

# Beyond Behaviorist Representational Harms: A Plan for Measurement and Mitigation

Jennifer Chien  
jjchien@ucsd.edu  
UC San Diego  
USA

David Danks  
ddanks@ucsd.edu  
UC San Diego  
USA

## ABSTRACT

Algorithmic harms are commonly categorized as either allocative or representational. This study specifically addresses the latter, examining current definitions of representational harms to discern what is included and what is not. This analysis motivates our expansion beyond behavioral definitions to encompass harms to cognitive and affective states. The paper outlines high-level requirements for measurement: identifying the necessary expertise to implement this approach and illustrating it through a case study. Our work highlights the unique vulnerabilities of large language models to perpetrating representational harms, particularly when these harms go unmeasured and unmitigated. The work concludes by presenting proposed mitigations and delineating when to employ them. The overarching aim of this research is to establish a framework for broadening the definition of representational harms and to translate insights from fairness research into practical measurement and mitigation praxis.

## CCS CONCEPTS

• **Applied computing** → **Psychology**; • **Human-centered computing** → *Human computer interaction (HCI)*; • **Computing methodologies** → **Machine learning**; • **Information systems** → **Collaborative and social computing systems and tools**;

### ACM Reference Format:

Jennifer Chien and David Danks. 2024. Beyond Behaviorist Representational Harms: A Plan for Measurement and Mitigation. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3630106.3658946>

## 1 INTRODUCTION

Algorithmic predictions are increasingly deployed across high-stakes domains such as hiring [3], finance [110], and scientific discovery [90, 124]. These systems can produce significant harms, leading to many efforts to either reduce the harms, or at least ensure that they occur fairly. That is, a necessary (though not sufficient) condition for deployment of a high-stakes system is that it not provide unintentionally disparate benefits, or produce disparate harms, across groups or individuals. As a result, many fairness evaluations have focused on *allocative* harms – i.e. disparities in

access to resources or other material benefits that make a group or individual worse-off. Allocative harms are relatively more easily measured, as they include physical, financial, social, and other observable resources (e.g., wages and jobs lost, loans not approved, houses not bought, or pharmaceutical interventions not pursued).

However, not all harms are allocative in nature. In particular, representational harms include the less tangible harms of algorithmic systems, such as behavioral, psychological, and cognitive effects on reputation, societal standing, and cultural harmony. For example, consider an algorithmic decision that someone is a credit risk, but still appropriate for a loan. This applicant has not suffered any allocative harm, as they received the resources. However, the applicant may suffer a representational harm, particularly if the algorithmic output incorrectly leads them to think negatively about themselves. As this example suggests, representational harms are significantly more difficult to measure, which may explain their relative neglect in the literature on algorithmic harms. Nonetheless, representational harms may have equal, if not more dire, downstream effects than allocative harms (e.g., hate crimes, lowered self-confidence, increased suicide rates, loss of group identification or social connection). We thus aim to bring renewed attention to the definition and measurement of representational harms so that they can be considered in design, development, deployment, and mitigation decisions.

Representational harms are particularly salient when we consider the potential negative impacts of large language models (LLMs), which have undergone a recent explosion in capabilities and applications. Products built with natural language interfaces (typically with foundation models) include outputs that span natural language, images, and videos, and therefore have the potential to produce representational harms across a vast variety of domains. From stereotype-reinforcing outputs [69] to outsider gaze perpetuating images [99] to consistent misgendering [87], LLMs threaten to silently and pervasively enact unmeasured and unmitigated harms to minority and majority groups alike, particularly when masquerading as human productions.

In this work, we examine existing definitions of representational harms (Section 2) and identify a notable limitation – a focus on ‘tangible’ harms. These definitions primarily focus on observing how one person acts toward another person or group, neglecting changes in internal cognitive states or reasoning, considerations of alternative possibilities, experiential fluctuations, self-identification, and similar psychological and social impacts. We develop and defend an expanded understanding of representational harms, offering illustrative themes and measures (Section 3). Improved definitions do not directly lead to improved practice, however. This motivates Section 4, in which we outline high-level requirements for practical

