

Measuring What Matters: Connecting AI Ethics Evaluations to System Attributes, Hazards, and Harms

Shalaleh Rismani ¹, Renee Shelby ^{*2}, Leah Davis ^{*1}, Negar Rostamzadeh ², AJung Moon ¹

¹ McGill University

² Google Research

Abstract

Over the past decade, an ecosystem of measures has emerged to evaluate the social and ethical implications of AI systems, largely shaped by high-level ethics principles. These measures are developed and used in fragmented ways, without adequate attention to how they are situated in AI systems. In this paper, we examine how existing measures used in the computing literature map to AI system components, attributes, hazards, and harms. Our analysis draws on a scoping review resulting in nearly 800 measures corresponding to 11 AI ethics principles. We find that most measures focus on four principles – fairness, transparency, privacy, and trust – and primarily assess model or output system components. Few measures account for interactions across system elements, and only a narrow set of hazards is typically considered for each harm type. Many measures are disconnected from where harm is experienced and lack guidance for setting meaningful thresholds. These patterns reveal how current evaluation practices remain fragmented, measuring in pieces rather than capturing how harms emerge across systems. Framing measures with respect to system attributes, hazards, and harms can strengthen regulatory oversight, support actionable practices in industry, and ground future research in systems-level understanding.

1 Introduction

How do we know whether an Artificial Intelligence (AI) system adheres to AI ethics principles, and how does such adherence relate to the evaluation of algorithmic harm? As AI systems become more deeply embedded in high-stakes domains, there is a growing need to go beyond technical performance and examine the adverse impact these systems have on individuals and communities. This need is reflected in recent regulatory and policy initiatives, including the European Union’s (EU) AI Act and the National Institute of Standards and Technology’s (NIST) AI Risk Management Framework (EU 2024b; NIST 2023). These initiatives emphasize establishing appropriate measurement protocols to identify and mitigate sociotechnical harm, define thresholds for unacceptable system behavior, and enable ongoing oversight and monitoring (EU 2024a, Article 9). These regulatory advances are supported by, and often built directly on,

the wide range of approaches developed by the AI ethics and responsible AI (RAI) communities to assess normative qualities of AI systems (e.g., fairness (Morley et al. 2023)).

Despite the proliferation of measures designed to assess AI ethics principles, current evaluation practices fall short in two key ways. First, from a measurement theory standpoint, many existing measures lack construct validity and reliability— they do not consistently or accurately capture the normative qualities they aim to assess (Jacobs and Wallach 2021; Hutchinson et al. 2022). Without such guarantees, these measures risk producing misleading or incomplete insights about an AI system’s attributes and performance (Jacobs 2021; Xiao et al. 2023; Salaudeen et al. 2025).

Second, most measures focus on isolated components— such as models or datasets— rather than the AI system as a whole or cross-component interactions. Yet, evaluating an individual component is not equivalent to assessing a systems-level property (Barocas et al. 2021; Vera Liao and Xiao 2023). This is especially problematic from a system safety perspective, where safety is understood as an *emergent* property arising from the interactions among technical, human, and organizational components (Leveson 2011; Dobbe 2022). Even when each component functions as intended, harms can emerge from inappropriate interactions.

In this paper, we build on this second criticism to critically examine the efficacy of measures used to assess an AI system’s adherence to AI ethics principles, hereby referred to as RAI measures. In particular, we reflect on the role that these measures play in identifying, preventing, and responding to sociotechnical harms stemming from complex systems. We address the following research questions:

- **RQ1:** How do measures of adherence to AI ethics principles map to different components of an AI system, and what attributes do they capture?
- **RQ2:** What types of hazards and harms do these measures signal?

We conducted a scoping review of computing literature that claims to measure adherence to AI ethics principles. We extracted a total of 791 measures from 257 academic articles and performed a reflexive analysis to examine how these measures relate to system components, attributes, hazards, and harms. We find that existing measures are unevenly distributed across principles. Fairness, transparency, privacy,

*These authors contributed equally.




Figure 1: Example of the relational sequencing between a system component, its attribute, measure, hazard, and harm.

and trust account for the vast majority of measures, while principles like dignity, responsibility, and sustainability remain underrepresented. Most measures assess properties of the model and output as individual components, rather than evaluating hazards that emerge from interactions across the AI system. There is little documentation on how the measures relate to harms experienced by users, and what thresholds should warrant changes in design or deployment decisions. Based on these findings, we call on the AI community to take on more temporally responsive, systems-level approaches to evaluating AI systems.

We contribute a multi-dimensional analysis of RAI measures that identifies which ethics principles, system components, and types of harm are prioritized in the literature, as well as key attributes and hazards evaluated for each harm type. Drawing on a system safety perspective, we distill cross-cutting themes that highlight gaps and opportunities for more comprehensive, integrated, and system-aware evaluation practices. We also provide our dataset of measures, categorized by system components, attributes, hazards, and harms, and an interactive visualization to support further exploration and engagement. These contributions advance a systems-level perspective on RAI evaluation, offering concrete implications for computing researchers, industry practitioners, and policymakers.

2 Background

Measurement theory defines a measurement as the process of assigning a number to a single aspect of an object (Hand 2004). As such, a measurement is often a simplified indicator (e.g., height in centimeters) for a more complex attribute (e.g., a person’s size). In AI system development, various measurements are taken across stages like dataset curation and model training to guide refinement or improvement. In this paper, we focus on measures used by the computing community to assess the ethical normative qualities of AI systems. Next, we briefly review RAI evaluation practices and challenges, and then introduce concepts from system safety that can inform these approaches.

2.1 RAI Evaluation Practices and Challenges

Traditional performance measures and benchmarks have long been used to assess the quality of AI systems and track progress (Reuel et al. 2024; Orr and Kang 2024). In response to growing societal concerns, the scope of AI evaluation has expanded beyond performance to include social and ethical impacts since the early 2010s (Barocas, Hardt, and Narayanan 2019). These newer RAI assessment approaches fall into two broad categories.

The first category includes largely qualitative assessments of impact, harm, and risk assessments designed to exam-

ine an AI system’s effect on the stakeholders within its deployment context (Sloane and Moss 2023; Watkins et al. 2021). These methods aim to produce a context-specific understanding of harms and risks, to support governance initiatives and responsible development. The second category includes quantitative measures that aim to signal an AI system’s normative qualities beyond simple performance metrics. Many of these measures were developed to operationalize AI ethics principles (e.g., fairness, privacy, and transparency) into practice (Jobin, Ienca, and Vayena 2019; Morley et al. 2023). Examples include user fairness affinity, self-reported trust, privacy budget, and embodied emissions (Franklin et al. 2022; Vereschak, Bailly, and Caramiaux 2021; Carvalho, Pereira, and Cardoso 2019). Such measures may be used in risk assessments, post-deployment audits, or to evaluate trade-offs between technical performance and ethical considerations (Gardner, Gursoy, and Kakadiaris 2022; Zhao, Kaafar, and Kourtellis 2020).

While these quantitative RAI efforts represent important progress, scholars have highlighted key limitations in their design and application. A central concern is that measures lack construct validity and reliability. For instance, Jacobs and Wallach (2021) use measurement theory to show how common fairness metrics often conflate theoretical constructs with their operationalizations. Xiao et al. (2023) and Blodgett et al. (2021) apply similar critiques to natural language generation models and fairness benchmark datasets, respectively, demonstrating that their metrics frequently fail to capture the concepts they aim to assess.

Furthermore, most existing measures have a narrow scope, focusing on isolated components, such as model outputs, rather than evaluating AI systems in context. Recent work emphasizes that meaningful evaluation requires a systems-level perspective that accounts for sources of harm across AI system components, and considers for whom the system’s effects are most consequential (Diaz and Madaio 2025; Weidinger et al. 2023). Barocas et al. (2021) distinguish between evaluating system-wide performance to reveal deployment-related harms and assessing individual components to identify their underlying causes. Likewise, Liao and Xiao (2023) critique current evaluation practices for failing to bridge the “sociotechnical gap” between model behavior and downstream human impact. Without evaluating harms across components and contexts, we risk a false sense of systems-level oversight. Lastly, measurement practices often remain ad hoc and fragmented, driven more by convenience, precedent, or tool availability rather than by systematic reasoning (Rismani et al. 2023b; Madaio et al. 2022). These challenges underscore the need for a more structured, systems-oriented approach to evaluation.

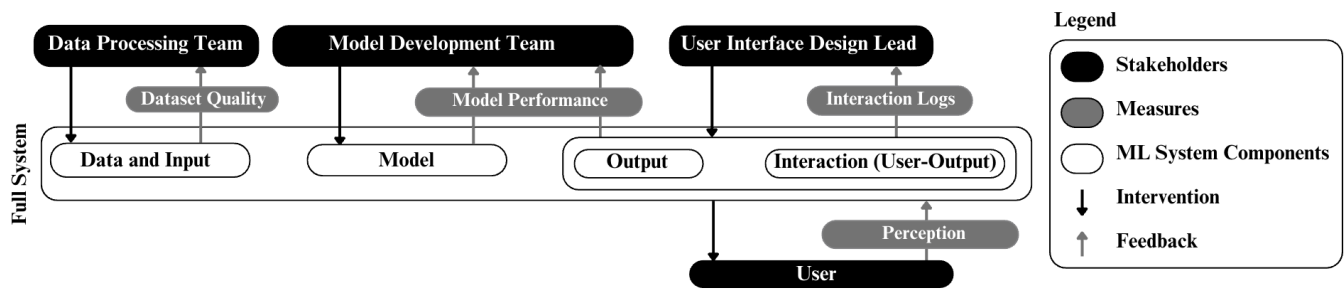


Figure 2: A conceptual model of feedback loops in a sociotechnical AI system. Measures provide signals about the status of system components, enabling stakeholders to make appropriate interventions.

