

# Understanding and Meeting Practitioner Needs When Measuring Representational Harms Caused by LLM-Based Systems

Emma Harvey\*<sup>◇</sup> Emily Sheng<sup>▣</sup> Su Lin Blodgett<sup>▣</sup> Alexandra Chouldechova<sup>▣</sup>  
Jean Garcia-Gathright<sup>▣</sup> Alexandra Olteanu<sup>▣</sup> Hanna Wallach<sup>▣</sup>

<sup>◇</sup>Cornell University <sup>▣</sup>Microsoft Research

evh29@cornell.edu

{emilysheng, sulin.blodgett, alexandrac,

jeang, alexandra.olteanu, wallach}@microsoft.com

## Abstract

The NLP research community has made publicly available numerous instruments for measuring representational harms caused by large language model (LLM)-based systems. These instruments have taken the form of datasets, metrics, tools, and more. In this paper, we examine the extent to which such instruments meet the needs of practitioners tasked with evaluating LLM-based systems. Via semi-structured interviews with 12 such practitioners, we find that practitioners are often unable to use publicly available instruments for measuring representational harms. We identify two types of challenges. In some cases, instruments are *not useful* because they do not meaningfully measure what practitioners seek to measure or are otherwise misaligned with practitioner needs. In other cases, instruments—even useful instruments—are *not used* by practitioners due to practical and institutional barriers impeding their uptake. Drawing on measurement theory and pragmatic measurement, we provide recommendations for addressing these challenges to better meet practitioner needs.

## 1 Introduction

Representational harms (Barocas et al., 2017; Crawford, 2017) occur when a system “represents some social groups in a less favorable light than others, demeans them, or fails to recognize their existence altogether” (Blodgett et al., 2020). Numerous studies have documented representational harms caused by large language model (LLM)-based systems (e.g., Sheng et al., 2019; Dev et al., 2021; Venkit et al., 2022; Kotek et al., 2023; Hofmann et al., 2024). It is important to measure and mitigate such harms—especially for systems that will be deployed in real-world contexts. However, this is known to be a challenging task. Like many other concepts related to the capabilities,

behaviors, and impacts of LLM-based systems, representational harms are abstract and can have contested meanings across use cases, languages, and cultures (Wallach et al., 2025). As a result, they are particularly difficult to define precisely and thus measure (Blodgett et al., 2020; Dev et al., 2022; Katzman et al., 2023; Wang et al., 2023).

To facilitate measuring representational harms, the NLP research community has produced numerous publicly available<sup>1</sup> measurement instruments, including datasets, metrics, tools, benchmarks,<sup>2</sup> and annotation instructions. In this paper, we investigate whether these instruments meet the needs of practitioners tasked with evaluating LLM-based systems. As studies examining the uptake of other responsible AI artifacts have found, practitioner needs as assumed in the research literature are often different from those actually voiced by practitioners (Holstein et al., 2019; Lee and Singh, 2021; Deng et al., 2022; Ojewale et al., 2024). This potential mismatch presents a critical opportunity for researchers and practitioners to engage with one another. If measurement instruments do not meet the needs of practitioners tasked with evaluating LLM-based systems, *those instruments will not be used by such practitioners*—causing developers and deployers of LLM-based systems either to not engage in measurement or to rely on bespoke instruments that might not be publicly understood.

Through a series of semi-structured interviews with 12 practitioners tasked with evaluating LLM-based systems for representational harms, we find that practitioners are often unable to use publicly available measurement instruments, despite a desire to do so. We identify two types of challenges that lead to this. In some cases, instruments are *not useful*: they do not meaningfully measure what practitioners seek to measure or are otherwise

\*Work conducted during a Microsoft Research internship.

<sup>1</sup>By *publicly available*, we mean available on the internet or via an academic publication for others to use or adapt.

<sup>2</sup>*Benchmarks* consist of datasets and metrics.

Instrument	Examples
Datasets	Fifty Shades of Bias (Hada et al., 2023), FairPrism (Fleisig et al., 2023), ToxiGen (Hartvigsen et al., 2022)
Metrics	WEAT (Caliskan et al., 2017), SEAT (May et al., 2019), $\alpha$ -Intersectional Fairness (Maheshwari et al., 2023)
Tools	Perspective API (Lees et al., 2022), Llama Guard (Inan et al., 2023), HateBERT (Caselli et al., 2021)
Benchmarks	BOLD (Dhamala et al., 2021), BBQ (Parrish et al., 2022), StereoSet (Nadeem et al., 2021)
Annotation instructions	Included as part of instruments like datasets (e.g., Fleisig et al., 2023) and benchmarks (e.g., Nadeem et al., 2021), or released as part of measurement frameworks (e.g., Magooda et al., 2023)
Other	Matched guise probing (Hofmann et al., 2024), DivDist (Bommasani and Liang, 2024)

Table 1: Examples of publicly available instruments for measuring representational harms.

misaligned with practitioner needs. In other cases, instruments—even useful instruments—are *not used* by practitioners due to practical or institutional barriers impeding their uptake. Although we focus on instruments for measuring representational harms, many of our findings apply broadly to cases where practitioners seek to use publicly available instruments to measure other abstract or contested concepts. However, our targeted focus allows us to identify cases where challenges are exacerbated by the specifics of measuring representational harms or evaluating LLM-based systems.

Developing measurement instruments that are simultaneously useful and used is not a new challenge. Measurement theory from the social sciences has long been concerned with designing instruments that are useful, i.e., those that meaningfully measure what they purport to measure (Adcock and Collier, 2001; Jacobs and Wallach, 2021; Wallach et al., 2025). Pragmatic measurement builds on measurement theory to focus on designing measurement instruments that are both useful and used in practice, i.e., that are aligned with practitioner needs and designed to overcome barriers impeding their uptake (Glasgow and Riley, 2013). Drawing on work from measurement theory and pragmatic measurement, we identify opportunities to improve measurement instruments and their uptake among practitioners.

## 2 Related Work

The NLP research community has made publicly available numerous instruments for measuring representational harms caused by LLM-based systems. We provide examples of such instruments in Table 1,<sup>3</sup> and point the reader to recent surveys by Sheng et al. (2021), Dev et al. (2022), and Gallegos et al. (2024) for more complete overviews.

Assessments of existing measurement instruments have identified potential limitations to their usefulness. Prior work has shown that the concepts such instruments seek to measure are often poorly motivated, unclear, or not meaningfully measured by those instruments (Blodgett et al., 2020, 2021; Goldfarb-Tarrant et al., 2023; Xiao et al., 2023; Delobelle et al., 2024; Porada et al., 2024; Zhao et al., 2024). Researchers have also found that instruments can be highly sensitive to implementation choices that should not affect measurement outcomes (Antoniak and Mimno, 2021; Delobelle et al., 2022; Seshadri et al., 2022; Sclar et al., 2023; Shu et al., 2024). Furthermore, instruments often fail to produce measurements that enable informed actions (Delobelle et al., 2024), or that are useful for measuring representational harms in real-world deployment contexts (Goldfarb-Tarrant et al., 2021; Cao et al., 2022; Delobelle et al., 2022).

In this paper, we explore the extent to which these assessments of existing measurement instruments reflect the needs of practitioners, which are often both implicit and impacted by practical constraints (Zhou et al., 2022). We build on prior work from HCI showing that practitioners often struggle to use artifacts developed by researchers, in part because these artifacts are misaligned with practitioner needs (Holstein et al., 2019; Lee and Singh, 2021; Richardson et al., 2021; Deng et al., 2022; Balayn et al., 2023; Ojewale et al., 2024). This work focused on the needs of practitioners who seek to measure allocative harms caused by predictive models. To our knowledge, we are the first to explore whether and to what extent publicly available measurement instruments meet the real-world needs of practitioners who seek to measure representational harms caused by LLM-based systems.

<sup>3</sup>Table 1 is not exhaustive, nor is it intended to suggest that we are specifically critiquing the instruments that are listed.

PID	Role	Employer
P01	Research engineer	Big tech company
P02	Applied scientist	Big tech company
P03	Research scientist	Big tech company
P04	Research engineer	Big tech company
P05	Consultant	AI startup
P06	Scientist	Big tech company
P07	Research scientist	AI startup
P08	Data engineer	Big non-tech company
P09	NLP specialist	Big non-tech company
P10	Research scientist	AI startup
P11	Researcher	AI nonprofit
P12	Researcher	Big tech company

Table 2: Participant details.

