AI systems, they share an assumption that harms can be articulated and evaluated independently of how systems are experienced in regular use.

Personalized generative AI systems do not produce fixed or memoryless outputs. Rather, they generate responses dynamically through interaction, adapting to user behavior and

preferences (Ning et al. 2025; Xu et al. 2025) or prior conversational history (Kowsher et al. 2024). As a result, harms cannot be fully characterized by isolated outputs or one-shot evaluations. Prior work has shown that such evaluations frequently fail to capture how systems truly behave in personalized contexts (Wang, Ho, and Koyejo 2025). Consequently, outputs from simulated audits may not correspond to the harms users actually experience, while significant real-world harms may remain entirely invisible to evaluators. Recent real-world incidents illustrate how severe harms can emerge through prolonged, situated interaction rather than isolated prompts. While extreme, real world cases such as the generation of CSAM (BBCNews 2026) or suicide (Chatterjee 2025) demonstrate consequential failure modes that may not be detectable through standard, static audit paradigms.

Thus, traditional evaluation of generative AI systems, often occurring in simulated, memoryless, or post-hoc audits, are abstracted from real-world use cases of these systems and its subsequent impact on users. Taken together, this suggests that existing auditing paradigms may fail to capture harms in personalized generative AI systems. By defining and evaluating harm outside of real-world interaction, these frameworks risk overlooking harms that emerge as systems adapt to users.

Auditing Approaches Presuppose Non-Pluralistic Definitions of Harm

Existing auditing approaches can further obscure asymmetries in how harm is distributed and experienced. Many fairness evaluations operationalize harm at an aggregate group level, commonly defined by broad demographic categories such as race or gender (Zack et al. 2023; Salinas, Haim, and Nyarko 2024; Thakur 2023; Ranjan, Gupta, and Singh 2024). However, these groupings to evaluate for harm implicitly assume that members of a group experience harms in similar ways. This assumption can mask misaligned, alienating, or harmful interactions experienced by subsets of users within a group. Even participatory approaches, aimed at directly understanding harms from groups themselves, can run the risk of obscuring asymmetries and complexities in how harm is experienced. Critical scholarship has long cautioned against equating participation with representation (Sloane et al. 2022; Harrington, Erete, and Piper 2019). Processes that elevate a small number of what may seem like representative voices, as in participatory design, may obscure internal disagreement, marginalize less powerful members, or reproduce dominant norms that exist within a community (Young 2002; Costanza-Chock 2020; Ovalle et al. 2023). Prior work has highlighted the challenges of pluralistic alignment, where generative AI systems overlook human social diversity between groups (Sorensen et al. 2024; Ali et al. 2025); we extend this idea to within-group diversity and highlight the pitfalls of documenting harms and preferences at the group level. Importantly, we do not focus on non-negotiable harms, such as the generation of slurs, explicit stereotypes, or overtly abusive content, for which there is broad, normative agreement of what constitutes harm. Instead, we examine outputs that occupy an ambiguous middle ground, where framing, assumptions, or nor-



Figure 1: We present an example of case in which different plausible interpretation of a chat output by two users in the same demographic group can cause harm to one and not the other. We also consider how intersectionality can affect the harm a user might experience, where additional identities (e.g. income status) can add new harm dimensions for a user in the same demographic group.

mative stance can be interpreted in fundamentally different ways depending on users' lived experiences.

Figure 1 illustrates how an example of an output from a personalized, generative AI system, while not overtly stereotypical or detectable as harmful in traditional auditing paradigms, can be interpreted in fundamentally different ways by users sharing a demographic identity. Two users within the same demographic group may interpret the response in two contrasting ways, each being rooted in their own lived experience. Psychological studies have documented this within-group pluralism of what constitutes as harm or bias (Brancalone et al. 2025), such as in the context of cultural appropriation (Mosley and Biernat 2025). Furthermore, as illustrated in the figure, additional identities (e.g. low-income status) can further shape how harm is perceived, experienced, or intensified. Thus, intersectionality adds another dimension of consideration for the multiplicity of interpretation of what can be determined as harmful: different forms of inequality manifest from various axes of identity (e.g. race, gender, class), creating unique experiences of privilege or discrimination (Crenshaw 2013a). Thus, trying to understand or audit for harms at broad group categories risk compounding harm by obscuring how multiple dimensions of identity and experience shape interpretation and definitions of harm (Costanza-Chock 2020; Collins et al. 2021; Crenshaw 2013b; Wang et al. 2025a; Ovalle et al. 2023). Prior work has shown that, due to the combinatorial explosion of subgroups, measuring intersectional fairness exhaustively can be intractable (Molina and Loiseau 2022). Similarly, enumerating all plausible interpretations of harm for a given system output may be impossible; even the type of harm (e.g. racism, sexism, ageism) should neither be measured or mitigated in the same way (Wang 2025).