2.2 A System Safety Approach to Evaluation

Recent work suggests that system safety engineering frameworks can help developers identify and mitigate potential sources of sociotechnical harms by examining the interconnected network of AI system components and actors (Qi et al. 2023; Dobbe and Wolters 2024; Rismani et al. 2023a). These works emphasize the vital role measurement plays in enabling safe design (Dahlgren Lindström et al. 2025; Rismani, Dobbe, and Moon 2024). Measures are necessary to quantify and monitor key *attributes* – observable properties of a system or a component – and assess whether they meet predefined safety criteria (Sánchez 2010). From this perspective, a *hazard* is a condition where one or more attributes deviate from expected parameters, creating the potential for a harmful outcome (Leveson 2011; Leveson and Thomas 2018). A *criterion* is the threshold or standard used to judge whether an attribute is within safe bounds; violating a criterion can signal a hazardous state (Firesmith 2004). If a system enters a hazardous state, it could lead to *harm* defined as the “adverse lived experiences resulting from a system’s deployment and operation in the world” (Shelby et al. 2023, p. 723). Figure 1 illustrates this relational chain.

System safety also frames measurement as a feedback mechanism that signals whether system elements are functioning within acceptable bounds, allowing actors to intervene when conditions drift toward harm (Leveson 2011). As shown in Figure 2, measures of dataset quality, model performance, interaction logs, and perception, act as feedback signals for higher-level components and actors. Dataset quality measures can help data teams refine or rebalance data; model performance measures can guide design choices; and interaction and perception-based evaluations can reveal user experiences, enabling improvements to system behavior and outcomes.

In this work, we leverage these system safety concepts to analyze existing measures designed to assess adherence to AI ethics principles. This lens allows us to assess the adequacy of current evaluation practices not only in terms of *what* they measure but also *how effectively* they help anticipate and mitigate harm across the entire sociotechnical system and its embedded context.

3 Methodology

Our methodology, as illustrated in Figure 3 consisted of three steps: (1) a scoping review to identify measures from recent computing literature; (2) a reflexive analysis to map each measure to specific attributes, hazards, and harms; and (3) compiling the results into a dataset and interactive visualization to support further exploration.

3.1 Scoping Review: Identifying the Measures

Our scoping review focused on measures related to the eleven AI ethics principles articulated by Jobin et al. (2019), given their comprehensive synthesis of ethical AI principles. We followed the five-stage scoping review process outlined by Arksey et al. (2005) given its suitability for mapping emerging fields (see also: Peters et al. (2021)), using the Covidence platform to manage the review. The lead author oversaw the entire process, with other team members contributing to various stages.

Identify research questions We posed the following foundational question to guide this scoping review: *What measurements and measurement processes are computing researchers proposing to assess an AI system’s adherence to ethical principles?* We focused on literature from computing-related venues, as technologists are the primary developers and implementers of these evaluation approaches.

Identify relevant studies We employed a three-phase search strategy combining database searches, citation-based reviews, and expert consultation. We queried the ACM Digital Library and IEEE Xplore via the Web of Science for papers published between January 2017 and August 2023. Our search queries incorporated terms related to measurement, AI, and the eleven principles stated by Jobin et al. (2019). Our initial search terms yielded over 8,000 results, warranting refinements. After testing 108 queries for relevance, we selected 22, which yielded 1,677 papers for screening. We then conducted a citation-based review on the papers selected for full-text review to identify 310 additional papers, including influential works from 2011–2016. Finally, three domain experts recommended an additional 40 papers for review. The collection of studies concluded in February 2024.

Study selection After removing 436 duplicates, 1,618 papers proceeded to title and abstract screening by four re-

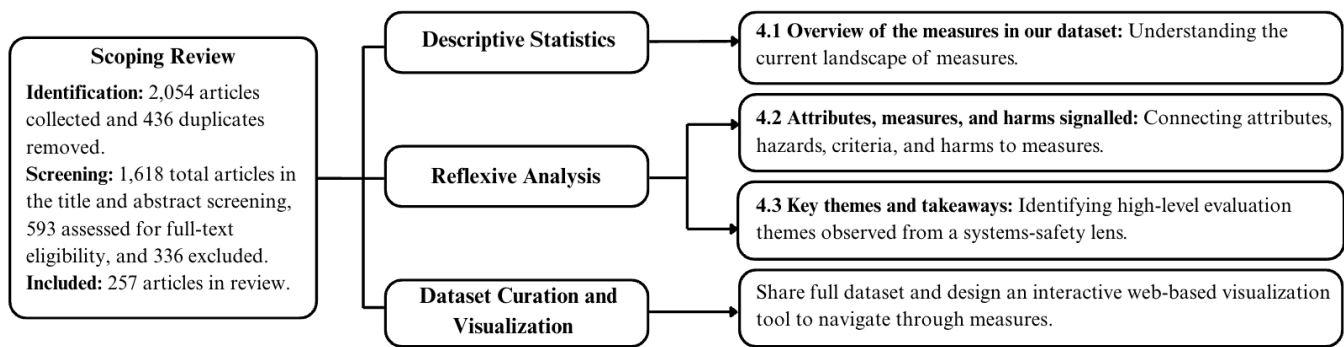


Figure 3: Overview of the primary methods conducted in this work.

searchers using predefined inclusion and exclusion criteria. We included any peer-reviewed studies that proposed a measure to assess how a non-embodied AI system adheres to one of the 11 ethics principles from Jobin et al. (2019). Each paper was initially screened by at least two of the four researchers. Papers with two “yes” votes and two “no” votes were included or excluded, respectively. Conflicts were resolved by the lead author. This stage resulted in 593 papers being passed to full-text review, from which 257 papers were included for data extraction. A PRISMA flow diagram is available in Appendix A of the supplemental materials in our GitHub repository (Rismani et al. 2025).

Data charting and summarization The data charting and summarization process involved two stages: extracting information from the papers and then separating multiple measures from each paper into distinct entries. A team of three researchers, including the lead author, performed this step. For each measure discussed in a paper, we used templates in Covidence and Google Sheets to extract key elements, including the ethics principle, measure name (e.g., demographic parity), measurement process (e.g., comparing positive classification rates), assessment type (e.g., statistical), the system component(s) being measured (e.g., output), the application area (e.g., criminal justice), and publication metadata. Drawing on the OECD’s definition of an AI system (OECD 2024), we mapped each measure to one or more system components: data/input, model, output, user-output interaction, or the full system (as shown in Figure 2). In addition, we categorized measures into four assessment types: mathematical (based on formal equations or logic), statistical (derived from data distributions and patterns), behavioral (based on observed user actions), and self-reported (based on users’ subjective responses). As many papers proposed multiple distinct measures, this process yielded a final corpus of 791 individual entries.

3.2 Reflexive Analysis: Identifying the Criterion, Attribute, Hazard, and Harm

For each of the 791 measures, we analyzed the measurement process taking place and the context provided in the source paper. We revisited source papers as needed to extract additional information. For each measure, we identified whether

the author(s) provided an explicit criterion (e.g., equal positive prediction rates across groups) for the measure (e.g., demographic parity), and subsequently categorized the measure’s attribute (e.g., distribution of system output), hazard (e.g., disparate distribution of output), and related primary and secondary sociotechnical harms (e.g., allocative harm through the denial of loans or representational harm through the reinforcement of stereotypes).

For criteria, attributes, and hazards, we developed inductive codes from our detailed review of the papers. For harm categorization, we applied the five categories from Shelby et al.’s sociotechnical harm taxonomy: allocative, representational, quality of service, interpersonal, and social system harms. Two researchers independently coded the first 100 measures and met to establish consensus. The remaining measures were coded collaboratively, with the lead author handling the majority of entries and consulting with the second researcher as needed to refine coding decisions. The full team then iteratively reviewed the completed dataset to identify key patterns connecting principles, attributes, hazards, and harms, using reflexive thematic analysis (Braun and Clarke 2006, 2021).

3.3 Dataset Compilation and Visualization

The final stage involved data pre-processing, statistical analysis, and visualization. After cleaning the raw data for lexical inconsistencies, we calculated the counts and proportions of the measures across the 11 ethical principles, system components, and harm categories. The final dataset includes 18 columns covering all 791 measures, their corresponding measurement process, system component(s), assessment factors (criterion, attribute, and hazard), primary harm, secondary harm (if present), and publication metadata. To facilitate further exploration of this dataset, we created an interactive Sunburst diagram. The description of all the features, the full dataset, and the visualization are included in our GitHub repository (Rismani et al. 2025).