### 3 Methods

To understand the extent to which publicly available measurement instruments meet the needs of practitioners, we conducted a series of semi-structured interviews with 12 practitioners tasked with evaluating LLM-based systems. Participants, whom we refer to throughout by IDs P1–P12, worked on LLM-based systems (e.g., search engines, chatbots) and content moderation tools for such systems. See Table 2 for participant details.<sup>4</sup>

**Recruitment.** We recruited participants through our professional networks, social media, cold emails, and snowball sampling (Morgan, 2008). Each interview was one hour long and conducted between June and August 2024. All participants provided informed consent prior to their interviews. Each participant received a \$75 gift card. The study was approved by Microsoft’s research IRB.

**Interviews.** To scaffold the interviews, we identified a set of desiderata for measurement instruments: validity, reliability, specificity, extensibility, scalability, interpretability, and actionability (see Table 3). These desiderata were identified based on our own experiences measuring representational harms caused by LLM-based systems and a systematic review of the NLP literature on assessing measurement instruments (see Appendix A).

We began each interview by asking the participant to describe their role and the LLM-based systems they worked on. Next, we asked them to walk us through an example of how they measured representational harms, noting the publicly available measurement instruments they used or considered using. We then asked them to reflect on their

<sup>4</sup>Participants were based in North America and Europe. They self-identified as white, Asian, and Middle Eastern/North African; and as men, women, and non-binary. To preserve anonymity, we do not report individual demographics.

Desideratum	Definition
Validity	Meaningfully measures what stakeholders think it measures
Reliability	Results in similar measurements when used in similar ways, especially over time
Specificity	Sufficiently specific to a system, its use cases, and its deployment contexts
Extensibility	Can be adapted for different systems, use cases, and deployment contexts
Scalability	Can scale to increasing workloads
Interpretability	Produces measurements that can be understood by stakeholders
Actionability	Produces measurements that can be acted upon by stakeholders

Table 3: Desiderata for measurement instruments. Measurement instruments that fail to meet these desiderata may be challenging to use. We used these desiderata to scaffold the interviews (see the guide in Appendix B).

experiences with those instruments, discussing any challenges they faced. We prompted them about whether they faced any challenges related to instruments failing to meet any of the desiderata listed in Table 3 and also asked open-ended questions about any other challenges they faced. Our semi-structured interview guide is in Appendix B.

We conducted interviews until saturation, i.e., until multiple consecutive interviews did not uncover any new challenges (Small, 2009). Nevertheless, we note that our recruitment efforts yielded a low response rate, which likely resulted in a skewed participant pool (we discuss this in our limitations section). Therefore, we are careful not to over-generalize our findings: our interviews were specifically intended to answer the question, “what challenges do practitioners face when trying to use publicly available measurement instruments?” but not questions like, “what is the prevalence of these challenges among all practitioners?”

**Thematic analysis.** We conducted a thematic analysis using an inductive–deductive coding approach (Braun and Clarke, 2006, 2019). We coded the interview transcripts for challenges mentioned by participants. Initially, the set of challenges that we focused on corresponded to the desiderata listed in Table 3 (i.e., challenges that arose because an instrument failed to meet a given desideratum). The first author coded each transcript and, at that time, identified additional challenges raised by participants that were not covered by our original set of codes. The first and second authors discussed these

uncategorized challenges to create an expanded set of codes, and the first author reread and re-coded each transcript. All authors discussed and synthesized the codes into themes based on how, why, and to what extent the associated challenges impacted practitioners’ abilities to use measurement instruments. Finally, at least one author other than the first author re-coded each transcript. Any coding disagreements identified through that process were resolved through discussion with all of the authors.

**Positionality.** We are a group of researchers and practitioners with expertise in the domains of NLP, machine learning, statistics, computational social science, software engineering, and responsible AI. Collectively, we are familiar with the NLP research community as well as the needs of practitioners working on NLP products and services. Many of us are employed in industry positions in which we have been tasked with measuring representational harms caused by LLM-based systems and have personally faced some of the same challenges raised by participants. Our professional experiences measuring representational harms had an impact on our research, in particular by influencing the set of desiderata that we used to scaffold the interviews.

## 4 Practitioner Challenges

Participants reported being aware of a range of publicly available measurement instruments, including tools designed to identify unsafe or toxic text (e.g., DeBERTa (He et al., 2021) and Llama Guard (Inan et al., 2023)), as well as datasets and benchmarks (e.g., StereoSet (Nadeem et al., 2021) and CrowS-Pairs (Nangia et al., 2020)). However, these measurement instruments were used by at most one or two participants—none were more widely used. More importantly, all participants discussed facing challenges that prevented them from using publicly available measurement instruments.

Specifically, participants experienced two types of challenges that left them unable to use publicly available measurement instruments, even when they had a desire to do so. First, instruments are *not useful* when they do not meaningfully measure what practitioners seek to measure or are otherwise misaligned with practitioner needs (§4.1). Second, instruments are *not used* when practitioners face barriers impeding their uptake—even if the instruments are useful (§4.2). These challenges are context-dependent and determined by practitioner needs; the same instrument may be useful for

one practitioner in one context but not for another.

We found that the desiderata identified in Table 3 aligned closely with considerations participants described as being central to their decisions about the usefulness of measurement instruments. Validity and specificity were primary considerations—all participants reported that they considered these desiderata and chose not to use measurement instruments that did not meet them. Participants also considered interpretability and actionability, but typically not until after they had already deemed measurement instruments sufficiently valid and specific. Although participants reported being concerned about the reliability and scalability of measurement instruments, most did not consider these desiderata when deciding whether to use an instrument. Participants did not report considering any additional desiderata.<sup>5</sup> However, participants identified additional challenges to using measurement instruments in the form of practical and institutional barriers impeding their uptake. Finally, participants also identified cases where challenges were exacerbated by the specifics of measuring representational harms or evaluating LLM-based systems (§4.3).

### 4.1 Measurement Instruments Are Not Useful

*“Does it result in valid measurements?... Is [it] going to translate well to my scenario?” – P2*

Every participant reported at least one experience in which they chose not to use a measurement instrument due to concerns related to one or more desiderata. In fact, multiple participants reported that they were unable to use any publicly available measurement instruments because of challenges related to their usefulness (P1, P2, P6, P8, P9).

#### **Practitioners cannot use measurement instruments that they perceive as lacking validity.**

Many participants reported feeling unable to use measurement instruments because the concepts that those instruments were intended to measure were not clearly defined or linked to any existing theoretical understandings of those concepts (P3, P6–P8, P10–P12). This makes it very difficult to pose the question of “how well” those concepts are being measured. Even when concepts were clearly defined, participants reported that instruments

<sup>5</sup>We do not discuss extensibility here because participants did not report explicitly considering extensibility when deciding whether to use measurement instruments. Instead, we discuss it in §5 as a desideratum that can be used to improve the usefulness and uptake of measurement instruments.



sometimes failed to meaningfully measure them (“There is a Jigsaw dataset that is supposed to be used to measure gender biases in hate speech detection systems. But when I [looked] into the data, the samples were often not about gender biases at all. They just [included gender] keywords. . . I concluded that it’s simply not up to the task.” – P11, also P3–P6, P12). As another example, participants reported frequently encountering datasets that contain mislabeled instances (“Every single public benchmark we use. . . has a couple of rows that we look at by eye, and we’re like, ‘that doesn’t make sense.’ And then that makes us question the entire validity of the benchmark.” – P7, also P3–P6, P11). These potential threats to validity caused participants to lose trust in measurement instruments.

**Data contamination exacerbates validity concerns.** Data contamination was an overarching concern about the validity of publicly available datasets and benchmarks. Because developers of LLMs seldom disclose their training data, it is difficult to tell whether an LLM-based system that performs well on a benchmark has simply been trained using the benchmark data. Half of the participants expressed discomfort with using publicly available benchmarks and datasets under any circumstances (P4–P6, P9, P10, P12).

**Practitioners cannot use instruments that lack specificity for their needs.** All participants reported choosing not to use publicly available measurement instruments because they were not sufficiently specific to their needs. Many reported creating their own instruments from proprietary data as a result (“we had to develop tests that were more suited to the sorts of scenarios that would happen in the workplace, not just a chat conversation that would happen outside of the workplace.” – P2, also P1, P3, P5, P6, P8–P12).