The pluralistic ways users within a group can experience harm are a problem for current personalization mechanisms in generative AI systems, which often rely on identity-based

proxies in attempts to personalize to users. For instance, Wang et al. (2025b) finds that GPT was more likely to recommend *Black Panther* to names associated with Black individuals and *Little Women* to names associated with women in personalized contexts, flattening the diversity of preferences of a group into stereotypes. Even if personalized systems were able to avoid explicit stereotypes, the fact that personalization can depend on identity-based proxies showcases how these systems can inadvertently elevate the preferences of a subset of users into implicit group standards, erasing the pluralistic ways of how harm can occur. Allowing for dominant members of groups to determine what constitutes harm can lead to its own harms. Psychological studies have documented cases of lateral violence, where members of oppressed groups can marginalize members within the same group (Whyman et al. 2021, 2022). Similarly, qualitative studies have documented intragroup discrimination, where intersectional identities can result in harm within the same marginalized group (MacCarthy et al. 2021). Consequently, personalized generative systems operate under conditions of pluralism in harm understanding, where users who share demographic identities may nonetheless have legitimate, conflicting expectations of what constitutes harmful system behavior.

Auditing Approaches Presuppose Harm as Static

Furthermore, we argue that many existing auditing approaches fail to account for how harms can evolve over the course of interaction. Interpretations of what constitutes harm may not be static even for a single individual. Psychology studies have found, for instance, that situational and dispositional factors can prime individuals to consider previously ambiguous statements as harmful (Bleske-Rechek et al. 2023) or as a microaggression (Holtgraves et al. 2025). Even an individual's level of historical knowledge affects their ability to perceive harms, such as racism (Bonam et al.

2019). Thus, it may be plausible, for example, that a user who initially appreciates culturally specific food recommendations may later come to experience them as reductive as they learn to recognize microaggressions, such as how food recommendations can become a mechanism for stereotyping. Thus, system behavior may shift from being perceived as helpful to harmful as users continue to evolve.

Traditional end-user auditing mechanisms rarely account for the interpretive fluidity of harm, treating user preferences and definitions of harm as stable rather than contextual and longitudinal. Empirical audits of large language models typically rely on one-off evaluations, such as static benchmark datasets (Blodgett et al. 2020), prompt sets (Salinas, Haim, and Nyarko 2024), or simulated personas (Santurkar et al. 2023), to assess harmful behaviors. Thus, such audits may fail to capture how harms vary across time, interaction history, and changing user understanding. Even participatory approaches can fall into the trap of treating community-driven insights as static artifacts, often consulted at design time and translated into system requirements without long term engagement (Harrington, Erete, and Piper 2019; Udoewa 2022). Unless evaluation and community connection continue beyond the creation of the system itself, harms and insights from the community when they put the system to use can be missed (Harrington, Erete, and Piper 2019; Nguyen 2022; Frauenberger et al. 2015).

Thus, in the context of personalized, generative AI, many auditing approaches and even participatory methods may fail to consider how harm might evolve over time as users themselves change or build history with personalization. Ultimately, without longitudinal audits or regular consultation of users or communities, it becomes difficult to account for how harms evolve as systems adapt, users change, and contexts shift.

Perfect Personalization Doesn't Close the Auditing Gap

A natural response to the limitations outlined above might be that if current personalization mechanisms don't adequately capture individual harm, then the solution could be deeper, more accurate personalization. One might argue that personalized, generative AI systems could eventually learn what constitutes harm for a given user by personalizing based on inferred user discomfort or past negative signals. In theory, if users interact with personalized AI systems long enough, the system may be able to perfectly capture fluid preferences and avoid harmful output. However, we argue that attempting to resolve this issue with further personalization with the current paradigm of personalized generative AI systems leads to an asymmetry of **labor** and **privacy** for users in the minority.

Asymmetry of Labor

Creating a perfectly personalized, generative AI system for a given user begs the question of what labor it would take to create such an experience. Without explicit mechanisms for users to articulate harm or transparency into whether such articulation shapes future behavior, users may be repeatedly

exposed to harmful interactions before a system adapts to them. Prior work has already found that LLMs often default to outputs that cater towards WEIRD (Western, Educated, Industrialized, Rich, Democratic) perspectives and values (Agarwal, Naaman, and Vashistha 2025; Santurkar et al. 2023). Thus, users who fall outside of that default would be more likely to encounter harmful outputs that need to be corrected. This dynamic distributes the labor of correction unevenly: users whose interpretations align with dominant norms may experience less friction for creating personalized experiences, while those whose experiences diverge must repeatedly intervene to correct system assumptions.

Additionally, even if users were to articulate harms as they occur, the lack of transparency in current AI systems (Worth et al. 2024) and the difficulty for users to steer AI system behavior (Vafa et al. 2025) make it unclear whether this labor would prevent harms in future outputs. Often, the goal of AI auditing is technical remediation by AI practitioners or model developers (Saleiro et al. 2019; Raji et al. 2020; Mökander 2023). However, this consolidates power in the hands of platforms, who can decide whether to address the harms users experience and how they should do so. Thus, users who document harms they experience directly to the AI system or through traditional auditing mechanisms are not guaranteed to experience less harm as a result of their labor.

Ultimately, repeated harmful outputs from the system can lead to reduced trust and adoption among marginalized groups, who have already been documented as less likely to be AI adopters (Zhou et al. 2025). This can ultimately exacerbate digital divides (Van Dijk and Hacker 2003) in who benefits from personalized AI systems (Zhou et al. 2026), concentrating the advantages of adaptive systems among users whose interpretations align more closely with dominant norms. Thus, minoritized users may encounter more harms and bear disproportionate labor in making their preferences legible to the system, creating a labor asymmetry.