3.4 Limitations of Our Study

We acknowledge several limitations. First, our scoping review captures a snapshot of a rapidly evolving field; many relevant measures may be unpublished or exist only in preprints. In addition, not all RAI measures are explicitly

framed in relation to the AI ethics principles outlined by Jobin et al. (2019). Second, our analysis is shaped by our team’s interpretive lens; it is inherently reflexive and introduces some subjectivity in selection, extraction, and categorization. Third, while we adopted a specific harm taxonomy, alternative frameworks exist that could offer complementary perspectives.

4 Findings

We begin with an overview of the measures in our dataset. Next, we examine how these measures relate to harms by analyzing the attributes they quantify and the hazards they represent. We close by discussing key themes from our reflexive analysis.

4.1 Overview of the Measures

In this section, we provide a high-level overview of how the measures identified in our corpus ($n = 791$) map onto AI ethics principles, components of AI systems, and types of sociotechnical harms. We observe clear patterns in which principles receive the most attention, which system components are most frequently measured, and how different categories of harms are represented. Notably, the sum of the calculated percentages exceeds 100% as measures can be coded under multiple principles/components. Aggregate percentage ratios use 791 as the denominator, whereas component or harm-type breakdowns use totals within each principle.

Four principles account for 90% of all measures We found a highly uneven distribution of measurement efforts across principles. Approximately 90% ($n = 717$) of measures coded fall under just four principles: fairness (46.4%, $n = 367$), transparency (21.1%, $n = 167$), privacy (14.2%, $n = 112$), and trust (10.6%, $n = 84$), as shown in Figure 4. We refer to these as *high-prevalence principles*. The remaining seven principles — freedom and autonomy, beneficence, non-maleficence, solidarity, dignity, sustainability, and responsibility — account for approximately 10% ($n = 79$) of the measures, which we term as *low-prevalence principles*. Notably, a small subset of measures (2.9%, $n = 23$) were coded with more than one principle, such as transparency/trust, fairness/privacy, and dignity/fairness.

Disproportionate focus on model and output components

We also observed a pronounced imbalance in which AI system components are evaluated. Across all measures, the vast majority target either the model (57.3%, $n = 453$) or the output (42.6%, $n = 337$). Far fewer assess the data/input (16.2%, $n = 128$) or the user-output interaction (25.4%, $n = 201$). This concentration is especially evident within the high-prevalence principles. For instance, 87.5% ($n = 321$) of all fairness-related measures assess either the model or the output. Similarly, 67.7% ($n = 113$) of transparency measures focus on the model or output, though a notable portion (32.3%, $n = 54$) also targets the user-output interaction. In contrast, privacy measures primarily assess attributes of the data/input (63.4%, $n = 71$), whereas trust measures overwhelmingly target the user-output interaction (96.4%, $n = 81$). Among the low-prevalence principles, the majority of

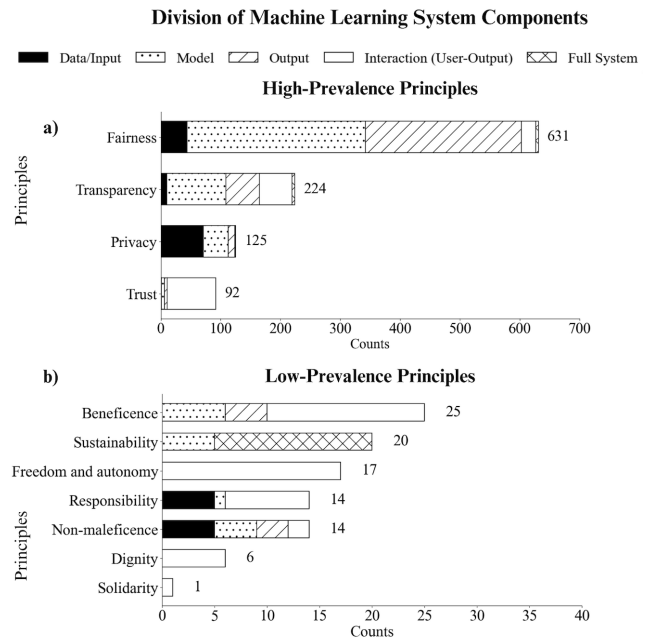


Figure 4: Division of machine learning system components categorized by (a) high-prevalence principles and (b) low-prevalence principles. The total count is above $n = 791$ as many measures are cross-listed with more than one principle and system component.

measures relate to the user-output interaction (62.0%, $n = 49$). Specifically, all measures related to the principles of solidarity, dignity, and freedom and autonomy are measures of user-output interaction. A key exception is sustainability, where 88.2% ($n = 15$) of measures assess the full system (e.g., total energy consumption (Henderson et al. 2020)), accounting for most of the full-system measures in our corpus.

Uneven representation of harm types across principles

Our analysis reveals that existing measures span all five types of sociotechnical harm outlined by Shelby et al. (2023), but their distribution is uneven. The most frequently addressed harms are interpersonal (29.2%, $n = 231$), involving instances where AI systems “adversely shape relations between people or communities” (Shelby et al. 2023, p. 727) and social system (25.9%, $n = 205$), reflecting “adverse macro-level effects” of AI systems (p. 731). Allocative harms account for 25.0% ($n = 198$), capturing cases of economic and/or opportunity loss due to AI deployment. The remaining measures correspond to representational harm (13.1%, $n = 104$) describing the reproduction of “unjust societal hierarchies”, and quality of service harms (6.7%, $n = 53$) which reflect performance disparities across identity groups.

Measures for each high-prevalence principle tend to concentrate on a few harm types (Figure 5). For example, 95.6% ($n = 351$) of fairness measures correspond to representational, allocative, or quality of service harms. The majority of transparency measures (83.2%, $n = 139$) map to social system harms. Privacy measures are almost exclusively linked to interpersonal harms (95.5%, $n = 107$), while nearly

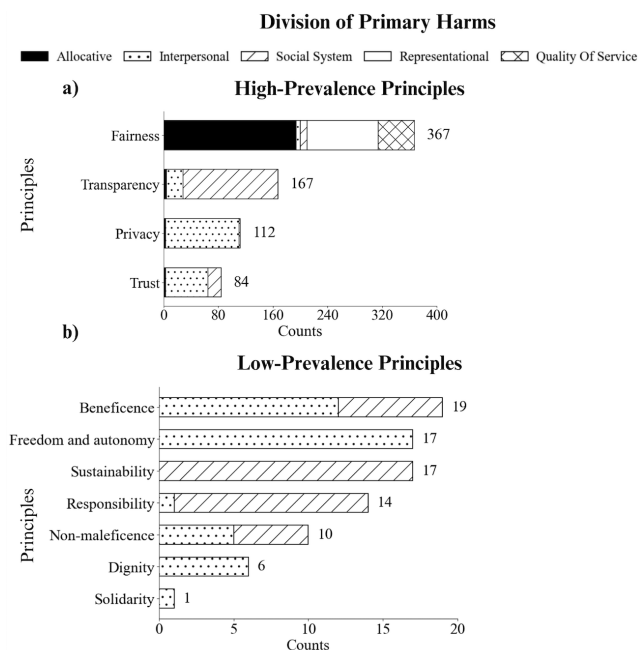


Figure 5: Division of primary harms categorized by (a) high-prevalence principles and (b) low-prevalence principles.

all of the trust measures (96.4%, $n = 81$) are linked to interpersonal and social system harms. In contrast, measures tied to low-prevalence principles are only linked to interpersonal (53.2%, $n = 42$) or social system harms (46.8%, $n = 37$).

4.2 Attributes, Hazards, and Harms Signaled

In this section, we examine the attributes of AI systems quantified by the measures and the hazards they signal for each of the five harm types. The total number of measures related to each harm type is presented in parentheses along the section headings. For each type of harm, we identify and describe the most common categories of measures, system components, attributes, and associated hazards. Note that many of the measures in our corpus can be interpreted through the lens of multiple types of harms. In such cases, we categorized the measures based on the original framing of the measures in the source articles. In addition, common hazards include disparities across system outcomes, errors, and performance between groups, which occur across representational, allocative, and quality-of-service harm types. While these hazards often involve similar measures, we assigned harm types according to the framing in the source paper. A more detailed and illustrative thematic summary of each of these mappings can be found in Appendix C in the supplementary materials on our GitHub (Rismani et al. 2025).

Representational harm (104 measures) Representational harms “occur when algorithmic systems reinforce the subordination of social groups along the lines of identity” (Shelby, 2023, p. 728, as cited in Katzman et al., 2021). Most measures signaling this harm are mathematical or statistical

assessments, quantifying attributes of the model or output (79.8% and 71.2%, respectively). A smaller share (19.2%) of measures target the data/input component, and none relate to the full system or user-output interaction stages. We identified three categories of measures signaling representational harm which we described below.