**The contextual nature of representational harms exacerbates specificity concerns.** Participants reported that many publicly available measurement instruments are not sufficiently specific to the representational harms they sought to measure. They attributed this to the fact that measuring representational harms requires cultural context; i.e., a specific understanding of who may be harmed and how (“it’s tough to come in and say, here’s this dataset with a bunch of stereotypes about race. Hopefully all of these stereotypes are gonna be present in this very specific system that we’re working on.” – P3, also P5, P6, P8, P9, P11).

**Practitioners struggle to use measurement instruments that lack interpretability.** Several participants reported experiencing concerns about whether instruments produce measurements that can be understood (P1, P2, P5, P9, P11, P12). In particular, participants felt that measurements produced by tools and benchmarks could not be interpreted without additional information. As P11 put it, after using a measurement instrument, “you end up with a number” and then must decide “when will this number become problematic?” Without more information about what measurements mean (e.g., comparisons to other measurements from the same instrument), it is challenging to understand them.

**Practitioners deprioritize measurement instruments that they perceive as lacking actionability.** Some participants reported experiencing concerns about whether instruments produce measurements that can be acted upon (P2, P3, P5–P7, P10). These concerns often stemmed from other issues, e.g., because measurement instruments are not sufficiently valid (P7) or interpretable (P10) for stakeholders to confidently act upon their outputs. Other participants reported that if they could not pre-identify a clear strategy for mitigating a harm, they often deprioritized measuring it (“You only have so many hours in a week. . . the [harms] that don’t have a clear mitigation strategy. . . they’re unlikely to be useful [to measure].” – P3, also P2, P5, P6, P10).

**While reliability is desirable in theory, it is often not considered in practice.** Participants did not report observing reliability issues in publicly available measurement instruments and then choosing not to use those instruments as a result—despite the importance of reliability to measurement (Jacobs and Wallach, 2021) and the well-documented lack of reliability exhibited by measurement instruments in the NLP literature (e.g. Sclar et al., 2023; Shu et al., 2024; Delobelle et al., 2024). We hypothesize that this mismatch may be due to the fact that validity and specificity were primary considerations for the participants we interviewed, followed by interpretability and actionability. If instruments failed to meet these desiderata, participants chose not to use them—never reaching the point of considering whether the instruments were sufficiently reliable.

**Scalability concerns do not typically prevent practitioners from using measurement instruments.** Most commonly, participants reported scalability challenges related to measurement instruments that required repeated calls to LLM-

based systems (e.g., instruments that rely on LLMs as judges, or instruments that require multiple responses from LLM-based systems to produce measurements) (P1–P3, P11). Although participants noted that repeated calls to LLM-based systems added time, costs, and environmental impacts to their measurement processes, they considered this an inevitable result of evaluating LLM-based systems. However, when participants sought to do online measurement for client-facing systems, they sometimes chose not to use such instruments due to the latency or token limits imposed by those instruments (P7, P10). We hypothesize that when practitioners exclusively do offline measurement, scalability of the instruments themselves is less of a concern; however, when online measurement is the goal, scalability concerns become more salient.

## 4.2 Measurement Instruments Are Not Used

*“[I]t can just be really hard to actually get the time allotted to go and find the resources that probably exist. So usually we end up making our own thing.” – P3*

Even when participants thought that publicly available measurement instruments might be useful, they were sometimes unable to use them due to practical and institutional barriers. These barriers arose not because the instruments failed to meet particular desiderata, but because of the challenges specific to measuring representational harms in practice. We highlight these barriers not to suggest that designers of measurement instruments are solely responsible for addressing them, but because it is important to identify barriers impeding the uptake of measurement instruments to provide a more holistic view of what can be changed and who can make those changes.

### **Practitioners face practical barriers to using measurement instruments when those instruments do not meet organizational requirements.**

Participants reported being unable to use publicly available instruments due to challenges related to organizational requirements, such as security (“We have this agreement not to disclose [our customers’ data] outside. With these publicly available tools, especially with the ones that are not open source, we have always these security issues.” – P9) and data licensing (“We have to use compliant datasets. In some case, we just created our own [version] of a dataset that we cannot use.” – P4, also P6).

### **Practitioners face institutional barriers when organizational culture impedes uptake of measurement instruments.**

Participants reported facing a lack of organizational support—and sometimes outright disincentives—for using publicly available instruments, regardless of whether those instruments might be useful. For example, participants reported that they sometimes created new measurement instruments due to a lack of time to find existing publicly available instruments (P3, P7). Some participants were not able to use measurement instruments if those instruments did not align with their organizations’ processes. For example, P10 reported that their organization had a preference for measurement processes that were similar to those used in software engineering, i.e., having a small set of curated test cases that a system must pass prior to deployment. It can be unclear how to translate large benchmarks into smaller sets of prioritized test cases, e.g., “the 30 cases that should not fail if we change anything” (P10). Finally, participants reported facing limited incentives to measure representational harms, especially compared to other kinds of harms (e.g., quality-of-service harms) (P5, P7, P8, P10).

## 4.3 Challenges Are Exacerbated by, but Extend Beyond, Representational Harms

Participants identified properties of representational harms that made them particularly challenging to measure. These included the fact that measuring representational harms requires additional information or expertise, such as cultural context or social science expertise (P2, P11). Additionally, some participants felt that the contestedness of representational harms made it more difficult to assess the validity or reliability of instruments intended to measure them (P4, P5, P9, P10, P12).

Although our interviews focused specifically on instruments for measuring representational harms, participants consistently voiced that their challenges applied broadly to instruments for measuring other abstract or contested concepts.<sup>6</sup> For example, participants also reported validity concerns about instruments for measuring disinformation (P11); specificity concerns about instruments for measuring the legality of text (because legal codes are specific to geographic jurisdictions) (P10); and interpretability and actionability con-

<sup>6</sup>These challenges did not apply to instruments intended to measure directly observable concepts, such as whether an LLM-based system generates phone numbers (P7).

cerns about machine translation benchmarks (P5).

Finally, we note that multiple participants lived in countries where English is not the primary language, but nevertheless focused on evaluating LLM-based systems in English. Although we did not specifically recruit practitioners who were multilingual or focused on low-resource languages, a third of the participants discussed speaking languages other than English and challenges related to measuring representational harms in those languages (P5, P9, P11, P12). P11, who speaks a low-resource European language, shared that they had considered trying to develop a measurement instrument in that language, but opted against it due to the low probability that their instrument would be widely adopted. We therefore hypothesize that practitioners who seek to measure representational harms in low-resource languages likely face additional challenges related to the availability of measurement instruments in those languages.

## 5 Addressing Practitioner Challenges

In this section, we draw on measurement theory from the social sciences and pragmatic measurement to improve the usefulness and uptake of measurement instruments. As we mentioned in §1, measurement theory has long been concerned with designing instruments that are useful, while pragmatic measurement builds on measurement theory to focus on designing measurement instruments that are both useful and used in practice. Designers of measurement instruments can draw on measurement theory to improve the validity and reliability of their instruments, and can draw on pragmatic measurement to improve other aspects of the usefulness and uptake of their instruments. This is not to suggest that designers of measurement instruments are solely responsible for meeting practitioner needs. Rather, practitioners who seek to use measurement instruments are responsible for adapting those instruments to meet their specific needs. However, designers can facilitate this by designing extensible measurement instruments, which can help address usefulness and uptake challenges, as we explain below. Regulators and organizations that develop and deploy LLM-based systems can also play a role in removing both practical and institutional barriers. Finally, we briefly discuss trade-offs inherent to attempting to meet multiple desiderata simultaneously.

**Measurement theory provides a framework with which to improve the usefulness of measurement instruments.** All participants reported concerns about whether measurement instruments meaningfully measure the concepts that they are intended to measure—often because those concepts were not clearly defined. These concerns were exacerbated by the fact that representational harms are abstract and contested, meaning that different measurement instruments may operationalize different definitions of representational harms. Measurement theory offers a framework with which these concerns can be better understood and addressed.