Asymmetry of Privacy

In addition to imposing unequal labor burdens, attempts to achieve perfect personalization in generative AI systems may also create an asymmetry of privacy. Personalized, generative AI systems often rely on the collection and retention of user data to adapt system behavior (Luera et al. 2025; Li et al. 2025). However, the amount and sensitivity of information required to avoid harms is not evenly distributed across users. For users whose preferences and interpretations of harm already align with dominant norms embedded in system defaults, effective personalization may require little additional disclosure. Their needs may be more likely to be met by generic system behavior, allowing them to benefit from the system without explicitly revealing sensitive aspects of their identity, experiences, or values. In contrast, users whose preferences diverge from these defaults may have to disclose more information to correct system assumptions. To steer system behavior away from alienating or harmful outputs, they may have to reveal information that is more personal, sensitive, or identity-linked than what is required of users of dominant groups. Even without explicitly

revealing aspects of their identity, personalized systems may infer their identity based on their discomfort or disagreement with system defaults.

Many marginalized communities, especially Indigenous communities, are already concerned about how data collection and disclosures can become a form of extraction that replicates colonial and extractive histories (Moudgalya and Swaminathan 2024; Spano and Zhang 2025). This concern, coupled with opaque data practices from proprietary AI systems (Vaassen 2022), can cause marginalized users, who may already have distrust for these technological systems (Mhasakar et al. 2025), to avoid these systems entirely, furthering the digital divide of who uses emergent technologies (Van Dijk and Hacker 2003).

Ultimately, users from marginalized groups may be placed in a position where avoiding harm may require increased self-disclosure. This creates a privacy asymmetry, where users may have to choose between tolerating repeated harm or surrendering additional personal data to achieve desired interactions. Even if users reveal information about themselves to these systems, they risk being stereotyped, rather than the system truly aligning with their preferences. Prior work has described this phenomenon as a “personalization double bind” for marginalized users, where refusing personalization may expose them to majority-oriented defaults, while accepting personalization risks stereotyping or overfitting to group identity (Wang et al. 2025b).

Towards Understanding Harm As User-Centered, Adaptive Processes

Taken together, we argue that harm in personalized generative AI systems cannot be adequately understood as a fixed property of model outputs, nor as a quantity that can be fully specified through existing auditing paradigms. Instead, harms in these systems emerge through ongoing interaction, shaped by personalization, user history, shifting social context, and evolving harm interpretations. As such, we argue that understanding and addressing harm in personalized generative AI systems requires reconceptualizing understanding harm as a user-centered, adaptive process. Personalization can amplify the interpretative nature of harm by exposing users to system behaviors that are indeterminate, context-dependent, and differentially interpreted. As a result, users should not merely be recipients of system behavior, but active participants in defining what constitutes harm for them and how this may shift over time.

This reconceptualization has implications for how we understand the goals of auditing. Existing auditing paradigms primarily aim to surface harms for model developers, often in service of downstream technical remediation (Saleiro et al. 2019; Raji et al. 2020; Mökander 2023). Users who experience harm rarely receive immediate feedback or assurance that their experiences will shape future system behavior. In personalized systems, where harms arise through interaction itself, auditing that remains external to user experience risks missing precisely the harms that are most impactful. We argue that auditing should shift toward supporting users directly in shaping less harmful interactions, rather

than solely producing retrospective assessments for developers. We do not claim this goal is absent from existing work as participatory and community-based auditing methods (e.g., Deng et al. (2025)) and longitudinal HCI scholarship (Harrington, Erete, and Piper (2019)) already treat harm as socially situated and contextual. Yet, even these approaches typically consult users or communities at discrete points, such as design time or post-hoc reporting, rather than continuously during interaction itself, and typically resolve community input into a single articulation of harm for system design rather than preserving ongoing plurality and revision. It is this gap, between situated auditing and continuously adaptive, plural auditing suited to the indeterminate and evolving behavior of personalized generative systems, that motivates our proposed directions below.

Auditing Infrastructure

We advocate for frameworks that move auditing practices for personalized generative AI systems away from silently inferring what constitutes harm and toward infrastructures that support explicit articulation, negotiation, and revision of harm over time. Rather than designers, auditors, or models unilaterally determining what is harmful for users, auditing structures should create space for users to express when outputs feel harmful, reductive, or misaligned, and to do so without assuming that harm is stable or universally agreed upon, even within shared demographic groups.

One possible direction is to design auditing structures that allow users to indicate moments of harm during interaction and, if they choose, contextualize why a response was problematic. These signals would not function as ground truth labels, but as expressions of user interpretation that can directly shape future interactions. Importantly, this reframes personalization not as the system “learning the user,” but as an ongoing process in which users retain agency over how system behavior adapts.

However, relying solely on individual-level articulation risks reproducing asymmetries of labor, disproportionately burdening users, particularly those from marginalized groups, with the work of identifying and explaining harm, as aforementioned. To mitigate this, individual feedback could be complemented by community-mediated signals, where users can encounter harm characterizations articulated by others with shared or adjacent experiences and choose whether they resonate. For example, users might opt into, reject, or modify shared descriptions of harmful behaviors, without being required to generate every explanation completely from scratch.