Representation of identities: These measures quantify how different social groups are represented in the data/input, model, and output of an AI system. Measures associated with the data/input component most often point to the hazard of *imbalanced group representation in the training data*. For instance, *person prominence* is a measure of how prominently individuals from different groups appear in visual training content (Wang, Narayanan, and Russakovsky 2020). Additionally, the *difference in reconstruction error* captures disparities in how accurately different groups are represented in compressed data (Pelegrina et al. 2022). The measures that observe attributes of a model signal the *reinforcement of stereotypes through learned embeddings* as a hazard. They typically quantify how different groups are represented within the model. Included in this category are measures like the *cosine similarity of images* and *word embedding association tests*, which quantify stereotypical associations by analyzing patterns in the model embedding space (Goyal et al. 2022; Malik et al. 2022). The measures that observe the output of an AI system quantify the presence of protected groups in its results. They signal an *underrepresentation or omission of specific protected groups in the outcomes of an AI system* as a hazard. Examples include the *diversity score* measure, which reflects the variety of representation across social attributes in the system’s outcomes, and the *percentage of content by group* measure, which quantifies the proportion of representation of specific identity groups (Solaiman et al. 2023; Ferraro, Serra, and Bauer 2021).

Reinforcement of social hierarchies: These measures focus on assessing who receives disproportionate errors or benefits, directing attention to whether the system’s outcomes are distributed in ways that reinforce existing social hierarchies. The measures and hazards are similar to those associated with allocative harm, which are further described in the following section under the category of *disparate outcome distribution*.

Misalignment with user identity: These measures assess the degree to which system outputs resonate with users’ identities, values, and cultural contexts. For example, the *steerability* measure quantifies how well an AI model can align its responses with the perspectives of specific demographic groups when prompted with contextual cues (Santurkar et al. 2023), and the *commonality* measure indicates whether users collectively encounter content that fosters a shared cultural experience (Ferraro et al. 2024). The key hazard associated with this perspective is that an AI system *creates outcomes that are irrelevant, unfamiliar, or misaligned with users’ identities and experiences*.

Allocative harm (198 measures) A hazard enables allocative harm when a system’s distribution of information, resources, or opportunities adversely affects histori-

cally marginalized groups (Shelby, 2023, p. 729, as cited in Barocas, Hardt, and Narayanan, 2019 and Richardson and Gilbert, 2021). The vast majority of measures associated with this harm type quantify attributes of the model (83.8%) and its output (69.7%) using statistical or mathematical approaches. A smaller proportion observes attributes of the data/input (11.6%). Measures of user-output interaction account for 7.1% of the total. Compared to representational harm, this set contains a higher proportion of self-reported and behavioral measures. Overall, we identified four categories of allocative harm measures as described below.

Disparate outcome distribution: These measures quantify how beneficial or erroneous system outputs are distributed across different social groups. For example, *disparate impact* measures the ratio of positive outcomes (e.g., loan approvals) between disadvantaged and advantaged groups (Besse et al. 2021). *Individual-user-to-individual-item fairness* indicates whether each user-item pair receives exposure aligned with a fairness target (Wu, Zhang, and Wu 2018). Both measures signal the hazard that *certain groups systematically receive fewer beneficial outcomes from the AI system*. Similarly, measures of error distribution, such as *treatment equality* (the ratio of false negatives to false positives) (Berk et al. 2021), help observe how model errors are distributed across groups. These measures signal the hazard that *error rates disproportionately affect specific groups*.

Disparate model performance: These measures capture *group-level differences in model performance*, leading to disparities in access to resources or opportunities. The measures and hazards are similar to those associated with the quality of service harms described under the category of *performance quality disparities* in the following section.

Perceived decision fairness: These measures capture users' fairness perception of AI systems. These are primarily measures of user-output interaction, which quantify a user's perceived understanding of algorithmic decisions and perceived fairness of the AI-based process taking place. Examples include self-reported understanding (e.g., agreement with "I understand the process by which the decision was made") (Binns et al. 2018) and perceived fairness ratings after viewing explanations of a training dataset (Anik and Bunt 2021). These measures signal the hazard that *users may not understand an AI system's reasoning or may perceive its decisions to be biased or unjust*.

Design-level disparities: These measures capture upstream design choices that lead to outcome disparities, such as the decisions surrounding features and labels in the input dataset and model design choices that incorporate fairness considerations. Examples include *class imbalance and conditional disparity in labels* measures, which quantify gaps in representation or label distribution across groups in a dataset (Hardt et al. 2021), and the *price of fairness* measure, which captures the accuracy-fairness trade-off by quantifying the increase in model loss required to reduce disparity (Berk et al. 2017). Both measures signal the hazard that *poor quality data or modeling strategies may produce or reinforce disparities in system outcomes*.

Quality of service harm (53 measures) Quality of service harms "occur when algorithmic systems disproportionately underperform for certain groups of people along social categories." (Shelby, 2023, p. 730). Most of the measures (90.6%) signaling this harm type assess the model or output components of an AI system using statistical or mathematical methods. Only five measures are classified at the user-output interaction, and none are found for the data/input or full system components. We observed two categories of measures, which we describe next.

Performance quality disparities: These measures capture differences in model performance across groups using measures like *precision*, *F1-score*, or *exposure*. Whereas allocative harm focuses on unequal access to resources, these measures focus on whether the system functions reliably and accurately for all users, signaling the hazard that *the model performs unevenly across certain groups*.

User experience quality: These measures capture the perceived quality of the user-output interaction among a user(s) and an AI system. For instance, the *Hellinger distance* measure can quantify the divergence between the distribution of recommended content across demographics (e.g., gender of artists) and a user's historical behavior (Ferraro, Serra, and Bauer 2021). A high Hellinger distance indicates that recommendations deviate significantly from a user's past preferences, signaling potential misalignment. The resulting hazard is that *system outputs are less useful and satisfactory for certain groups*.

Interpersonal harm (231 measures) Interpersonal harm refers to "instances when algorithmic systems adversely shape relations between people or communities" (Shelby et al. 2023, p. 730). The associated measures provide observations of privacy violations, loss of agency, and diminished well-being. Notably, we did not identify any measures for the technology-facilitated violence harm sub-type. Overall, the measures in this harm category assess the user-output interaction (49.4%), data/input (31.2%), and model (22.5%) components, with a few measures relating to the output (7.4%). Among this harm type, measurement approaches are diverse, with 50.2% being mathematical or statistical, 38.5% self-reported, and 12.1% behavioral. We identified five different categories, which we describe next.

Design-level privacy preservation: These measures indicate the risk of re-identification or the level of privacy protection applied to a dataset and/or model. Example measures include the *inverse trace of the Fisher Information Matrix*, which quantifies the difficulty of reconstructing sensitive variables from released data (Farokhi 2021); the *K-value*, representing the minimum group size required to make individual records indistinguishable (González-Zelaya et al. 2023); and the differential privacy parameters ϵ and δ , which quantify the maximum privacy loss and the probability that this guarantee fails (Jarín and Eshete 2022). These measures signal the hazards that *individuals can be identified due to patterns in the dataset or privacy protections are either too weak to be effective or so strong that they undermine data utility*.

Privacy attack risk: These measures reflect the model's

susceptibility to adversarial inference and the amount of sensitive information that can be extracted. For instance, the *adversary accuracy* measure estimates the likelihood that an attacker correctly infers a sensitive label (Xiao and Khisti 2019), while *mutual information* measures how much private information is leaked through model outputs (Zhang and Venkatasubramaniam 2018). These measures relate to the hazard of *successful privacy attacks that expose user-level information and inference methods that allow significant reconstruction of individual data*.

Influence on user behavior: This category of measures captures the AI system’s influence on the user. For example, *acceptance rates* quantify how often users follow AI-generated suggestions (Zhang et al. 2023), and *decision threshold shifts* measure how much a user’s decision boundary changes when aided by the AI system (Poursabzi-Sangdeh et al. 2021). These signal the hazard that users may be *unduly influenced or become over-reliant on AI system outputs*.

Perceived level of influence: Related to the previous category, these measures quantify users’ perceived control and perceived autonomy of the AI system. *Perceived control Likert scales*, for example, capture a users’ sense of influence over outcomes, while *mind perception agency scales* reflect how much autonomy and independent decision-making users attribute to the AI system (Kim et al. 2023; Gonzalez et al. 2022). These measures signal the hazard that *users feel disempowered or incorrectly believe the system acts with more autonomy than it actually possesses*.

Impact on user well-being: These measures observe the impact of a system’s output on user well-being, quantifying attributes like the harmfulness of generated content, perceived quality of the system output, and perceived respectfulness. The *HONEST score*, computed as the proportion of completions containing offensive content, is an example measure of the harmfulness of generated content (Nozza, Bianchi, and Hovy 2021). *Perceived helpfulness and risk ratings* capture users’ subjective impressions of an AI system’s utility or danger (Przybya and Soto 2021), and *dehumanization scales* offer a self-reported measure of whether users feel objectified or treated as less than fully human in their interactions with AI (Formosa et al. 2022). These measures signal hazards such as an *exposure to harmful content, distorted perceptions of system utility or risk, and a lack of perceived dignity or respect in AI-mediated experiences*.

Social system harm (205 measures) Social system harms “reflect the adverse macro-level effects of new and reconfigurable algorithmic systems, such as systematizing bias and inequality and accelerating the scale of harm.” (Shelby 2023, p. 731) All measures in this category relate to information loss or environmental harms; none reflect the harm subtypes of cultural, socio-economic, and political/civic harm. The measures are distributed across multiple system components: model (52.2%), output (31.2%), user–output interaction (33.2%), full system (10.7%), and data/input (6.3%). About three-quarters are mathematical or statistical; the remainder are behavioral or self-reported assessments. We identified two main categories as described below.