Measurement theory provides a framework for obtaining measurements (e.g., scores calculated using a benchmark) of abstract concepts (e.g., “stereotypes”) through the processes of systematization, operationalization, application, and interrogation (Adcock and Collier, 2001; Wallach et al., 2025). *Systematization* is the process of formulating a specific, often theoretically grounded, definition of the concept of interest. This definition—the *systematized concept*—is, as noted by multiple participants and in prior work, often absent from instruments for measuring representational harms. Instead, designers of measurement instruments often jump straight to *operationalization*, which is the process of developing one or more instruments for measuring the concept of interest.<sup>7</sup>

**Designers of measurement instruments should draw on measurement theory to ensure that their instruments are valid and reliable.** Jumping straight to operationalization not only creates uncertainty concerning what precisely measurement instruments are intended to measure, but also renders it impossible to pose the question of “how well” a concept of interest is being measured when that concept has multiple competing meanings. For example, stereotypes can be negative, neutral, or positive in sentiment. An instrument that performs well at measuring stereotypes with negative sentiment may perform poorly if its validity is interrogated with respect to a broader definition of stereotyping that encompasses other sentiments. We therefore recommend that designers of measurement instruments do not skip systematization (or conflate it with operationalization) and clearly

<sup>7</sup>We note that *application* is the process of using the resulting measurement instruments to obtain measurements of the concept of interest, while *interrogation* is the process of interrogating the validity of the systematized concept, the measurement instruments, and their resulting measurements.

document systematized concepts as part of making measurement instruments publicly available.

Once a concept has been systematized, questions of how accurately the systematized concept is being measured can be answered using different lenses of validity and reliability (Jacobs and Wallach, 2021). Because participants reported perceived threats to the validity of measurement instruments, we recommend that designers of measurement instruments use these lenses to rigorously interrogate the validity and reliability of their instruments, providing evidence from these interrogations when making their instruments publicly available (see Xiao et al., 2023; Van Der Wal et al., 2024). Lastly, we recommend that designers of measurement instruments release resources (e.g., guidelines, code) that practitioners can use to interrogate validity and reliability in different contexts.

**Pragmatic measurement suggests additional ways to improve the usefulness of measurement instruments.** All participants struggled to use measurement instruments that were misaligned with their needs. We argue that drawing on pragmatic measurement, which builds on measurement theory and has historically been concerned with improving the real-world effectiveness of clinical research, can help designers of measurement instruments to improve aspects of the usefulness of their instruments beyond validity and reliability (Glasgow and Riley, 2013).

**Designers of measurement instruments should draw on pragmatic measurement to ensure that their instruments are interpretable, actionable, and scalable in different contexts.** Pragmatic measurement suggests ways to improve the interpretability, actionability, and scalability of measurement instruments, all of which contribute to their usefulness. Many participants reported struggling to use measurement instruments that lack interpretability. Pragmatic measurement offers several ways to improve the interpretability of measurement instruments. For example, where possible, designers of measurement instruments should publish distributions of measurements produced by their instruments for known datasets (Lewis et al., 2021).<sup>8</sup> Additionally, designers of measurement instruments should publish information about how to interpret the measurements produced by their

instruments (e.g., what does a measurement of 0.3 mean vs. a measurement of 0.8?) (Stanick et al., 2021). These suggestions are simple, evidence-based ways to improve interpretability. Furthermore, they are complementary to the recommendations for improving actionability recently proposed by Delobelle et al. (2024). See Estabrooks et al. (2012), Martinez et al. (2014), and Lewis and Dorsey (2020) for other ways to further improve the usefulness of measurement instruments.

**Practitioners are responsible for adapting measurement instruments to meet their specific needs—but designers of measurement instruments can help by ensuring that their instruments are extensible.** Pragmatic measurement emphasizes the importance of *extensibility*, i.e., whether an instrument can be adapted for different systems, use cases, and deployment contexts.<sup>9</sup> In pragmatic measurement, extensibility is a low-effort, high-impact way to make measurement instruments useful: rather than designing multiple measurement instruments to meet different needs, designers can focus on extensible instruments that practitioners can individually tailor (Powell et al., 2015; Waltz et al., 2015). Based on our findings in §4 and the pragmatic measurement literature, we identify key criteria that need to be met for measurement instruments to be considered extensible. First, we recommend making them open source or otherwise modifiable. Second, we recommend that they be *modular*, i.e., composed of discrete, interconnected pieces (Baldwin and Clark, 2000).

**Designers of measurement instruments and practitioners should both ensure that measurement instruments remain valid and reliable when adapted to new contexts.** A measurement instrument may be sufficiently valid and reliable in one context, but not in another. To ensure that extensible measurement instruments remain valid and reliable, we recommend that designers of measurement instruments be explicit about what is being measured (the systematized concept) and how it is being measured. Separating systematization from operationalization can enable measurement instruments to be more easily modified. Designers of measurement instruments should also be explicit about the impacts on validity and reliability of modifying different aspects of their instruments when adapting them to new contexts (Powell et al., 2015; Waltz et al., 2015). We also emphasize that

<sup>8</sup>This is different from the practice of publishing the accuracy of measurement instruments on known datasets, which does not help practitioners interpret individual measurements.

<sup>9</sup>Extensibility is sometimes also called *adaptability*.



any time a measurement instrument is to be used in a new context, its validity and reliability must be re-interrogated (Martinez et al., 2014). To facilitate this, we reiterate that designers of measurement instruments should release resources (e.g., guidelines, code) that practitioners can use to interrogate validity and reliability in different contexts. Designers should also make it easy for practitioners to compare measurements produced by instruments with and without adaptations (Martinez et al., 2014).

**Extensibility can help overcome practical barriers impeding the uptake of measurement instruments.** Participants noted practical barriers to using measurement instruments, including security and data licensing. Extensible measurement instruments that can be adapted to support local use may help mitigate some security concerns. It should also be possible to adapt extensible measurement instruments to incorporate newly obtained, proprietary, or other sources of data in order to overcome data licensing issues.

**Trade-offs.** Designing measurement instruments that are simultaneously valid, reliable, specific, extensible, scalable, interpretable, and actionable and are actually used by practitioners is a tall order. Trade-offs are inevitable: for example, as measurement instruments become more specific to a particular context, they are likely to become less extensible. We emphasize the importance of interrogating validity and reliability, and, beyond that, we caution against over-indexing on any particular desideratum. For example, designers of measurement instruments who optimize for scalability without considering other desiderata may face unwanted trade-offs. By analogy, blocklists, which are commonly used to mitigate representational harms, are relatively low effort to deploy and scale, but do not wholly and fully capture the contextual nature of representational harms. As a result, they produce false positives when social groups engage in actions like reclaiming slurs, which can lead to over-moderation or erasure of those groups (Vashishtha et al., 2023).

## 6 Conclusion

We found that practitioners are often unable to use publicly available instruments for measuring representational harms and identified two types of challenges. In some cases, instruments are *not useful* because they do not meaningfully measure what practitioners seek to measure or are otherwise mis-

aligned with practitioner needs. In other cases, instruments—even useful instruments—are *not used* by practitioners due to practical and institutional barriers impeding their uptake. Furthermore, both types of challenges can be exacerbated by the specifics of measuring representational harms or evaluating LLM-based systems. Drawing on measurement theory and pragmatic measurement, we provided recommendations for addressing these challenges to better meet practitioner needs.

## Limitations

The primary limitation of our paper is that our recruitment efforts yielded a low response rate. As is the case with many studies targeting technology workers (Scheuerman, 2024), it was challenging to identify and recruit potential participants. For example, in addition to relying on our professional networks and social media, we cold-emailed 73 practitioners whom we identified as being potential participants—based on LinkedIn profiles, company websites, and publications at ACL venues—and whose contact information was available online. We were ultimately able to interview one such practitioner. Potential participants often declined to speak with us due to NDAs or other confidentiality concerns. We interviewed a total of 12 practitioners, some of whom declined to answer certain questions in order to remain in compliance with their organizations’ NDAs. Due to these recruitment challenges, it is likely that our participant pool is not broadly representative of all practitioners. For example, it is likely skewed toward practitioners who faced challenges when trying to use publicly available instruments for measuring representational harms caused by LLM-based systems. Therefore, we are careful not to over-generalize our findings. For example, our findings do not enable us to answer questions about the prevalence of the challenges that we identified. We believe there is potential to expand on our findings in future work by surveying practitioners to answer such questions.

Because qualitative research of this type is not typical in ACL venues, we clarify that the above limitation is not the same as having a too-small sample size. Our sample size is, in fact, typical of HCI interview studies (Caine, 2016). Indeed, we conducted interviews until saturation, i.e., until multiple consecutive interviews did not uncover any new challenges (Small, 2009; Hennink and Kaiser, 2022). We emphasize that in interview

studies like ours, the goal is not to interview a target number of participants. Rather, the process of establishing saturation determines the sample size and whether it is appropriate (see [Small, 2009](#)).