Crucially, community signals should not be treated as authoritative or homogenizing. Instead, they should make visible disagreement, variation, and internal plurality within groups, resisting the flattening of identities into normative standards. Users should retain the ability to refuse personalization altogether, selectively adopt shared interpretations, or revise prior signals as their understanding evolves. Across both individual and community layers, participation must remain voluntary, legible, and non-punitive. Safety should not be contingent on user labor, nor should avoiding harm require excessive self-disclosure.

More broadly, these directions suggest a shift in the purpose of auditing itself: from identifying harms for external remediation toward creating infrastructure through which harm can be expressed and addressed situationally. While these approaches raise open questions, such as how to balance conflicting signals, prevent misuse, and avoid reinforcing dominant norms, they underscore the central idea that understanding harm in personalized generative AI requires mechanisms for more user control and harm articulation, not just improved metrics or more granular personalization.

Limitations and Open Challenges

These directions for future auditing raise several limitations and open challenges. Surfacing harm at the interaction level requires trust. Users, particularly those from marginalized communities, may hesitate to disclose harm in systems operated by large institutions with opaque data practices or histories of extraction (Mhasakar et al. 2025). Harm-reporting mechanisms themselves can be misused, ignored, or weaponized, especially when power remains centralized with platform providers. Surfacing end-user harms would be most ideal in systems that are governed by communities themselves. Acknowledging this, we make recommendations that fit into the current structure of personalized systems, but advocate for longer term, structural change that consolidates power towards end-users, not just model developers. We therefore view these suggestions as incremental steps, rather than a substitute for necessary structural change in the governance and development of these systems.

Furthermore, treating harm as adaptive complicates evaluation: pluralism, disagreement, and change over time resist clean measurement and raise difficult questions about accountability and governance. Technical, pluralistic alignment across different demographic groups remains an open challenge in AI systems (Sorensen et al. 2024; Ali et al. 2025), and pluralistic alignment within groups is therefore a challenge as well. Our design recommendations do not address these technical challenges, and rather focus on concrete ways for users to identify and articulate harms within the constraints of personalized AI platforms. Ultimately, understanding and preventing harms in personalized generative AI systems cannot be solved through design or auditing alone. It requires sustained governance, institutional reflexivity, and meaningful power-sharing with the communities most affected by these systems.

Conclusion

As generative AI systems become increasingly personalized and embedded in everyday interaction, prevailing auditing approaches, static, simulated, and group-aggregated, are ill-suited to capture the harms users actually experience. We argue that many consequential harms in personalized generative AI systems arise at the level of interaction, shaped by evolving user history, interpretation, and preferences, and therefore cannot be fully specified or evaluated outside real-world use. Moreover, attempting to close this gap through deeper personalization risks imposing asymmetric burdens of labor and privacy on marginalized users. Taken together,

these challenges motivate a shift from treating harm as a static property of model outputs to understanding harm as an adaptive, user- and community-centered process. Recognizing harm as something that emerges, is contested, and evolves through interaction is essential for developing auditing and design practices that more faithfully reflect lived user experience in personalized AI systems.

Acknowledgments

I am grateful to Emily Tseng and Solon Barocas for their invaluable feedback and guidance.

References

- Agarwal, D.; Naaman, M.; and Vashistha, A. 2025. AI suggestions homogenize writing toward western styles and diminish cultural nuances. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–21.
- Ali, D.; Zhao, D.; Koenecke, A.; and Papakyriakopoulos, O. 2025. Operationalizing Pluralistic Values in Large Language Model Alignment Reveals Trade-offs in Safety, Inclusivity, and Model Behavior. *arXiv preprint arXiv:2511.14476*.
- Asplund, J.; Eslami, M.; Sundaram, H.; Sandvig, C.; and Karahalios, K. 2020. Auditing Race and Gender Discrimination in Online Housing Markets. *Proceedings of the International AAAI Conference on Web and Social Media*, 14(1): 24–35.
- Barocas, S.; and Selbst, A. D. 2016. Big data’s disparate impact. *Calif. L. Rev.*, 104: 671.
- BBCNews. 2026. Changes to Grok’s AI safeguards and concerns over generated sexualised imagery.
- Birhane, A.; Steed, R.; Ojewale, V.; Vecchione, B.; and Raji, I. D. 2024. AI auditing: The Broken Bus on the Road to AI Accountability. *arXiv:2401.14462*.
- Bleske-Rechek, A.; Deaner, R. O.; Paulich, K. N.; Axelrod, M.; Badenhorst, S.; Nguyen, K.; Seyoum, E.; and Lay, P. S. 2023. In the eye of the beholder: Situational and dispositional predictors of perceiving harm in others’ words. *Personality and Individual Differences*, 200: 111902.
- Blodgett, S. L.; Barocas, S.; Daumé Iii, H.; and Wallach, H. 2020. Language (technology) is power: A critical survey of “bias” in nlp. *arXiv preprint arXiv:2005.14050*.
- Bonam, C. M.; Nair Das, V.; Coleman, B. R.; and Salter, P. 2019. Ignoring history, denying racism: Mounting evidence for the Marley hypothesis and epistemologies of ignorance. *Social Psychological and Personality Science*, 10(2): 257–265.
- Bouguettaya, A.; Stuart, E. M.; and Aboujaoude, E. 2025. Racial bias in AI-mediated psychiatric diagnosis and treatment: a qualitative comparison of four large language models. *npj Digital Medicine*, 8(1): 332.
- Brancaleone, P. J.; Cofresí, R. U.; Volpert-Esmond, H. I.; Amodio, D. M.; Ito, T. A.; and Bartholow, B. D. 2025. Within-person dynamics of attention to race and expression of race bias: a real-time test of the self-regulation of prejudice model. *Social cognitive and affective neuroscience*, 20(1): nsaf026.

- Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, 77–91. PMLR.
- Chatterjee, R. 2025. Their Teenage Sons Died by Suicide. Now, They Are Sounding an Alarm About AI Chatbots.
- Chen, L.; Ma, R.; Hannák, A.; and Wilson, C. 2018. Investigating the impact of gender on rank in resume search engines. In *Proceedings of the 2018 chi conference on human factors in computing systems*, 1–14.
- Cheng, M.; Durmus, E.; and Jurafsky, D. 2023. Marked personas: Using natural language prompts to measure stereotypes in language models. *arXiv preprint arXiv:2305.18189*.
- Collins, P. H.; da Silva, E. C. G.; Ergun, E.; Furseth, I.; Bond, K. D.; and Martínez-Palacios, J. 2021. Intersectionality as critical social theory: Intersectionality as critical social theory, Patricia Hill Collins, Duke University Press, 2019.
- Contemporary Political Theory, 20(3): 690.
- Costanza-Chock, S. 2020. *Design justice: Community-led practices to build the worlds we need*. The MIT Press.
- Cramer, H.; Garcia-Gathright, J.; Springer, A.; and Reddy, S. 2018. Assessing and addressing algorithmic bias in practice. *Interactions*, 25(6): 58–63.
- Crenshaw, K. 2013a. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. In *Feminist legal theories*, 23–51. Routledge.
- Crenshaw, K. W. 2013b. Mapping the margins: Intersectionality, identity politics, and violence against women of color. In *The public nature of private violence*, 93–118. Routledge.
- Dam, S. K.; Hong, C. S.; Qiao, Y.; and Zhang, C. 2024. A complete survey on llm-based ai chatbots. *arXiv preprint arXiv:2406.16937*.
- Deng, W. H.; Claire, W.; Han, H. Z.; Hong, J. I.; Holstein, K.; and Eslami, M. 2025. WeAudit: Scaffolding User Auditors and AI Practitioners in Auditing Generative AI. *arXiv:2501.01397*.
- DeVos, A.; Dhabalia, A.; Shen, H.; Holstein, K.; and Eslami, M. 2022. Toward User-Driven Algorithm Auditing: Investigating users’ strategies for uncovering harmful algorithmic behavior. In *Proceedings of the 2022 CHI conference on human factors in computing systems*, 1–19.
- Eapen, J.; and Adhithyan, V. 2023. Personalization and customization of llm responses. *International Journal of Research Publication and Reviews*, 4(12): 2617–2627.
- Eslami, M.; Vaccaro, K.; Karahalios, K.; and Hamilton, K. 2017. “Be careful; things can be worse than they appear”: Understanding biased algorithms and users’ behavior around them in rating platforms. In *Proceedings of the international AAAI conference on web and social media*, volume 11, 62–71.
- Eslami, M.; Vaccaro, K.; Lee, M. K.; Elazari Bar On, A.; Gilbert, E.; and Karahalios, K. 2019. User attitudes towards algorithmic opacity and transparency in online reviewing platforms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–14.
- Frauenberger, C.; Good, J.; Fitzpatrick, G.; and Iversen, O. S. 2015. In pursuit of rigour and accountability in participatory design. *International journal of human-computer studies*, 74: 93–106.
- Gao, G.; Taymanov, A.; Salinas, E.; Mineiro, P.; and Misra, D. 2024. Aligning llm agents by learning latent preference from user edits. *Advances in Neural Information Processing Systems*, 37: 136873–136896.
- Geng, J.; Chen, H.; Liu, R.; Ribeiro, M. H.; Willer, R.; Neubig, G.; and Griffiths, T. L. 2025. Accumulating Context Changes the Beliefs of Language Models. *arXiv preprint arXiv:2511.01805*.
- Ghosh, S.; Lutz, N.; and Caliskan, A. 2024. “I Don’t See Myself Represented Here at All”: User Experiences of Stable Diffusion Outputs Containing Representational Harms across Gender Identities and Nationalities. In *Proceedings of the AAAI/ACM conference on AI, ethics, and society*, volume 7, 463–475.
- Ghosh, S.; Venkit, P. N.; Gautam, S.; Wilson, S.; and Caliskan, A. 2024. Do generative AI models output harm while representing non-Western cultures: Evidence from a community-centered approach. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 476–489.
- Grudin, J. 1988. Why CSCW applications fail: problems in the design and evaluation of organizational interfaces. In *Proceedings of the 1988 ACM conference on Computer-supported cooperative work*, 85–93.
- Guo, K.; and Li, D. 2024. Understanding EFL students’ use of self-made AI chatbots as personalized writing assistance tools: A mixed methods study. *System*, 124: 103362.
- Ha, J.; Jeon, H.; Han, D.; Seo, J.; and Oh, C. 2024. CloChat: Understanding how people customize, interact, and experience personas in large language models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–24.
- Harrington, C.; Erete, S.; and 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.
- Hofmann, V.; Kalluri, P. R.; Jurafsky, D.; and King, S. 2024. AI generates covertly racist decisions about people based on their dialect. *Nature*, 633(8028): 147–154.
- Holtgraves, T.; Sarin, R.; Wood, R.; Cronk, E.; and Nogueira Mourão, A. J. 2025. Interpreting Microaggressions: The Role of Discourse Context, Recipient Status, and Observers’ Political Orientation. *Personality and Social Psychology Bulletin*, 01461672251377809.
- Hopkins, A.; Boggust, A.; and Suresh, H. 2025. Chatbot Evaluation Is (Sometimes) Ill-Posed: Contextualization Errors in the Human-Interface-Model Pipeline.
- Joachims, T. 2002. Optimizing search engines using click-through data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 133–142.