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

FAccT '24, June 03–06, 2024, Rio de Janeiro, Brazil

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0450-5/24/06

<https://doi.org/10.1145/3630106.3658946>

Concept	Reference	Application
Social Stereotyping	[7, 13, 23, 31, 42, 48, 57, 65, 87, 101, 111, 115, 119, 123, 128, 134, 135]	Image Captioning, Generative AI, ML, Image Cropping, Story Generation, Image Tagging, Knowledge Distillation
Reification of Social Groups or Use of Essentialist Categories	[7, 65, 111, 119, 123, 128, 135]	Image Captioning, Image Cropping, Image Tagging
Inaccurate/Skewed Representations	[44, 82, 119, 128, 135, 137]	Misinformation, Image Search, Image Cropping
Demeaning/Derogatory Language	[65, 82, 111, 123, 128, 135]	Image Captioning, Image Cropping, Image Tagging, Misinformation
Denial of Self-Identification	[65, 87, 111, 123, 128]	Image Captioning, Image Tagging, Generative AI
(Hyper)Attention/Exposure/Erasure	[44, 48, 65, 111, 115]	Image Search, Image Tagging
Discrimination/Incite (and Normalize) Hate/Violence	[7, 48, 115, 128]	LLMs
Outsider Gaze	[13, 23, 87]	Generative AI, Story Generation
Hierarchies & Marginalization	[23, 111, 115]	Image Tagging, Story Generation

**Table 1: Types of Representational Harms and Associated Applications.**

implementation, supported by a case study. These contributions collectively establish a framework that can be expanded as additional types of representational harms are identified and characterized. We explore the heightened susceptibility of LLM-powered products to propagate representational harms (Section 5). Lastly, we suggest potential mitigations to reduce the likelihood and severity of representational harms, including conditions under which one might deploy them (Section 6). Collectively, we aim to bring representational harms into discussions, analyses, and practices that have previously been dominated by considerations of allocative harms.

## 2 REPRESENTATIONAL HARMS: DEFINITION AND LIMITATIONS

In order to understand the current landscape of representational harm definitions, we conducted a brief survey via snowball sampling using the keyword “representational harms” in the ACM Digital Library and in Google Scholar.<sup>1</sup> Table 1 provides the key types of representational harm that repeatedly arose across the different papers. Although there is significant diversity in the types of representational harms, there are some clear patterns in the characterizations.

*Most definitions are behaviorist.* Most definitions focus solely on behavioral evidence of harm, rather than the harm itself. In particular, most of these types of harm exclude aspects of conscious experience or phenomenology [125]. All definitions of representational harm surveyed (except for denial of self-identification) point to observable behaviors, either by individuals directly or by society as a whole. Social stereotyping, demeaning/derogatory language, discrimination, and hate/violence all point to explicit undesired behaviors enacted *towards* groups of people.

While observable behaviors obviously can provide important evidence about representational harms, a behaviorist approach fails

to provide insight into the relationship between exposure to biased outputs and corresponding increases in undesired behaviors. Instead, this focus implies simply that, for instance, an increased exposure to negative group associations causes their acceptance as truth. This thereby disregards the extent to which human critical reasoning, belief propagation, and personal experiences interact with an individual’s own beliefs about the world [52]. Moreover, a behaviorist approach potentially encourages a focus on local interactions, rather than including broader, more long-term manifestations of underlying representational harms.

*Ambiguity of correlation without rigorous requirement of causation.* Similarly, inaccurate/skewed representations, outsider gaze, reification of social groups, and hierarchies all provide diagnoses of representational harms (and broader societal behavioral problems). However, this characterization does not provide any guidance or insight about *how* depictions or algorithmic outputs reproduce or further entrench societal inequities. Instead, they merely describe the association between outputs and social challenges, with the causation left implicit. This does not elucidate the extent to which undesirable representations actively influence and shape our reality. Instead, this work accepts ambiguity between a reflection and the causation of currently existing inequalities.

*Association alone is not always problematic.* Even more importantly, the focus on behavioral signals fails to account for contexts in which, for instance, negative representations of a group might be morally permissible. That is, the mere existence of an association between negative representations (in algorithmic output) and a particular group may not always be morally problematic. For instance, the depiction of less-educated individuals as less creditworthy may be permissible as an empirical description of the current state of the economy. In fact, this algorithmic output might be valuable for designing better public policy to ensure more widespread, equitable access to financial resources. In the extreme, removal of these associations in an algorithm may actually lead to harms, as it could lead to people being approved for a loan without appropriate

<sup>1</sup>We acknowledge that this is not a comprehensive review, but is intended to provide a broad perspective on how representational harms have been characterized, demonstrated, measured, and embodied.