Understanding of the system: These measures quantify the interpretability and understandability of the system, and the quality of the post-hoc method used to explain model predictions. The *perceived understanding* measure, for example, captures users’ self-reported comprehension and ability to anticipate the model’s behavior (Wang and Yin 2022). An example of post-hoc method quality measures includes the *fidelity score*, which reflects how much a model’s performance drops when features identified as influential are masked (Bonet et al. 2022). *Explanation completeness* is another example, which calculates the fraction of the overall explanation retained in a reduced subset (Mariotti, Alonso-Moral, and Gatt 2022). These measures signal hazards that *users are unable to meaningfully interpret the system or are misled by low-fidelity explanations*, raising the risk that AI systems obscure underlying logic, or contribute to misinformed decision-making.

Resource intensity: These measures signal the hazard that *AI systems consume excessive resources and generate high emissions*. Energy consumption and carbon emissions are two attributes observed in these measures. For example, *average energy consumption* can be estimated from hardware-specific multiply-and-accumulate operations, while the *energy per sample* measure is calculated by dividing the real-time power usage by the model’s processing throughput (Yang et al. 2017). Carbon impact is captured through measures such as the *estimated greenhouse gas emissions*, which factors in regional electricity carbon intensity, and the *CO₂-equivalent output*, calculated using tools like the ML Emissions Calculator based on training time, hardware type, and location (Savazzi et al. 2023).

4.3 Key Themes and Takeaways

Our analysis reveals five cross-cutting themes about the current state of RAI measurement.

Measures are fragmented across system components

Most evaluation measures target isolated components of AI systems, particularly the model or its outputs, while overlooking how harms emerge from interactions across components. This lack of system-spanning evaluation has been noted before. For example, Black et al. (2023) highlight that most fairness interventions focus on the modeling stage, leaving earlier and later stages under-examined. A fairness measure of a dataset does not equate to the fairness performance of an AI system trained on it (Barocas et al. 2021). Likewise, a high-performing AI model that leaves users feeling unduly influenced can lead to unsafe AI system use.

Notably, a small subset of measures in our corpus explicitly engage with component interdependencies or aim to evaluate the system holistically. One example is *MARS-Gym*, a benchmark framework for reinforcement learning-based recommender systems that incorporates fairness measures, such as disparate treatment and mistreatment, to trace how sensitive inputs influence model decisions and affect outputs. Its design enables evaluation across the data and input, model, and outcome components, offering a more interconnected view of system behavior (Santana et al. 2020). A few sustainability-related measures, such as those by

Savazzi et al. (2023) and Ollivier et al. (2023), provide observations about the full system. These measures estimate environmental impact by combining energy consumption across various stages of development. While they do not model inter-component effects directly, they integrate attributes across stages to reflect overall system performance. To identify harms at the systems-level, assessments must go beyond modular checkpoints to examine how one component's attributes affect another, capturing dependencies and dynamics that could create unsafe outcomes.

Measurements are taken far from where harm is experienced There is a large gap between where normative measurements are taken and where harms are ultimately experienced. While many papers in our corpus are motivated to prevent potential sociotechnical harms in a general sense, the process of quantifying the attributes into measures or taking measurements from components of a system is often decoupled from a specific harm and the conditions under which it occurs. Critically, we found no instances where a component-level measurement was empirically traced to a downstream harm experienced by users. This disconnect reflects criticisms raised about the validity of common measures (Jacobs and Wallach 2021; Blodgett et al. 2021).

In our corpus, many measures are designed to evaluate system-level hazards instead of direct harm assessments. For example, the *HONEST score* measure quantifies the presence of offensive language in model outputs (Nozza, Bianchi, and Hovy 2021). While it observes a user-facing attribute, it does not directly assess the lived experience of the users. Similarly, Shen et al.'s (2023) *correlated sensitivity metric* estimates privacy leak risk based on dataset composition. While it highlights a potential privacy hazard, using it as a measure of harm is questionable, as it does not assess the experience of privacy violation. Crucially, the disconnect lies not in the measures themselves, but in how they are applied and interpreted within evaluation practices. As Barocas et al. (2021) argue, component-level measures can help identify hazards but offer limited insight into the lived experience of harm. We argue that evaluation practices should distinguish between measures that assess hazards and those that evaluate harm as lived experiences. The mapping presented in this paper offers an initial step toward this goal.

Measures lack clear criteria for identifying hazards Only a small subset of measures – approximately 26.9% ($n = 213$) – were accompanied by explicit criteria for interpretation. While not every measure requires a predefined criterion, such guidance becomes important when a measure is used to identify potential hazards or signal when system behavior may be unacceptable. From the measures that had stated criteria, few provided a discussion on appropriate thresholds. This is not surprising as thresholds are inherently context-dependent and may vary based on application domains, user populations, or risk tolerances. Most of the measures with stated criteria relate to fairness and focus on concepts such as group- or individual-level parity (e.g. *demographic parity criterion ensures that individuals with a protected attribute have the same likelihood of being classified as positive (e.g., recidivist) as those in a reference group*, as

paraphrased from (Tolan et al. 2019)). However, achieving parity is rarely feasible or expected. As a result, practitioners rely on judgment calls or heuristics such as the disparate impact rule, which states that the selection rate for a disadvantaged group should be at least 80% of that for the most advantaged group. Threshold ambiguity was also present in transparency measures. Example criteria include *good explanations should be sparse* or *good explanations should be stable*, as paraphrased from (Yuan et al. 2022). Yet these criteria are under-specified for operationalization. For example, determining how much detail is needed in a feature description or what level of consistency qualifies as stable are inherently subjective decisions, often left to the discretion of practitioners. From a systems safety perspective, thresholds are essential for distinguishing tolerable variation from conditions that signal a hazard. Evaluation practices can greatly benefit from outlining and disclosing what is deemed to be unacceptable behavior of a system component or a full system in relation to an attribute that is being measured. This clarity enables both those designing evaluations and those applying them to make informed, safety-aligned decisions.

Measures overlook how harms accumulate over time Although sociotechnical harms stemming from AI systems often emerge gradually, many of the measures are taken at a single point in time. While some statistical measures related to models and outputs are integrated into development and maintenance pipelines, many are not applied over time. Notably, none of the papers in our corpus discussed how their proposed measures could support long-term monitoring. A few papers related to environmental harm were exceptions, tracking energy consumption or carbon emissions throughout training and deployment (Henderson et al. 2020). These were among the only examples in our dataset where time was meaningfully integrated into the construction of the measure itself. From a system safety perspective, time-aware evaluation is critical for detecting emergent hazards and triggering timely intervention (Leveson and Thomas 2018). Studying how component-level attributes change over time, and when those changes signal a hazard and harm, is a valuable direction for future evaluation practice.

The role of perception-based measures in signaling harms remains unclear Perception-based measures composed 17.3% ($n = 137$) of the entire corpus. They are particularly prevalent in measures of interpersonal and social-system harms. These self-reported measures are central to evaluating user experiences. However, in our reflexive analysis, it was difficult to determine what hazards and harms these measures signal. Human perception is subjective, and self-reported measures have well-established drawbacks (Lai et al. 2023). Therefore, it is not surprising when users' perceptions of a system (e.g., perceived utility) do not align with other measures (e.g., actual system efficacy). While optimizing for a perception-based measure is inappropriate for designing safe AI, ignoring these measures could mean overlooking potential indications of harm experienced by users. The RAI community needs to pay more attention to unpacking the role perception-based measures can play in harm and risk assessments of AI systems.

5 Discussion

5.1 Implications for Governance and Policy

Measures are central to emerging AI governance and evaluation frameworks, whether they are principle-based, risk-based, or harm-based (Jacobs 2021; Brundage et al. 2020). Our analysis highlights several critical gaps. First, existing measures disproportionately focus on a narrow set of principles (recorded as high-prevalence principles). Second, measures tend to concentrate on a limited set of well-established hazards which risks narrowing the scope of what gets measured and prioritized. This raises a deeper question: what hazards should we actually be paying attention to when it comes to understanding and mitigating harm? Without reflection and consensus on which hazards matter across contexts, important dimensions of harm may be overlooked (Dobbe and Wolters 2024; Selbst et al. 2019). This resulting fragmentation is further compounded by the lack of clear evaluation criteria for which levels of measurement indicate the presence or severity of a hazard. For normative quality assessments of AI systems (e.g., AI ethics audits) to become a meaningful part of AI governance frameworks, a more comprehensive approach to identifying and monitoring harms and assessing adherence to a variety of ethical principles is needed (Ojewale et al. 2025). By making explicit the connections between what is being measured, why it matters, and how it relates to potential harms, evaluation practices can become more traceable and accountable (Kroll 2021). Our corpus of measures and the accompanying visualization can serve as a practical and reflexive resource to springboard such an effort.