- Juhn, Y. J.; Ryu, E.; Wi, C.-I.; King, K. S.; Malik, M.; Romero-Brufau, S.; Weng, C.; Sohn, S.; Sharp, R. R.; and Halamka, J. D. 2022. Assessing socioeconomic bias in machine learning algorithms in health care: a case study of the HOUSES index. *Journal of the American Medical Informatics Association*, 29(7): 1142–1151.
- Kang, W.-C.; Ni, J.; Mehta, N.; Sathiamoorthy, M.; Hong, L.; Chi, E.; and Cheng, D. Z. 2023. Do llms understand user preferences? evaluating llms on user rating prediction. *arXiv preprint arXiv:2305.06474*.
- Kim, T. S.; Lee, Y.; Park, Y.; Kim, J.; Kim, Y.-H.; and Kim, J. 2025. CUPID: Evaluating Personalized and Contextualized Alignment of LLMs from Interactions. *arXiv preprint arXiv:2508.01674*.
- Kowsher, M.; Panditi, R.; Prottasha, N. J.; Bhat, P.; Bairagi, A. K.; and Arefin, M. S. 2024. Token trails: Navigating contextual depths in conversational ai with chatllm. In *International Conference on Applications of Natural Language to Information Systems*, 56–67. Springer.
- Lacmanovic, S.; and Skare, M. 2025. Artificial intelligence bias auditing—current approaches, challenges and lessons from practice. *Review of Accounting and Finance*, (ahead-of-print).
- Lewis, A. A.; et al. 2025. Unpacking Cultural Bias in AI Language Learning Tools: An Analysis of Impacts and Strategies for Inclusion in Diverse Educational Settings. *International Journal of Research and Innovation in Social Science*, 9(1): 1878–1892.
- Li, H.; Yang, C.; Zhang, A.; Deng, Y.; Wang, X.; and Chua, T.-S. 2025. Hello again! llm-powered personalized agent for long-term dialogue. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 5259–5276.
- Lim, S.; and Pérez-Ortiz, M. 2024. The african woman is rhythmic and soulful: An investigation of implicit biases in llm open-ended text generation. *arXiv preprint arXiv:2407.01270*.
- Lü, L.; Medo, M.; Yeung, C. H.; Zhang, Y.-C.; Zhang, Z.-K.; and Zhou, T. 2012. Recommender systems. *Physics reports*, 519(1): 1–49.
- Luera, R.; Rossi, R.; Dernoncourt, F.; Siu, A.; Kim, S.; Yu, T.; Zhang, R.; Chen, X.; Lipka, N.; Zhang, Z.; et al. 2025. Personalizing Data Delivery: Investigating User Characteristics and Enhancing LLM Predictions. In *Companion Proceedings of the ACM on Web Conference 2025*, 1167–1171.
- MacCarthy, S.; Bogart, L. M.; Galvan, F. H.; and Pantalone, D. W. 2021. Inter-group and intraminority-group discrimination experiences and the coping responses of Latino sexual minority men living with HIV. *Annals of LGBTQ public and population health*, 2(1): 1.
- Mack, K. A.; Qadri, R.; Denton, R.; Kane, S. K.; and Bennett, C. L. 2024. “They only care to show us the wheelchair”: disability representation in text-to-image AI models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–23.
- Melero Lázaro, M.; García Ull, F. J.; et al. 2023. Gender stereotypes in AI-generated images. *El Profesional de la información*, 32(5).
- Metaxa, D.; Park, J. S.; Robertson, R. E.; Karahalios, K.; Wilson, C.; Hancock, J.; and Sandvig, C. 2021. Auditing Algorithms: Understanding Algorithmic Systems from the Outside In. *Foundations and Trends® in Human–Computer Interaction*, 14(4): 272–344.
- Mhasakar, M.; Baker-Ramos, R.; Carter, B.; Helekahi-Kaiwi, E.-B.; and Hester, J. 2025. “I Would Never Trust Anything Western”: Kumu (Educator) Perspectives on Use of LLMs for Culturally Revitalizing CS Education in Hawaiian Schools. *arXiv:2501.17942*.
- Mim, N. J.; Nandi, D.; Khan, S. S.; Dey, A.; and Ahmed, S. I. 2024. In-between visuals and visible: The impacts of text-to-image generative ai tools on digital image-making practices in the global south. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Mittermaier, M.; Raza, M. M.; and Kvedar, J. C. 2023. Bias in AI-based models for medical applications: challenges and mitigation strategies. *NPJ Digital Medicine*, 6(1): 113.
- Mökander, J. 2023. Auditing of AI: Legal, ethical and technical approaches. *Digital Society*, 2(3): 49.
- Molina, M.; and Loiseau, P. 2022. Bounding and approximating intersectional fairness through marginal fairness. *Advances in Neural Information Processing Systems*, 35: 16796–16807.
- Mosley, A. J.; and Biernat, M. 2025. Social identity and the psychology of cultural appropriation. *Handbook of Social Identity Research*, 326–344.
- Moudgalya, S. K.; and Swaminathan, S. 2024. Toward Data Sovereignty: Justice-oriented and Community-based AI Education. In *Proceedings of the 2024 on RESPECT Annual Conference*, RESPECT 2024, 94–99. New York, NY, USA: Association for Computing Machinery. ISBN 9798400706264.
- Mysore, S.; Lu, Z.; Wan, M.; Yang, L.; Sarrafzadeh, B.; Menezes, S.; Baghaee, T.; Gonzalez, E. B.; Neville, J.; and Safavi, T. 2024. Pearl: Personalizing large language model writing assistants with generation-calibrated retrievers. In *Proceedings of the 1st Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U)*, 198–219.
- Narayanan, S.; Koppaka, L.; Edala, N.; Loritz, D.; and Daley, R. 2004. Adaptive interface for personalizing information seeking. *CyberPsychology & Behavior*, 7(6): 683–688.
- Neumann, A. T.; Yin, Y.; Sowe, S.; Decker, S.; and Jarke, M. 2024. An llm-driven chatbot in higher education for databases and information systems. *IEEE Transactions on Education*.
- Nguyen, Q. 2022. Evaluation in participatory design—The whys and the nots. In *Proceedings of the Participatory Design Conference 2022-Volume 2*, 161–166.