Finally, because English is the only language shared by all members of the research team, we were only able to conduct interviews in English. Our findings are therefore centered on practitioners who focused (either primarily or exclusively) on evaluating LLM-based systems in English, although some participants did touch on low-resource languages in §4. Nevertheless, we hypothesize that practitioners who seek to measure representational harms in low-resource languages likely face additional challenges ([Bender, 2019](#)).

## Ethical Considerations

We do not anticipate any risks to society or the general public associated with our findings as we simply identified practitioners’ challenges to using publicly available instruments for measuring representational harms caused by LLM-based systems.

Our research involved interviewing humans.<sup>10</sup> As a result, there is an inherent risk to our participants: the possibility of identification, as we did collect personally identifiable information (in order to obtain informed consent and record interviews). We have taken the following steps to reduce this risk. First, we created de-identified interview transcripts and then deleted the original recordings. We saved identified consent forms separately from the de-identified interview transcripts (prior to publication, we maintained a separate file linking participants’ names and IDs to ensure that if a participant contacted us asking for their data to be removed, we would be able to do so). We also asked participants not to share confidential information with us during the interview process. Additionally, we allowed all participants to review direct quotes so they could request that we remove any information that might enable them to be identified.

## Acknowledgments

We thank members of Microsoft Research’s Sociotechnical Alignment Center (Chad Atalla, Emily Corvi, Alex Dow, Nick Pangakis, Stefanie Reed, Dan Vann, Matt Vogel, and Hannah Washington) for their invaluable feedback and participation in

our pilot studies. We also thank members of Microsoft Research’s FATE group. Finally, we thank the practitioners who agreed to be interviewed.

## References

- Robert Adcock and David Collier. 2001. [Measurement Validity: A Shared Standard for Qualitative and Quantitative Research](#). *American Political Science Review*, 95(3):529–546.
- Maria Antoniak and David Mimno. 2021. [Bad seeds: Evaluating lexical methods for bias measurement](#). In *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)*, pages 1889–1904, Online. Association for Computational Linguistics.
- Agathe Balayn, Mireia Yurrita, Jie Yang, and Ujwal Gadiraju. 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*, AIES, pages 482–495, Montreal QC Canada. ACM.
- Carliss Y. Baldwin and Kim B. Clark. 2000. [Design Rules, Volume 1: The Power of Modularity](#). The MIT Press.
- Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. 2017. The problem with bias: Allocative versus representational harms in machine learning. In *Proceedings of SIGCIS*, Philadelphia, PA.
- Emily Bender. 2019. [The #benderrule: On naming the languages we study and why it matters](#). The Gradient.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. 2021. [Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets](#). In *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)*, pages 1004–1015, Online. Association for Computational Linguistics.
- Rishi Bommasani and Percy Liang. 2024. Trustworthy social bias measurement. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, pages 210–224.

<sup>10</sup>Our semi-structured interview guide is in Appendix B, and we provide details about participant recruitment in §3. The study was approved by Microsoft’s research IRB.

- Virginia Braun and Victoria Clarke. 2006. [Using thematic analysis in psychology](#). *Qualitative Research in Psychology*, 3(2):77–101.
- Virginia Braun and Victoria Clarke. 2019. [Reflecting on reflexive thematic analysis](#). *Qualitative Research in Sport, Exercise and Health*, 11(4):589–597.
- Kelly Caine. 2016. [Local Standards for Sample Size at CHI](#). In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 981–992, San Jose California USA. ACM.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- Yang Trista Cao, Yada Pruksachatkun, Kai-Wei Chang, Rahul Gupta, Varun Kumar, Jwala Dhamala, and Aram Galstyan. 2022. [On the intrinsic and extrinsic fairness evaluation metrics for contextualized language representations](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 561–570, Dublin, Ireland. Association for Computational Linguistics.
- Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2021. [HateBERT: Retraining BERT for abusive language detection in English](#). In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, pages 17–25, Online. Association for Computational Linguistics.
- Kate Crawford. 2017. The trouble with bias. In *Conference on Neural Information Processing Systems (invited speaker)*.
- Pieter Delobelle, Giuseppe Attanasio, Debora Nozza, Su Lin Blodgett, and Zeerak Talat. 2024. [Metrics for what, metrics for whom: Assessing actionability of bias evaluation metrics in NLP](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21669–21691, Miami, Florida, USA. Association for Computational Linguistics.
- Pieter Delobelle, Ewoenam Tokpo, Toon Calders, and Bettina Berendt. 2022. [Measuring fairness with biased rulers: A comparative study on bias metrics for pre-trained language models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1693–1706, Seattle, United States. Association for Computational Linguistics.
- Wesley Hanwen Deng, Manish Nagireddy, Michelle Seng Ah Lee, Jatinder Singh, Zhiwei Steven Wu, Kenneth Holstein, and Haiyi Zhu. 2022. [Exploring How Machine Learning Practitioners \(Try To\) Use Fairness Toolkits](#). In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT*, pages 473–484, Seoul Republic of Korea. ACM.
- Sunipa Dev, Masoud Monajatipoor, Anaelia Ovalle, Arjun Subramonian, Jeff Phillips, and Kai-Wei Chang. 2021. [Harms of gender exclusivity and challenges in non-binary representation in language technologies](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1968–1994, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Sunipa Dev, Emily Sheng, Jieyu Zhao, Aubrie Amstutz, Jiao Sun, Yu Hou, Mattie Sanseverino, Jiin Kim, Akihiro Nishi, Nanyun Peng, and Kai-Wei Chang. 2022. [On measures of biases and harms in NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, pages 246–267, Online only. Association for Computational Linguistics.
- Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta. 2021. [BOLD: Dataset and Metrics for Measuring Biases in Open-Ended Language Generation](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT*, pages 862–872, Virtual Event Canada. ACM.
- Yupei Du, Qixiang Fang, and Dong Nguyen. 2021. [Assessing the reliability of word embedding gender bias measures](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10012–10034, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Paul A. Estabrooks, Maureen Boyle, Karen M. Emons, Russell E. Glasgow, Bradford W. Hesse, Robert M. Kaplan, Alexander H. Krist, Richard P. Moser, and Martina V. Taylor. 2012. [Harmonized patient-reported data elements in the electronic health record: supporting meaningful use by primary care action on health behaviors and key psychosocial factors](#). *Journal of the American Medical Informatics Association: JAMIA*, 19(4):575–582.
- Eve Fleisig, Aubrie Amstutz, Chad Atalla, Su Lin Blodgett, Hal Daumé III, Alexandra Olteanu, Emily Sheng, Dan Vann, and Hanna Wallach. 2023. [Fair-Prism: Evaluating fairness-related harms in text generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6231–6251, Toronto, Canada. Association for Computational Linguistics.
- Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2024. [Bias and fairness in large language models: A survey](#). *Computational Linguistics*, 50(3):1097–1179.
- Sebastian Gehrmann, Elizabeth Clark, and Thibault Selam. 2023. [Repairing the Cracked Foundation: A Survey of Obstacles in Evaluation Practices for Generated Text](#). *Journal of Artificial Intelligence Research*, 77:103–166.