Type	Instance of Harm	Measure of Harm
Reduction in Options Considered (Individual)	Consideration of only a narrower subset of possibilities, such as for an occupation (e.g. gender gaps in nursing [121] as male nurses face questions about professional competency and masculinity), hobbies (e.g. impaired performance for girl soccer players believing they perform worse at soccer tasks [55]), or social connections (e.g. friend prejudices [1, 102]).	Decision Regret Scale [34], Quality of Life Enjoyment and Satisfaction Questionnaire [38]
Increased Stress (Individual)	“[U]nhealthy and debilitating” reactions to racial stereotypes, including heightened determination to defy or uphold stereotypes, result in emotional distress, stress, and a neglect of self-care and other basic needs [80].	Talbieh Brief Distress Inventory [103, 104], Brief COPE Scale [21]
Increased Social Anxiety and Other Mental Barriers (Individual)	Increased fear of stereotype confirmation and social anxiety, i.e. the fear of negative evaluation, resulting in behavioral avoidance [62].	Social Phobia Inventory (SPIN) [25], Social Responsiveness Scale (SRS) [26]
Increased Actual and Psychological Conflicts (Within Group)	Negative responses to a deviation from group norms resulting in feelings of rejection, lack of belonging, and/or increased pressure to conform. However, in the presence of alternative groups, a deviant was less likely to conform to the current group, and may even abandon the group after receiving an angry reaction [51].	Group Environment Questionnaire [20], Perceived Cohesion Scale [16], Rahim Organizational Conflict Inventory (ROCI) [100]
Increased Actual and Psychological Conflicts (Between Groups)	Harmful inaction (i.e. avoidance, withdrawal, and/or disregard) in intergroup conflict is uniquely associated with group-based contempt, facilitating intergroup disconnection and harm [36].	Semin and Fiedler’s linguistic category model for Linguistic Intergroup Bias [15, 76, 109]*
Reductions in Supportive or Other Healthy Behaviors (Indirect)	Normalization or decreased likelihood of seeking psychological services among refugees or others who have experienced trauma [37, 47].	Self-Identification as Having Mental Illness – Scale (SELF-I) [107]

**Table 2: Non-Exhaustive List of Types and Measures of Representational Harms.**

\* This is technically a measure of the phenomenon (i.e. the tendency to describe positive in-group and negative out-group behaviours in more abstract terms than negative in-group and positive out-group behaviours) rather than the harm (i.e. the frequency of harmful inaction), as the measurable effects of this type of harm may be quite broad and/or diffuse, depending on contextual factors.

supports in place. Removal or mitigation of this association in an algorithm can thus worsen societal inequities, even if it is done in the interests of reducing representational harms.

### 3 EXPANDING REPRESENTATIONAL HARMS

As the name suggests, representational harms are fundamentally grounded in the representations—both external and internal—of particular groups. Changes in those representations (due to algorithmic outputs) may lead to the behavioral changes outlined in Table 1, but those downstream effects are only diagnostic of representational harms, not constitutive of them. We thus propose an expansion of the definition of representational harms to include changes in cognitive (attention, engagement, workload, beliefs, desires, intentions, knowledges, thoughts, judgements, beliefs), affective (positive/negative valence, action-promoting/not), and emotional dimensions of mental representations<sup>2</sup> that potentially result in harm (physical, psychological, social, or other) to an individual.

<sup>2</sup>We do not engage with the large philosophical literature on the nature of mental representations, both due to space limitations and since the philosophical details do not impact our account of representational harms. Standard, intuitive views of representations will suffice for our purposes.

On this characterization, representational harms can arise across multiple temporal and social granularities (individual, intergroup, and intragroup). Table 2 provides just a few types of representational harm (with examples), as well as psychometrically validated measures/scales for that type of harm.<sup>3</sup>

Of course, the behaviorist harms that were identified in prior work may be proxies for cognitive, affective, and emotional changes in representations. For instance, incitement of hate or violence might result from narrowed option consideration, increased stress, pressure to conform to group norms, or increased disconnection between groups. However, as this example demonstrates, the behavioral effect or proxy potentially underdetermines the underlying changes in representations or mental processes, thereby failing to draw important distinctions between different types of representational harms. Thus, we need to adopt a more comprehensive

<sup>3</sup>In this table, we deliberately select harms that have established measures, though we recognize that some metrics may perform better along some dimensions than others (e.g., if they were developed and tested in selected demographics or localities). We provide an expanded table of references and examples in Table 3 and Table 4 in the Appendix.

characterization that is based in changes in mental representations, rather than the impacts of those changes.