- Ning, L.; Liu, L.; Wu, J.; Wu, N.; Berlowitz, D.; Prakash, S.; Green, B.; O'Banion, S.; and Xie, J. 2025. User-llm: Efficient llm contextualization with user embeddings. In *Companion Proceedings of the ACM on Web Conference 2025*, 1219–1223.
- Obermeyer, Z.; Powers, B.; Vogeli, C.; and Mullainathan, S. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464): 447–453.
- Ovalle, A.; Subramonian, A.; Gautam, V.; Gee, G.; and Chang, K.-W. 2023. Factoring the matrix of domination: A critical review and reimagination of intersectionality in ai fairness. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 496–511.
- Purificato, E.; Boratto, L.; and Luca, E. W. D. 2024. User Modeling and User Profiling: A Comprehensive Survey. *arXiv:2402.09660*.
- Qadri, R.; Diaz, M.; Wang, D.; and Madaio, M. 2025. The Case for "Thick Evaluations" of Cultural Representation in AI. *arXiv:2503.19075*.
- Raji, I. D.; Smart, A.; White, R. N.; Mitchell, M.; Gebru, T.; Hutchinson, B.; Smith-Loud, J.; Theron, D.; and Barnes, P. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 33–44.
- Ranjan, R.; Gupta, S.; and Singh, S. N. 2024. A comprehensive survey of bias in llms: Current landscape and future directions. *arXiv preprint arXiv:2409.16430*.
- Ricci, F.; Rokach, L.; and Shapira, B. 2010. Introduction to recommender systems handbook. In *Recommender systems handbook*, 1–35. Springer.
- Saleiro, P.; Kuester, B.; Hinkson, L.; London, J.; Stevens, A.; Anisfeld, A.; Rodolfa, K. T.; and Ghani, R. 2019. Aequitas: A Bias and Fairness Audit Toolkit. *arXiv:1811.05577*.
- Salinas, A.; Haim, A.; and Nyarko, J. 2024. What's in a name? Auditing large language models for race and gender bias. *arXiv preprint arXiv:2402.14875*.
- Sandvig, C.; Hamilton, K.; Karahalios, K.; and Langbort, C. 2014. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry*, 22(2014): 4349–4357.
- Santurkar, S.; Durmus, E.; Ladhak, F.; Lee, C.; Liang, P.; and Hashimoto, T. 2023. Whose opinions do language models reflect? In *International Conference on Machine Learning*, 29971–30004. PMLR.
- Seyyed-Kalantari, L.; Zhang, H.; McDermott, M. B.; Chen, I. Y.; and Ghassemi, M. 2021. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature medicine*, 27(12): 2176–2182.
- Shen, H.; DeVos, A.; Eslami, M.; and Holstein, K. 2021. Everyday algorithm auditing: Understanding the power of everyday users in surfacing harmful algorithmic behaviors. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2): 1–29.
- Shen, H.; Kneare, T.; Ghosh, R.; Alkiek, K.; Krishna, K.; Liu, Y.; Petridis, S.; Peng, Y.-H.; Qiwei, L.; Si, C.; et al. 2024. Position: Towards Bidirectional Human-AI Alignment. In *The Thirty-Ninth Annual Conference on Neural Information Processing Systems Position Paper Track*.
- Sloane, M.; Moss, E.; Awomolo, O.; and Forlano, L. 2022. Participation Is not a Design Fix for Machine Learning. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394772.
- Sorensen, T.; Moore, J.; Fisher, J.; Gordon, M.; Miresghalah, N.; Rytting, C. M.; Ye, A.; Jiang, L.; Lu, X.; Dziri, N.; et al. 2024. A roadmap to pluralistic alignment. *arXiv preprint arXiv:2402.05070*.
- Spano, I.; and Zhang, Y. 2025. Indigenous data sovereignty in intangible cultural heritage governance: A complementary approach to public-private partnerships. *International Journal of Cultural Property*, 1–27.
- Sweeney, L. 2013. Discrimination in online ad delivery. *Communications of the ACM*, 56(5): 44–54.
- Thakur, V. 2023. Unveiling gender bias in terms of profession across LLMs: Analyzing and addressing sociological implications. *arXiv preprint arXiv:2307.09162*.
- Udoewa, V. 2022. An introduction to radical participatory design: decolonising participatory design processes. *Design Science*, 8: e31.
- Vaassen, B. 2022. AI, opacity, and personal autonomy. *Philosophy & Technology*, 35(4): 88.
- Vafa, K.; Bentley, S.; Kleinberg, J.; and Mullainathan, S. 2025. What's Producible May Not Be Reachable: Measuring the Steerability of Generative Models. *ArXiv:2503.17482 [cs]*.
- Van Dijk, J.; and Hacker, K. 2003. The digital divide as a complex and dynamic phenomenon. *The information society*, 19(4): 315–326.
- Wang, A. 2025. Identities are not Interchangeable: The Problem of Overgeneralization in Fair Machine Learning. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 485–497.
- Wang, A.; Bai, X.; Barocas, S.; and Blodgett, S. L. 2025a. Measuring Machine Learning Harms from Stereotypes Requires Understanding Who Is Harmed by Which Errors in What Ways. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 746–762.
- Wang, A.; Beeghly, E.; Koyejo, S.; and Ho, D. E. 2025b. Personalization Double Binds: When User Preferences Meet Group-Based Chatbot Behaviors. *arXiv preprint*.
- Wang, A.; Ho, D. E.; and Koyejo, S. 2025. The inadequacy of offline large language model evaluations: A need to account for personalization in model behavior. *Patterns*, 6(12): 101397.
- Wang, A.; Morgenstern, J.; and Dickerson, J. P. 2025. Large language models that replace human participants can harmfully misportray and flatten identity groups. *Nature Machine Intelligence*, 1–12.