- Russell E. Glasgow and William T. Riley. 2013. [Pragmatic Measures](#). *American Journal of Preventive Medicine*, 45(2):237–243.
- Seraphina Goldfarb-Tarrant, Rebecca Marchant, Ricardo Muñoz Sánchez, Mugdha Pandya, and Adam Lopez. 2021. [Intrinsic bias metrics do not correlate with application bias](#). In *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)*, pages 1926–1940, Online. Association for Computational Linguistics.
- Seraphina Goldfarb-Tarrant, Eddie Ungless, Esma Balkir, and Su Lin Blodgett. 2023. [This prompt is measuring <mask>: evaluating bias evaluation in language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2209–2225, Toronto, Canada. Association for Computational Linguistics.
- Rishav Hada, Agrima Seth, Harshita Diddee, and Kalika Bali. 2023. [“fifty shades of bias”: Normative ratings of gender bias in GPT generated English text](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1862–1876, Singapore. Association for Computational Linguistics.
- Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. [ToxiGen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3309–3326, Dublin, Ireland. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [{DEBERTA}: {DECODING}-{enhanced} {bert} {with} {disentangled} {attention}](#). In *International Conference on Learning Representations*.
- Monique Hennink and Bonnie N. Kaiser. 2022. [Sample sizes for saturation in qualitative research: A systematic review of empirical tests](#). *Social Science & Medicine*, 292:114523.
- Valentin Hofmann, Pratyusha Ria Kalluri, Dan Jurafsky, and Sharese King. 2024. [Ai generates covertly racist decisions about people based on their dialect](#). *Nature*, 633(8028):147–154.
- Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé, Miro Dudik, and Hanna Wallach. 2019. [Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need?](#) In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI, pages 1–16, Glasgow Scotland Uk. ACM.
- Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madian Khabza. 2023. [Llama guard: Llm-based input-output safeguard for human-ai conversations](#). *Preprint*, arXiv:2312.06674.
- Abigail Z. Jacobs and Hanna Wallach. 2021. [Measurement and Fairness](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT, pages 375–385, New York, NY, USA. Association for Computing Machinery.
- Jared Katzman, Angelina Wang, Morgan Scheuerman, Su Lin Blodgett, Kristen Laird, Hanna Wallach, and Solon Barocas. 2023. [Taxonomizing and measuring representational harms: a look at image tagging](#). In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’23/IAAI’23/EAAI’23. AAAI Press.
- Hadas Kotek, Rikker Dockum, and David Sun. 2023. [Gender bias and stereotypes in large language models](#). In *Proceedings of The ACM Collective Intelligence Conference*, CI ’23, page 12–24, New York, NY, USA. Association for Computing Machinery.
- Michelle Seng Ah Lee and Jat Singh. 2021. [The Landscape and Gaps in Open Source Fairness Toolkits](#). In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI, pages 1–13, Yokohama Japan. ACM.
- Alyssa Lees, Vinh Q. Tran, Yi Tay, Jeffrey Sorensen, Jai Gupta, Donald Metzler, and Lucy Vasserman. 2022. [A new generation of perspective api: Efficient multi-lingual character-level transformers](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD ’22, page 3197–3207, New York, NY, USA. Association for Computing Machinery.
- Cara C. Lewis and Caitlin Dorsey. 2020. [Advancing Implementation Science Measurement](#). In Bianca Albers, Aron Shlonsky, and Robyn Mildon, editors, *Implementation Science 3.0*, pages 227–251. Springer International Publishing, Cham.
- Cara C Lewis, Kayne D Mettert, Cameo F Stanick, Heather M Halko, Elspeth A Nolen, Byron J Powell, and Bryan J Weiner. 2021. [The psychometric and pragmatic evidence rating scale \(PAPERS\) for measure development and evaluation](#). *Implementation Research and Practice*, 2:26334895211037391.
- Ahmed Magooda, Alec Helyar, Kyle Jackson, David Sullivan, Chad Atalla, Emily Sheng, Dan Vann, Richard Edgar, Hamid Palangi, Roman Lutz, Hongliang Kong, Vincent Yun, Eslam Kamal, Federico Zarfati, Hanna Wallach, Sarah Bird, and Mei Chen. 2023. [A Framework for Automated Measurement of Responsible AI Harms in Generative AI Applications](#). *arXiv preprint*. ArXiv:2310.17750 [cs].
- Gaurav Maheshwari, Aurélien Bellet, Pascal Denis, and Mikaela Keller. 2023. [Fair without leveling down](#).



- A new intersectional fairness definition. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9018–9032, Singapore. Association for Computational Linguistics.
- Ruben G. Martinez, Cara C. Lewis, and Bryan J. Weiner. 2014. [Instrumentation issues in implementation science](#). *Implementation science: IS*, 9:118.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. [On measuring social biases in sentence encoders](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 622–628, Minneapolis, Minnesota. Association for Computational Linguistics.
- David L Morgan. 2008. Snowball sampling. *The SAGE encyclopedia of qualitative research methods*, 2:815–16.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pretrained language models](#). In *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)*, pages 5356–5371, Online. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for NLG](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252, Copenhagen, Denmark. Association for Computational Linguistics.
- Victor Ojewale, Ryan Steed, Briana Vecchione, Abeba Birhane, and Inioluwa Deborah Raji. 2024. [Towards AI Accountability Infrastructure: Gaps and Opportunities in AI Audit Tooling](#). *arXiv preprint*.
- Matthew J Page, Joanne E McKenzie, Patrick M Bossuyt, Isabelle Boutron, Tammy C Hoffmann, Cynthia D Mulrow, Larissa Shamseer, Jennifer M Tetzlaff, Elie A Akl, Sue E Brennan, Roger Chou, Julie Glanville, Jeremy M Grimshaw, Asbjørn Hróbjartsson, Manoj M Lalu, Tianjing Li, Elizabeth W Loder, Evan Mayo-Wilson, Steve McDonald, Luke A McGuinness, Lesley A Stewart, James Thomas, Andrea C Tricco, Vivian A Welch, Penny Whiting, and David Moher. 2021. [The prisma 2020 statement: an updated guideline for reporting systematic reviews](#). *BMJ*, 372.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel Bowman. 2022. [BBQ: A hand-built bias benchmark for question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2086–2105, Dublin, Ireland. Association for Computational Linguistics.
- Ian Porada, Alexandra Olteanu, Kaheer Suleman, Adam Trischler, and Jackie Cheung. 2024. [Challenges to evaluating the generalization of coreference resolution models: A measurement modeling perspective](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15380–15395, Bangkok, Thailand. Association for Computational Linguistics.
- Byron J Powell, Thomas J Waltz, Matthew J Chinman, Laura J Damschroder, Jeffrey L Smith, Monica M Matthieu, Enola K Proctor, and JoAnn E Kirchner. 2015. [A refined compilation of implementation strategies: results from the Expert Recommendations for Implementing Change \(ERIC\) project](#). *Implementation Science*, 10(1):21.
- Ehud Reiter. 2018. [A structured review of the validity of BLEU](#). *Computational Linguistics*, 44(3):393–401.
- Brianna Richardson, Jean Garcia-Gathright, Samuel F. Way, Jennifer Thom, and Henriette Cramer. 2021. [Towards Fairness in Practice: A Practitioner-Oriented Rubric for Evaluating Fair ML Toolkits](#). In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–13, Yokohama Japan. ACM.
- Morgan Klaus Scheuerman. 2024. [In the Walled Garden: Challenges and Opportunities for Research on the Practices of the AI Tech Industry](#). In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 456–466, Rio de Janeiro Brazil. ACM.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2023. Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. *arXiv preprint arXiv:2310.11324*.
- Preethi Seshadri, Pouya Pezeshkpour, and Sameer Singh. 2022. [Quantifying social biases using templates is unreliable](#). In *Workshop on Trustworthy and Socially Responsible Machine Learning, NeurIPS 2022*.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2021. [Societal biases in language generation: Progress and challenges](#). In *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)*, pages 4275–4293, Online. Association for Computational Linguistics.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. [The woman worked as a babysitter: On biases in language generation](#). In

- Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3407–3412, Hong Kong, China. Association for Computational Linguistics.
- Bangzhao Shu, Lechen Zhang, Minje Choi, Lavinia Dunagan, Lajanugen Logeswaran, Moontae Lee, Dallas Card, and David Jurgens. 2024. [You don’t need a personality test to know these models are unreliable: Assessing the reliability of large language models on psychometric instruments](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5263–5281, Mexico City, Mexico. Association for Computational Linguistics.
- Mario Luis Small. 2009. [‘how many cases do i need?’: On science and the logic of case selection in field-based research](#). *Ethnography*, 10(1):5–38.
- Cameo F. Stanick, Heather M. Halko, Elspeth A. Nolen, Byron J. Powell, Caitlin N. Dorsey, Kayne D. Mettert, Bryan J. Weiner, Melanie Barwick, Luke Wolfenden, Laura J. Damschroder, and Cara C. Lewis. 2021. [Pragmatic measures for implementation research: development of the Psychometric and Pragmatic Evidence Rating Scale \(PAPERS\)](#). *Translational Behavioral Medicine*, 11(1):11–20.
- Kaiser Sun, Adina Williams, and Dieuwke Hupkes. 2023. [The validity of evaluation results: Assessing concurrence across compositionality benchmarks](#). In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 274–293, Singapore. Association for Computational Linguistics.
- Oskar Van Der Wal, Dominik Bachmann, Alina Leiding, Leendert Van Maanen, Willem Zuidema, and Katrin Schulz. 2024. [Undesirable Biases in NLP: Addressing Challenges of Measurement](#). *Journal of Artificial Intelligence Research*, 79:1–40.
- Aniket Vashishtha, S Sai Prasad, Payal Bajaj, Vishrav Chaudhary, Kate Cook, Sandipan Dandapat, Sunayana Sitaram, and Monojit Choudhury. 2023. [Performance and risk trade-offs for multi-word text prediction at scale](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2226–2242, Dubrovnik, Croatia. Association for Computational Linguistics.
- Pranav Narayanan Venkit, Mukund Srinath, and Shomir Wilson. 2022. [A study of implicit bias in pretrained language models against people with disabilities](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1324–1332, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Hanna Wallach, Meera Desai, A. Feder Cooper, Angelina Wang, Chad Atalla, Solon Barocas, Su Lin Blodgett, Alexandra Chouldechova, Emily Corvi, P. Alex Dow, Jean Garcia-Gathright, Alexandra Olteanu, Nicholas Pangakis, Stefanie Reed, Emily Sheng, Dan Vann, Jennifer Wortman Vaughan, Matthew Vogel, Hannah Washington, and Abigail Z. Jacobs. 2025. [Position: Evaluating generative ai systems is a social science measurement challenge](#). *Preprint*, arXiv:2502.00561.
- Thomas J. Waltz, Byron J. Powell, Monica M. Matthieu, Laura J. Damschroder, Matthew J. Chinman, Jeffrey L. Smith, Enola K. Proctor, and JoAnn E. Kirchner. 2015. [Use of concept mapping to characterize relationships among implementation strategies and assess their feasibility and importance: results from the Expert Recommendations for Implementing Change \(ERIC\) study](#). *Implementation Science*, 10(1):109.
- Angelina Wang, Xuechunzi Bai, Solon Barocas, and Su Lin Blodgett. 2023. [Measuring stereotype harm from machine learning errors requires understanding who is being harmed by which errors in what ways](#). In *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*.
- Ziang Xiao, Susu Zhang, Vivian Lai, and Q. Vera Liao. 2023. [Evaluating evaluation metrics: A framework for analyzing NLG evaluation metrics using measurement theory](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10967–10982, Singapore. Association for Computational Linguistics.
- Dora Zhao, Jerone T. A. Andrews, Orestis Papakyriakopoulos, and Alice Xiang. 2024. [Position: Measure Dataset Diversity, Don’t Just Claim It](#). In *The Forty-first International Conference on Machine Learning, ICML, Vienna, Austria*. arXiv. Version Number: 1.
- Kaitlyn Zhou, Su Lin Blodgett, Adam Trischler, Hal Daumé III, Kaheer Suleman, and Alexandra Olteanu. 2022. [Deconstructing NLG evaluation: Evaluation practices, assumptions, and their implications](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 314–324, Seattle, United States. Association for Computational Linguistics.

## A Systematic Literature Review to Identify Desiderata

To identify the set of desiderata that we used to scaffold our interviews, we conducted a systematic review of the NLP literature on assessing measurement instruments. To do this, we followed PRISMA guidelines (Page et al., 2021). We provide our PRISMA flow diagram in Figure 1.

### A.1 Search Strategy

**ACL Anthology Search.** Using the ACL Anthology API, we identified all papers published in ACL venues through 2024 that met the following criteria:

- Title contains: ‘eval\*’ OR ‘measur\*’
- Abstract contains: ‘meta’ OR ‘survey\*’ OR ‘review\*’ OR ‘assess\*’ OR ‘audit\*’

This search produced 1,075 results. We reviewed these search results for title and abstract relevance (see §A.1.1) in order to identify ten papers that were about assessing measurement instruments. We then used these ten papers as “seed papers” to conduct an additional citation mapping search.

**Citation Mapping Search.** To ensure that we captured relevant work that was not published in ACL venues or that may have been missed by our keyword search, we used the Semantic Scholar API to identify all papers that were referenced by or cited the ten seed papers. We then selected all papers that were referenced by or cited at least two of the seed papers, resulting in an additional 59 papers, which we reviewed for title and abstract relevance.

#### A.1.1 Selection Strategy

To screen the results of our ACL Anthology keyword search, we focused on only those papers that were published in ACL conferences or journals (i.e., not front matter, tutorials, or workshop papers) and that had not been marked as deleted. After screening the papers, the first author iteratively reviewed each remaining paper for title relevance and, if the title was deemed relevant, abstract relevance. Titles and abstracts were considered relevant if they indicated that the papers appeared to be about 1) assessing measurement instruments, 2) instruments for measuring abstract or contested concepts from text (e.g., not about information retrieval, not about visual or audio tasks), and 3) measurement instruments (e.g., not solely about annotator reliability or researchers’

misuse of measurement instruments). We used the same title and abstract relevance criteria to screen the results of the citation mapping search. Finally, the first author evaluated all remaining papers for full-text relevance, again considering papers relevant if they were about the topics listed above.

#### A.1.2 Annotation Strategy

The first author read each paper and extracted all passages of text that appeared to identify a particular desideratum of measurement instruments (e.g., text identifying a quality that measurement instruments should have, or text identifying a quality such that instruments lacking that quality are challenging to use). All authors other than the first author developed an initial list of desiderata (validity, reliability, specificity, extensibility, scalability, interpretability, and actionability; see Table 3) based on their experiences measuring representational harms caused by LLM-based systems. In discussion with the other authors, the first author mapped each passage to these desiderata. While conducting this exercise, we did not identify any additional desiderata. In other words, we were able to map all extracted passages to one of the desiderata listed in Table 3.

#### A.1.3 Results

We list the papers mentioning each desideratum below. Some papers mention the desideratum explicitly. Others mention the desideratum implicitly (e.g., not using the exact terminology we do, but describing the concept captured by the desideratum).

**Mentioned validity.** Blodgett et al. (2020, 2021); Delobelle et al. (2022, 2024); Du et al. (2021); Gehrmann et al. (2023); Goldfarb-Tarrant et al. (2021, 2023); Novikova et al. (2017); Reiter (2018); Sun et al. (2023); Van Der Wal et al. (2024); Xiao et al. (2023); Zhou et al. (2022).

**Mentioned reliability.** Blodgett et al. (2021); Delobelle et al. (2022, 2024); Du et al. (2021); Gehrmann et al. (2023); Goldfarb-Tarrant et al. (2023); Novikova et al. (2017); Seshadri et al. (2022); Sun et al. (2023); Van Der Wal et al. (2024); Xiao et al. (2023); Zhou et al. (2022).

**Mentioned specificity.** Delobelle et al. (2024); Du et al. (2021); Gehrmann et al. (2023); Van Der Wal et al. (2024); Zhou et al. (2022).

**Mentioned extensibility.** Gehrmann et al. (2023); Reiter (2018); Zhou et al. (2022).

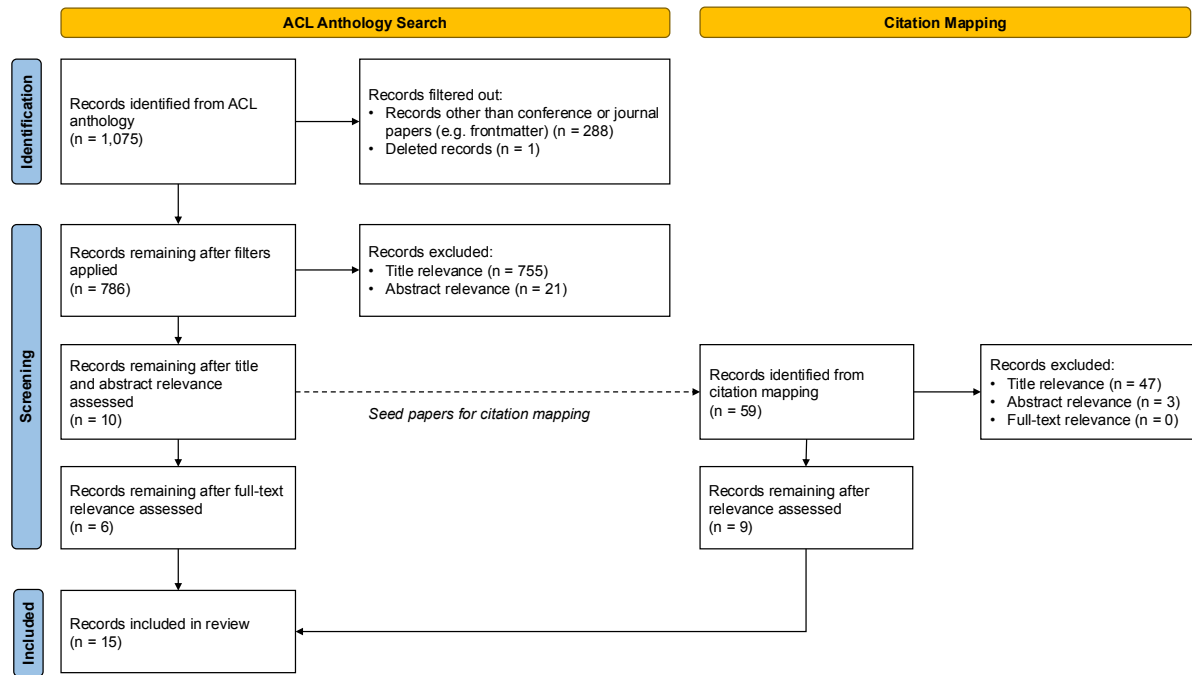


Figure 1: Our PRISMA flow diagram.