5.2 Implications for Industry and Practice

Our analysis offers two key implications for AI practitioners. First, our mapping of measures to principles, system components, assessment types, hazards, and harms can help practitioners select or design measures in a more deliberate and structured way. Rather than creating ad hoc measures, teams can use the visualization of our dataset to explore a variety of measures aligned with their evaluation goals, system attributes, or regulatory needs. Second, with an understanding that measurement selection, application, and implementation can be resource-intensive, this work's framing supports more strategic prioritization (Madaio et al. 2022). By making explicit what a given measure captures— and what it does not— teams can better weigh trade-offs, justify their evaluation choices, and understand why one measure may be more appropriate than another (Barocas et al. 2021). This can improve internal coordination across product, compliance, and management roles while making evaluation efforts more transparent and defensible (Balayn et al. 2023).

5.3 Implications for Researchers

We find two key implications for the research community. First, there is a pressing need for more structured and transparent reporting when new measures are introduced. Many papers lack clarity on what component is assessed, which attribute is measured, or how the measure relates to

a specific harm or hazard. Without this scaffolding, measures are difficult to interpret or apply beyond their original context (Olteanu et al. 2025). Second, addressing the evaluation gaps identified (i.e., a disproportionate emphasis on certain system components and a narrow concentration on well-established hazards) requires stronger interdisciplinary collaboration (Kawakami et al. 2025). Limitations such as narrow principle coverage, limited attention to cross-component interactions, and the disconnect between where measures are applied and where harm is experienced highlight the need to move beyond computing-centric approaches. Engaging with fields such as psychology, media studies, and sociology can help develop measures that better reflect user experiences, integrate perception-based data more meaningfully, and bring systems-level behavior into closer alignment with the experiences of harm.

5.4 Future Work

Looking ahead, there are several promising directions for future work. While our analysis focused on application-specific systems, the growing use of general-purpose AI models introduces a new layer of complexity. These systems are harder to assess, with broader and less predictable use cases. They require more adaptive, multi-layered approaches to evaluation (Wallach et al. 2024; Weidinger et al. 2023). In addition, many organizations contribute to the development and deployment of different components within a product that uses general-purpose AI systems. Applying a system safety perspective to these evaluations can illuminate which system components are being assessed, where measures are missing, and which organizations are responsible for each part. Future work could also encourage more community reporting on how measures are used in practice across different types of AI systems (Dev et al. 2021). In parallel, we see value in expanding how we conceptualize system components, beyond the five we used, to better reflect the complexity of real-world deployments. Finally, the visualization tool we developed could be refined and tested for use within organizations. Its usefulness will likely vary across sectors, company sizes, and governance structures. These directions point to a broader shift: treating evaluation not as an isolated research task but as a core part of the sociotechnical infrastructure needed to support RAI.

6 Conclusion

This paper offers a systems-level analysis of how AI measures relate to ethical principles, system components, and sociotechnical harms. By applying a systems safety lens, we find that current practices lack measurements that capture cross-component dynamics and that a significant gap remains between where measurements are taken and where real-world harms occur. Our analysis, dataset, and visualization tool aim to support more grounded, transparent, and context-sensitive evaluation practices. There is a pressing need for measurement approaches that reflect how harms emerge within sociotechnical systems. This work takes a step toward addressing this need by offering both a critical mapping of the current landscape and a foundation for building more comprehensive evaluation infrastructures.

Positionality Statement

Our team, based in the United States and Canada, brings expertise from computer science, software engineering, HCI, robotics, sociology, and STS, with experience across industry and academia. This positioning shaped our approach to measurement, interpretation, and analysis. While we followed a rigorous process with multiple quality checks, we acknowledge the inherent subjectivity in our decisions, from selecting and extracting data to framing our analysis. Our familiarity with industry and academic measurement practices informed our understanding, particularly when resolving inconsistencies or gaps in the reviewed work. We aimed to critically engage with these measures while ensuring our interpretations remained faithful to the original sources.

Acknowledgments

We gratefully acknowledge Ava Gilmor and Bonam Mingole for their assistance with the scoping review, and Fernando Diaz for his insightful feedback throughout this project. This project was financially supported by Google.

References

- Anik, A. I.; and Bunt, A. 2021. Data-Centric Explanations: Explaining Training Data of Machine Learning Systems to Promote Transparency. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, 1–13. ACM.
- Arksey, H.; and O'Malley, L. 2005. Scoping studies: towards a methodological framework. *Int. J. Soc. Res. Methodol.*, 8(1): 19–32.
- Balayn, A.; Yurrita, M.; Yang, J.; and Gadiraju, U. 2023. “fairness toolkits, A checkbox culture?” on the factors that fragment developer practices in handling algorithmic harms. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, volume 45, 482–495. New York, NY, USA: ACM.
- Barocas, S.; Guo, A.; Kamar, E.; Krones, J.; Morris, M. R.; Vaughan, J. W.; Wadsworth, W. D.; and Wallach, H. 2021. Designing Disaggregated Evaluations of AI Systems: Choices, Considerations, and Tradeoffs. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 368–378. New York, NY, USA: Association for Computing Machinery.
- Barocas, S.; Hardt, M.; and Narayanan, A. 2019. Fairness and Machine Learning. <http://www.fairmlbook.org>.
- Berk, R.; Heidari, H.; Jabbari, S.; Joseph, M.; Kearns, M.; Morgenstern, J.; Neel, S.; and Roth, A. 2017. A Convex Framework for Fair Regression. In *Proceedings of the Workshop on Fairness, Accountability, and Transparency in Machine Learning (FATML)*.
- Berk, R.; Heidari, H.; Jabbari, S.; Kearns, M.; and Roth, A. 2021. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 50(1): 3–44.
- Besse, P.; del Barrio, E.; Gordaliza, P.; Loubes, J.-M.; and Risser, L. 2021. A Survey of Bias in Machine Learning Through the Prism of Statistical Parity. *The American Statistician*, 76(2).
- Binns, R.; Van Kleek, M.; Veale, M.; Lyngs, U.; Zhao, J.; and Shadbolt, N. 2018. “It’s Reducing a Human Being to a Percentage”: Perceptions of Justice in Algorithmic Decisions. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, 377. Preprint also available via SocArXiv: <https://doi.org/10.31235/osf.io/9wqxr>.
- Black, E.; Naidu, R.; Ghani, R.; Rodolfa, K.; Ho, D.; and Heidari, H. 2023. Toward operationalizing pipeline-aware ML fairness: A research agenda for developing practical guidelines and tools. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–11. New York, NY, USA: ACM.
- Blodgett, S. L.; Lopez, G.; Olteanu, A.; Sim, R.; and Wallach, H. 2021. Stereotyping Norwegian Salmon: An Inventory of Pitfalls in Fairness Benchmark Datasets. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 1004–1015. Online: Association for Computational Linguistics.
- Bonet, E. R.; Do, T. H.; Qin, X.; Hofman, J.; Manna, V. P. L.; Philips, W.; and Deligiannis, N. 2022. Explaining Graph Neural Networks With Topology-Aware Node Selection: Application in Air Quality Inference. *IEEE Transactions on Signal and Information Processing over Networks*, 8: 499–513.
- Braun, V.; and Clarke, V. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3(2): 77–101.
- Braun, V.; and Clarke, V. 2021. One Size Fits All? What Counts as Quality Practice in (Reflexive) Thematic Analysis? *Qualitative research in psychology*, 18(3): 328–352.
- Brundage, M.; Avin, S.; Wang, J.; Belfield, H.; Krueger, G.; Hadfield, G.; Khlaaf, H.; Yang, J.; Toner, H.; Fong, R.; Maharaj, T.; Koh, P. W.; Hooker, S.; Leung, J.; Trask, A.; Bluemke, E.; Lebensold, J.; O’Keefe, C.; Koren, M.; Ryffel, T.; Rubinovitz, J. B.; Besiroglu, T.; Carugati, F.; Clark, J.; Eckersley, P.; de Haas, S.; Johnson, M.; Laurie, B.; Ingerman, A.; Krawczuk, I.; Askill, A.; Cammarota, R.; Lohn, A.; Krueger, D.; Stix, C.; Henderson, P.; Graham, L.; Prunkl, C.; Martin, B.; Seger, E.; Zilberman, N.; heigeartaigh, S.; Kroeger, F.; Sastry, G.; Kagan, R.; Weller, A.; Tse, B.; Barnes, E.; Dafoe, A.; Scharre, P.; Herbert-Voss, A.; Rasser, M.; Sodhani, S.; Flynn, C.; Gilbert, T. K.; Dyer, L.; Khan, S.; Bengio, Y.; and Anderljung, M. 2020. Toward trustworthy AI development: Mechanisms for supporting verifiable claims. *arXiv [cs.CY]*.
- Carvalho, D. V.; Pereira, E. M.; and Cardoso, J. S. 2019. Machine learning interpretability: A survey on methods and metrics. *Electron. ETF*.
- Dahlgren Lindström, A.; Methnani, L.; Krause, L.; Ericson, P.; de Rituerto de Troya, I. M.; Coelho Mollo, D.; and Dobbe, R. 2025. Helpful, harmless, honest? Sociotechnical limits of AI alignment and safety through Reinforcement Learning from Human Feedback. *Ethics Inf. Technol.*, 27(2): 28.