- Wen, A.; Patil, T.; Saxena, A.; Fu, Y.; O'Brien, S.; and Zhu, K. 2025. FAIRE: Assessing Racial and Gender Bias in AI-Driven Resume Evaluations. *arXiv preprint arXiv:2504.01420*.
- Whyman, T.; Adams, K.; Carter, A.; and Jobson, L. 2021. Lateral violence in indigenous peoples. *Australian Psychologist*, 56(1): 1–14.
- Whyman, T.; Murrup-Stewart, C.; Carter, A.; Young, U. M.; and Jobson, L. 2022. Ngarratja Kulpaana: Talking together about the impacts of lateral violence on aboriginal social and emotional well-being and identity. *Cultural diversity & ethnic minority psychology*, 28(2): 290.
- Wilson, K.; and Caliskan, A. 2024. Gender, race, and intersectional bias in resume screening via language model retrieval. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 1578–1590.
- Worth, S.; Snaith, B.; Das, A.; Thuermer, G.; and Simperl, E. 2024. Ai data transparency: an exploration through the lens of ai incidents. *arXiv preprint arXiv:2409.03307*.
- Xu, Y.; Zhang, J.; Salemi, A.; Hu, X.; Wang, W.; Feng, F.; Zamani, H.; He, X.; and Chua, T.-S. 2025. Personalized generation in large model era: A survey. *arXiv preprint arXiv:2503.02614*.
- Yang, Y. 2025. Racial bias in AI-generated images. *AI & SOCIETY*, 1–13.
- Young, I. M. 2002. *Inclusion and democracy*. OUP Oxford.
- Young, M.; Magassa, L.; and Friedman, B. 2019. Toward inclusive tech policy design: a method for underrepresented voices to strengthen tech policy documents. *Ethics and Information Technology*, 21(2): 89–103.
- Zack, T.; Lehman, E.; Suzgun, M.; Rodriguez, J. A.; Celi, L. A.; Gichoya, J.; Jurafsky, D.; Szolovits, P.; Bates, D. W.; Abdunour, R.-E. E.; et al. 2023. Coding inequity: assessing GPT-4's potential for perpetuating racial and gender biases in healthcare. *MedRxiv*, 2023–07.
- Zhang, L.; Liao, X.; Yang, Z.; Gao, B.; Wang, C.; Yang, Q.; and Li, D. 2024. Partiality and misconception: Investigating cultural representativeness in text-to-image models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–25.
- Zhou, K.; Gligorić, K.; Cheng, M.; Lam, M. S.; Raman, V.; Aminu, B.; Woo, C.; Brockman, M.; Cha, H.; and Jurafsky, D. 2025. Attention to Non-Adopters. *arXiv preprint arXiv:2510.15951*.
- Zhou, K.; Gligorić, K.; Cheng, M.; Lam, M. S.; Raman, V.; Aminu, B.; Woo, C.; Brockman, M.; Cha, H.; and Jurafsky, D. 2026. Attention to Non-Adopters. In *Findings of the Association for Computational Linguistics: ACL 2026*, 1336–1366.
- Zou, L.; and Khern-am nuai, W. 2023. AI and housing discrimination: The case of mortgage applications. *AI and Ethics*, 3(4): 1271–1281.