**Mentioned scalability.** Gehrmann et al. (2023); Novikova et al. (2017); Xiao et al. (2023); Zhou et al. (2022).

**Mentioned interpretability.** Delobelle et al. (2024); Du et al. (2021); Gehrmann et al. (2023); Van Der Wal et al. (2024); Xiao et al. (2023).

**Mentioned actionability.** Delobelle et al. (2024).

## B Semi-Structured Interview Guide

Our semi-structured interview guide is shown below. As is typical of semi-structured interviews, not every participant was asked exactly the same questions in exactly the same order, and some participants were asked additional follow-up or clarifying questions based on the answers they provided. The interview questions were supplemented with a set of slides containing definitions of key terms that we screenshared with participants. The definitions are included in the script below.

### B.1 Introductions [5 min]

Welcome! Thank you so much for taking the time for this interview. Before we get started, I just want to quickly introduce myself, talk about the goals of this study, and give you a chance to ask any questions you might have. This research study is intended to understand gaps between research and practice in evaluating large language model

(LLM)-based systems, with a focus on measuring harms, adverse impacts, or other undesirable behaviors. In this interview, I'll ask you to share your experiences with and opinions on such evaluations, without discussing confidential information. I will also record this interview for the purpose of creating a deidentified transcript. If you prefer that your video not be recorded, please feel free to turn your camera off at this time. In addition, if at any point you would like to skip a question, take a break, or end the interview, please feel free to do so.

Do you have any questions before we get started?

### B.2 Background [5 min]

[Q1] To start, please briefly describe your role, focusing on your professional experience as it relates to LLM-based systems.

[Q2] Can you briefly describe the LLM-based system(s) that you have previously evaluated, currently evaluate, or plan to evaluate?

### B.3 Experience with measurement instruments for representational harms [15 min]

[Q3] Throughout this interview, I will be focusing primarily on representational harms, which occur when “a system represents some social groups in a less favorable light than it represents other groups by stereotyping them,



demeaning them, or failing to recognize their existence altogether.”

What examples of representational harms caused by LLM-based systems are you aware of?

*If interviewee was not familiar with representational harms, we provided the following examples:*

- LLMs might reinforce stereotypes, for example, by using the word “nurse” to refer to a female healthcare provider and the word “doctor” to refer to a male healthcare provider in otherwise identical contexts.
- LLMs might generate slurs or derogatory language about a social group.
- LLMs might erase a social group, for example, by only listing male athletes when a user asks for examples of talented soccer players, thus failing to recognize the existence of non-male soccer players.

[Q4] Do your previous, current, or planned evaluation(s) of LLM-based system(s) involve measuring representational harms?

[Q5] What types of representational harms are you measuring?

[Q6] Can you walk me through, from start to finish, an example of how you measure representational harms? I’m especially interested in hearing about how you decided on your approach, whether you relied on existing, publicly available tools, benchmarks, datasets, metrics, annotation guidelines, and so on, or whether you decided to develop your own.

*To allow for open-ended discussion, we did not provide participants with a specific definition of ‘measurement instruments’; rather, we provided the following examples of instruments:*

- An example of a tool is Perspective API.
- An example of a benchmark is StereoSet, which includes a dataset of prompts that could elicit stereotyping content with corresponding metrics that measure the extent to which a language model produces stereotypes.

- An example of a dataset is WildChat, which is a corpus of 1 million real user-ChatGPT interactions.
- Examples of metrics are the Word and Sentence Embedding Association Tests (WEAT and SEAT), which measure whether “attribute words” (e.g. male, female) are disproportionately associated with a set of “target words” (e.g. different professions).
- Annotation instructions are sets of instructions and examples for humans to use when annotating system outputs for particular properties.
- An example of another type of instrument is Matched Guide Probing, a method adapted from sociolinguistics.

*For each instrument mentioned, we asked the following questions:*

[Q7] What type(s) of representational harms are you measuring with [this instrument]?

[Q8] How did you decide to use [this instrument]?

[Q9] How do you use [this instrument] in your evaluation(s)?

[Q10] Where did [this instrument] come from? Did you develop it yourself, modify an existing [instrument], or use an existing [instrument] as-is?

*If applicable, for one instrument that the interviewee developed themselves, we asked the following questions:*

[Q11] Why did you decide to develop [this instrument] yourself?

[Q12] What, if any, actions have you taken or plan to take upon seeing the measurements obtained using [this instrument]?

*If applicable, for one instrument that the interviewee adapted from an existing instrument, we asked the following questions:*

[Q13] Why did you decide to start with this existing [instrument]?

[Q14] Why did you decide to modify [this instrument] rather than using it as-is?

[Q15] What, if any, actions have you taken or plan to take upon seeing the measurements obtained using [this instrument]?

*If applicable, for one instrument that the interviewee used as-is, we asked the following questions:*

[Q16] Why did you decide to use this existing [instrument] as-is?

[Q17] What, if any, actions have you taken or plan to take upon seeing the measurements obtained using [this instrument]?

#### **B.4 Challenges with measurement instruments for representational harms [15 min]**

[Q18] Were there any other existing, publicly available [instruments] that you investigated using instead?

[Q19] *For each instrument mentioned:* Why did you decide not to use [this instrument]?

[Q20] *For each of the challenges defined below, say either:*

“It sounds like you mentioned an issue to do with [challenge]. Is that correct?”, *or*

“I don’t think you mentioned [challenge]. Did you experience any issues with this?”

*We provided interviewees with the following set of challenges related to measurement instruments:*

- Whether it results in valid measurements – i.e., meaningfully measures what stakeholders think it measures
- Whether it results in similar measurements when used in similar ways, especially over time
- Whether it is sufficiently specific to the system being evaluated and its particular use cases and deployment contexts
- Whether it can scale to increasing workloads
- Whether it can be adapted for different systems, use cases, and deployment contexts
- Whether its resulting measurements can be understood by stakeholders
- Whether its resulting measurements can be acted upon by stakeholders

[Q21] *For each challenge experienced:* What, if anything, did you do to address this issue?

[Q22] Did you experience any other issues that we haven’t discussed?

[Q23] *If applicable:* What, if anything, did you do to address this issue?

#### **B.5 Comparing measurement of representational harms to other harms [5 min]**

[Q24] Do your previous, current, or planned evaluation(s) of LLM-based system(s) involve measuring harms, adverse impacts, or other undesirable behaviors other than representational harms?

[Q25] *If yes:* What types of harms, adverse impacts, or other undesirable behaviors?

[Q26] *If yes:* Are your experiences measuring these types of harms, adverse impacts, or other undesirable behaviors similar to your experiences measuring representational harms? What, if anything, is similar and what, if anything, is different about your experiences? I’m especially interested in hearing about how the [instruments] you use to measure these types of harms, adverse impacts, or other undesirable behaviors are similar to or different from the [instruments] you use to measure representational harms.

#### **B.6 Desired improvements to measuring representational harms [5 min]**

[Q27] Putting aside any time or budget constraints, what, if anything, would you improve about the way that you previously, currently, or plan to measure representational harms?

[Q28] What do you need, that you don’t currently have, in order to make those improvements?

#### **B.7 Closing [5 min]**

[Q29] Is there anything else you would like to tell us about your previous, current, or planned evaluation(s) of LLM-based system(s)?