- Dev, S.; Sameki, M.; Dhamala, J.; and Hsieh, C.-J. 2021. Measures and best practices for responsible AI. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. New York, NY, USA: ACM.
- Diaz, F.; and Madaio, M. 2025. Scaling Laws Do Not Scale. In *Proceedings of the 2024 AAAI/ACM Conference on AI, Ethics, and Society*, 341–357. London, England: AAAI Press.
- Dobbe, R. 2022. System Safety and Artificial Intelligence. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 1584. New York, NY, USA: Association for Computing Machinery.
- Dobbe, R.; and Wolters, A. 2024. Toward Sociotechnical AI: Mapping Vulnerabilities for Machine Learning in Context. *Minds and Machines*.
- EU. 2024a. Article 9: Risk Management System. <https://artificialintelligenceact.eu/article/9/>. Accessed: 2025-5-23.
- EU. 2024b. Artificial Intelligence Act - Regulation - EU - 2024/1689 - EN - EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689>. Accessed: 2025-5-23.
- Farokhi, F. 2021. Privacy-Preserving Public Release of Datasets for Support Vector Machine Classification. *IEEE Transactions on Big Data*, 7(5): 893–899.
- Ferraro, A.; Ferreira, G.; Diaz, F.; and Born, G. 2024. Measuring Commonality in Recommendation of Cultural Content to Strengthen Cultural Citizenship. *ACM Transactions on Recommender Systems*, 2(1): 1–32.
- Ferraro, A.; Serra, X.; and Bauer, C. 2021. Break the Loop: Gender Imbalance in Music Recommenders. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*, CHIIR '21, 249–254. ACM.
- Firesmith, D. 2004. Engineering Safety Requirements, Safety Constraints, and Safety-Critical Requirements. *J. Object Technol.*, 3(3).
- Formosa, P.; Rogers, W.; Griep, Y.; Bankins, S.; and Richards, D. 2022. Medical AI and human dignity: Contrasting perceptions of human and artificially intelligent (AI) decision making in diagnostic and medical resource allocation contexts. *Computers in Human Behavior*, 133: 107296.
- Franklin, J. S.; Bhanot, K.; Ghalwash, M.; Bennett, K. P.; McCusker, J.; and McGuinness, D. L. 2022. An Ontology for Fairness Metrics. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES '22)*, 265–275. Association for Computing Machinery.
- Gardner, J. W.; Gursoy, F.; and Kakadiaris, I. A. 2022. Accuracy-Fairness Tradeoff in Parole Decision Predictions: A Preliminary Analysis. In *Proceedings of the IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT)*, 284–287. IEEE/ACM.
- Gonzalez, M. F.; Liu, W.; Shirase, L.; Tomczak, D. L.; Lobbe, C. E.; Justenhoven, R.; and Martin, N. R. 2022. Allying with AI? Reactions toward human-based, AI/ML-based, and augmented hiring processes. *Computers in Human Behavior*, 130: 107179.
- González-Zelaya, V.; Salas, J.; Megías, D.; and Missier, P. 2023. Fair and Private Data Preprocessing through Microaggregation. *ACM Transactions on Knowledge Discovery from Data*, 18(3): 1–24.
- Goyal, P.; Soriano, A. R.; Hazirbas, C.; Sagun, L.; and Usunier, N. 2022. Fairness Indicators for Systematic Assessments of Visual Feature Extractors. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 70–88. ACM.
- Hand, D. J. 2004. *Measurement Theory and Practice*. John Wiley & Sons.
- Hardt, M.; Chen, X.; Cheng, X.; Donini, M.; Gelman, J.; Gollaprolu, S.; He, J.; Larroy, P.; Liu, X.; McCarthy, N.; Rathi, A.; Rees, S.; Siva, A.; Tsai, E.; Vasist, K.; Yilmaz, P.; Zafar, M. B.; Das, S.; Haas, K.; Hill, T.; and Kenthapadi, K. 2021. Amazon SageMaker Clarify: Machine Learning Bias Detection and Explainability in the Cloud. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '21, 2974–2983. ACM.
- Henderson, P.; Hu, J.; Romoff, J.; Brunskill, E.; Jurafsky, D.; and Pineau, J. 2020. Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning. *ArXiv*, abs/2002.05651.
- Hutchinson, B.; Rostamzadeh, N.; Greer, C.; Heller, K.; and Prabhakaran, V. 2022. Evaluation Gaps in Machine Learning Practice. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 1859–1876. New York, NY, USA: Association for Computing Machinery.
- Jacobs, A. Z. 2021. Measurement as governance in and for responsible AI. *arXiv [cs.CY]*.
- Jacobs, A. Z.; and Wallach, H. 2021. Measurement and Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: ACM.
- Jarin, I.; and Eshete, B. 2022. DP-UTIL: Comprehensive Utility Analysis of Differential Privacy in Machine Learning. In *Proceedings of the Twelfth ACM Conference on Data and Application Security and Privacy*, CODASPY '22, 41–52. ACM.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9): 389–399.
- Katzman, J.; Barocas, S.; Blodgett, S. L.; Laird, K.; Scheuerman, M. K.; and Wallach, H. 2021. Representational Harms in Image Tagging. In *Beyond Fair Computer Vision Workshop at CVPR 2021*. Workshop paper.
- Kawakami, A.; Blodgett, S. L.; Barocas, S.; Chouldechova, A.; Jacobs, A.; Sheng, E.; Vaughan, J. W.; Wallach, H.; Winecoff, A.; Wang, A.; Zhu, H.; and Holstein, K. 2025. AI Measurement as a Stakeholder-Engaged Design Practice: FAccT 2025 Translation Tutorial. Tutorial presented at the ACM Conference on Fairness, Accountability, and Transparency (FAccT). Thursday, 26 June, 3:15 PM – 4:15 PM, New Stage C.
- Kim, T.; Molina, M. D.; Rheu, M. M.; Zhan, E. S.; and Peng, W. 2023. One AI Does Not Fit All: A Cluster Analysis of

- the Laypeople's Perception of AI Roles. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, 1–20. ACM.
- Kroll, J. A. 2021. Outlining Traceability: A Principle for Operationalizing Accountability in Computing Systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, 758–771. New York, NY, USA: Association for Computing Machinery.
- Lai, V.; Chen, C.; Smith-Renner, A.; Liao, Q. V.; and Tan, C. 2023. Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 1369–1385. New York, NY, USA: Association for Computing Machinery.
- Leveson, N. 2011. *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press.
- Leveson, N. G.; and Thomas, J. P. 2018. *STPA Handbook*. MIT Partnership for Systems Approaches to Safety and Security (PSASS), Cambridge, Massachusetts, U.S. Freely available via MIT PSASS website.
- Madaio, M.; Egede, L.; Subramonyam, H.; Wortman Vaughan, J.; and Wallach, H. 2022. Assessing the Fairness of AI Systems: AI Practitioners' Processes, Challenges, and Needs for Support. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW1): 1–26.
- Malik, V.; Dev, S.; Nishi, A.; Peng, N.; and Chang, K.-W. 2022. Socially Aware Bias Measurements for Hindi Language Representations. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1041–1052. Association for Computational Linguistics.
- Mariotti, E.; Alonso-Moral, J. M.; and Gatt, A. 2022. Measuring Model Understandability by means of Shapley Additive Explanations. In *2022 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 1–8. IEEE.
- Morley, J.; Kinsey, L.; Elhalal, A.; Garcia, F.; Ziosi, M.; and Floridi, L. 2023. Operationalising AI ethics: barriers, enablers and next steps. *AI Soc.*, 38(1): 411–423.
- NIST. 2023. AI Risk Management Framework Playbook. <https://airc.nist.gov/airmf-resources/playbook/>. Accessed: 2025-5-23.
- Nozza, D.; Bianchi, F.; and Hovy, D. 2021. HONEST: Measuring Hurtful Sentence Completion in Language Models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- OECD. 2024. Explanatory memorandum on the updated OECD definition of an AI system. *OECD Artificial Intelligence Papers*, No. 8, OECD Publishing, Paris.
- Ojewale, V.; Steed, R.; Vecchione, B.; Birhane, A.; and Raji, I. D. 2025. Towards AI accountability infrastructure: Gaps and opportunities in AI audit tooling. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–29. New York, NY, USA: ACM.
- Ollivier, S.; Li, S.; Tang, Y.; Cahoon, S.; Caginalp, R.; Chaudhuri, C.; Zhou, P.; Tang, X.; Hu, J.; and Jones, A. K. 2023. Sustainable AI Processing at the Edge. *IEEE Micro*, 43(1): 19–28.
- Olteanu, A.; Blodgett, S. L.; Balayn, A.; Wang, A.; Diaz, F.; Calmon, F. d. P.; Mitchell, M.; Ekstrand, M.; Binns, R.; and Barocas, S. 2025. Rigor in AI: Doing rigorous AI work requires a broader, responsible AI-informed conception of rigor. *arXiv [cs.CY]*.
- Orr, W.; and Kang, E. B. 2024. AI as a sport: On the competitive epistemologies of benchmarking. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: ACM.
- Pelegrina, G. D.; Brotto, R. D. B.; Duarte, L. T.; Attux, R.; and Romano, J. M. T. 2022. Analysis of Trade-offs in Fair Principal Component Analysis Based on Multi-objective Optimization. In *2022 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Peters, M. D. J.; Marnie, C.; Colquhoun, H.; Garrity, C. M.; Hempel, S.; Horsley, T.; Langlois, E. V.; Lillie, E.; O'Brien, K. K.; Tuncalp, O.; Wilson, M. G.; Zarin, W.; and Tricco, A. C. 2021. Scoping reviews: reinforcing and advancing the methodology and application. *Syst. Rev.*, 10(1): 263.
- Poursabzi-Sangdeh, F.; Goldstein, D. G.; Hofman, J. M.; Wortman Vaughan, J. W.; and Wallach, H. 2021. Manipulating and Measuring Model Interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, 1–52. ACM.
- Przybya, P.; and Soto, A. J. 2021. When classification accuracy is not enough: Explaining news credibility assessment. *Information Processing & Management*, 58(5): 102653.
- Qi, Y.; Dong, Y.; Khastgir, S.; Jennings, P. A.; Zhao, X.; and Huang, X. 2023. STPA for Learning-Enabled Systems: A Survey and a New Practice. In *26th IEEE International Conference on Intelligent Transportation Systems (ITSC 2023)*. Accepted by ITSC 2023; also available as arXiv preprint arXiv:2302.10588v2.
- Reuel, A.; Hardy, A.; Smith, C.; Lamparth, M.; Hardy, M.; and Kochenderfer, M. J. 2024. BetterBench: Assessing AI benchmarks, uncovering issues, and establishing best practices. *arXiv [cs.AI]*.
- Richardson, B.; and Gilbert, J. E. 2021. A Framework for Fairness: A Systematic Review of Existing Fair AI Solutions.
- Rismani, S.; Dobbe, R.; and Moon, A. 2024. From Silos to Systems: Process-oriented Hazard Analysis for AI systems. *arXiv [cs.AI]*.
- Rismani, S.; Shelby, R.; Davis, L.; Rostamzadeh, N.; and Moon, A. 2025. Measuring What Matters Connecting AI Ethics Evaluations to System Attributes Hazards and Harms. Dataset archived at Zenodo: <https://doi.org/10.5281/zenodo.16786835>.
- Rismani, S.; Shelby, R.; Smart, A.; Delos Santos, R.; Moon, A.; and Rostamzadeh, N. 2023a. Beyond the ML Model: Applying Safety Engineering Frameworks to Text-to-Image

- Development. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 70–83. New York, NY, USA: Association for Computing Machinery.
- Rismani, S.; Shelby, R.; Smart, A.; Jatho, E.; Kroll, J. A.; Moon, A.; and Rostamzadeh, N. 2023b. From Plane Crashes to Algorithmic Harm: Applicability of Safety Engineering Frameworks for Responsible ML. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, volume 1. Association for Computing Machinery.
- Salaudeen, O.; Reuel, A.; Ahmed, A.; Bedi, S.; Robertson, Z.; Sundar, S.; Domingue, B.; Wang, A.; and Koyejo, S. 2025. Measurement to meaning: A validity-centered framework for AI evaluation. *arXiv [cs.CY]*.
- Santana, M. R. O.; Melo, L. C.; Camargo, F. H. F.; Brandao, B.; Soares, A.; Oliveira, R. M.; and Caetano, S. 2020. MARS-Gym: A Gym framework to model, train, and evaluate Recommender Systems for Marketplaces. In *2020 International Conference on Data Mining Workshops (ICDMW)*, 189–197. IEEE.
- Santurkar, S.; Durmus, E.; Ladhak, F.; Lee, C.; Liang, P.; and Hashimoto, T. 2023. Whose Opinions Do Language Models Reflect? *ArXiv*, abs/2303.17548.
- Savazzi, S.; Rampa, V.; Kianoush, S.; and Bennis, M. 2023. An Energy and Carbon Footprint Analysis of Distributed and Federated Learning. *IEEE Transactions on Green Communications and Networking*, 7(1): 248–264.
- Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19, 59–68. New York, NY, USA: ACM.
- Shelby, R.; Rismani, S.; Henne, K.; Moon, A.; Rostamzadeh, N.; Nicholas, P.; Yilla-Akbari, N.; Gallegos, J.; Smart, A.; Garcia, E.; and Virk, G. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 723–741. New York, NY, USA: Association for Computing Machinery.
- Shen, H.; Li, J.; Wu, G.; and Zhang, M. 2023. Data release for machine learning via correlated differential privacy. *Information Processing & Management*, 60(3): 103349.
- Sloane, M.; and Moss, E. 2023. Assessing the assessment: Comparing algorithmic impact assessments and AI audits. *SSRN Electron. J.*
- Solaiman, I.; Talat, Z.; Agnew, W.; Ahmad, L.; Baker, D.; Blodgett, S. L.; Daumé, H., III; Dodge, J.; Evans, E.; Hooker, S.; Jernite, Y.; Luccioni, A. S.; Lusoli, A.; Mitchell, M.; Newman, J.; Png, M.-T.; Strait, A.; and Vassilev, A. 2023. Evaluating the Social Impact of Generative AI Systems in Systems and Society. *arXiv [cs.CY]*.
- Sánchez, D. 2010. A methodology to learn ontological attributes from the Web. *Data & Knowledge Engineering*, 69(6): 573–597.
- Tolan, S.; Miron, M.; Gómez, E.; and Castillo, C. 2019. Why Machine Learning May Lead to Unfairness: Evidence from Risk Assessment for Juvenile Justice in Catalonia. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, ICAIL '19. ACM.
- Vera Liao, Q.; and Xiao, Z. 2023. Rethinking Model Evaluation as Narrowing the Socio-Technical Gap. *arXiv [cs.HC]*.
- Vereschak, O.; Bailly, G.; and Caramiaux, B. 2021. How to evaluate trust in AI-assisted decision making? A survey of empirical methodologies. *Proc. ACM Hum. Comput. Interact.*, 5(CSCW2): 1–39.
- Wallach, H.; Desai, M.; Pangakis, N.; Cooper, A. F.; Wang, A.; Barocas, S.; Chouldechova, A.; Atalla, C.; Blodgett, S. L.; Corvi, E.; Dow, P. A.; Garcia-Gathright, J.; Olteanu, A.; Reed, S.; Sheng, E.; Vann, D.; Vaughan, J. W.; Vogel, M.; Washington, H.; and Jacobs, A. Z. 2024. Evaluating generative AI systems is a social science measurement challenge. *arXiv [cs.CY]*.
- Wang, A.; Narayanan, A.; and Russakovsky, O. 2020. *RE-VERSE: A Tool for Measuring and Mitigating Bias in Visual Datasets*, 733–751. Springer International Publishing. ISBN 9783030585808.
- Wang, X.; and Yin, M. 2022. Effects of Explanations in AI-Assisted Decision Making: Principles and Comparisons. *ACM Transactions on Interactive Intelligent Systems*, 12(4): 1–36.
- Watkins, E. A.; Moss, E.; Metcalf, J.; Singh, R.; and Elish, M. C. 2021. Governing Algorithmic Systems with Impact Assessments: Six Observations. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 1010–1022. New York, NY, USA: Association for Computing Machinery.
- Weidinger, L.; Rauh, M.; Marchal, N.; Manzini, A.; Hendricks, L. A.; Mateos-Garcia, J.; Bergman, S.; Kay, J.; Griffin, C.; Bariach, B.; Gabriel, I.; Rieser, V.; and Isaac, W. 2023. Sociotechnical Safety Evaluation of Generative AI Systems. *arXiv [cs.AI]*.
- Wu, Y.; Zhang, L.; and Wu, X. 2018. On Discrimination Discovery and Removal in Ranked Data using Causal Graph. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, 2536–2544. ACM.
- Xiao, T.; and Khisti, A. 2019. Maximal Information Leakage based Privacy Preserving Data Disclosure Mechanisms. In *2019 16th Canadian Workshop on Information Theory (CWIT)*, 1–6. IEEE.
- Xiao, Z.; Zhang, S.; Lai, V.; and Vera Liao, Q. 2023. Evaluating Evaluation Metrics: A Framework for Analyzing NLG Evaluation Metrics using Measurement Theory. *arXiv [cs.CL]*.
- Yang, T.-J.; Chen, Y.-H.; Emer, J.; and Sze, V. 2017. A method to estimate the energy consumption of deep neural networks. In *2017 51st Asilomar Conference on Signals, Systems, and Computers*. IEEE.
- Yuan, H.; Yu, H.; Gui, S.; and Ji, S. 2022. Explainability in Graph Neural Networks: A Taxonomic Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–19.

Zhang, G.; Chong, L.; Kotovsky, K.; and Cagan, J. 2023. Trust in an AI versus a Human teammate: The effects of teammate identity and performance on Human-AI cooperation. *Computers in Human Behavior*, 139: 107536.

Zhang, R.; and Venkatasubramanian, P. 2018. Mutual-Information-Private Online Gradient Descent Algorithm. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2077–2081. IEEE.

Zhao, B. Z. H.; Kaafar, M. A.; and Kourtellis, N. 2020. Not One but Many Tradeoffs: Privacy Vs. Utility in Differentially Private Machine Learning. In *Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop (CCSW 2020)*, 15–26. Association for Computing Machinery. Virtual event, November 9, 2020.