## 4 RECIPES FOR REAL-WORLD MEASUREMENT

The right-most column of Table 2 provides examples of measures for each type of representational harm. We are not thereby implying that these are the *only* measures, but rather highlighting the fact that potential measures already exist. We do not always need to invent new measures for representational harms, but can adapt existing psychological research. That being said, many of these measures cannot be used “off the shelf” in real-world deployments of algorithmic systems. We thus turn to high-level requirements for monitoring and implementing various measures. This aims to ensure that identification of representational harms is properly aligned with mitigation of those harms. This section is intended to help practitioners understand the composition of expertise required to translate these research findings into praxis.

### 4.1 High-Level Requirements

First, it is crucial to pinpoint the class, instance, and types of harms most pertinent to a given system or deployment application. Psychologists and social scientists are likely well-suited for this task, given their familiarity with the latest measures and methods for identifying harms across a range of contexts. In addition, the active involvement of members from potentially affected communities is essential to identifying potential harms. Through their lived experiences, these stakeholders can thereby ensure that the measurement is culturally and logistically appropriate and sustainable (see e.g. Participatory Research [60, 95]).<sup>4</sup>

This advice—engage directly with potentially impacted communities and use a multidisciplinary approach—has also been frequently proposed in the context of allocative harms. Representational harms are different, however, as the subtle, diffuse nature of the underlying changes may make identification, measurement, and monitoring significantly more difficult. Many potential allocative harms can be identified with relatively little engagement with key communities; representational harms are likely to be much more elusive. Therefore, determining the appropriate granularity for measurement presents an ongoing statistical and measurement challenge. Evaluating harms at both individual and group levels may necessitate separate considerations and measures. For instance, minute effects on an individual’s likelihood of applying to a given job may not warrant immediate mitigation. However, the aggregation of these individual effects at the group level in terms of applications, hiring outcomes, and workplace resilience may imply a need for significant intervention. Conversely, the harms of one individual being significantly discouraged from applying to a job may not be significant at the group level, even though that individual has been harmed because of their (individual) representational changes. More generally, we contend that addressing representational harms

<sup>4</sup>We acknowledge that determining who counts as a stakeholder reflects legal, social, moral, pragmatic, and normative judgements. Each of these value judgements may contribute to the ultimate decision of whether the harm is escalated to the point of mitigation deployment.

will require significantly deeper engagement with diverse disciplinary perspectives, and potentially impacted communities.

Once the relevant granularities and groups are identified, we must also determine the corresponding measures. This poses another challenge as even if suitable measures exist, they must first be appropriately calibrated (and may need to be regularly re-calibrated based on cultural and temporal shifts, see e.g. [56, 75, 83, 105]). In addition, they may be difficult or impractical to implement at scale. Requesting users to complete surveys, particularly after every interaction, could introduce numerous biases (e.g., response bias, non-response bias, survivorship sampling bias) and decrease satisfaction [106]. Instead, we will frequently need to use passive measures, such as shifts in behavioral dispositions. In practice, this may lead to using some of the behavioral measures that have already been identified, but for the purpose of inferring unobserved representational changes (rather than thinking that they provide direct measures of representational harm). There is also a clear need for future research efforts to develop and validate passive measures, both individual and group, for the relevant representational changes. In what follows, we demonstrate these ideas in a case study.

### 4.2 Case Study

Consider the application of LLMs in a conversational setting in which a user strikes up a conversation about their favorite sport, soccer/fútbol, as summarized in Fig. 1 and Fig. 2.<sup>5</sup> In these figures, we provide both the back-and-forth with the LLM (left column) alongside an analysis of potential representational harms that could result from the LLM responses (right column). One key moral is that different types of (potential) representational harms are implicated in each response from the LLM, perhaps depending on the particular user. At the same time, there are (arguably) no direct allocative harms, as the LLM provides informative, if biased, answers to all of the queries.

While individual recognition of soccer skills may not immediately appear as a significant harm, it holds importance beyond fan recognition (e.g., diverse stakeholders are involved in determining awards). Acknowledging these skills becomes critical for a coach in determining team placement, access to specialized training, fostering inclusion, and impacting downstream diversity in professional soccer. Moreover, it plays a role in shaping individual self-esteem and a sense of capabilities [71]. Therefore, it may be prudent to consider measuring individual recognition of skills, self-esteem, and resilience. Potential measures for these harms could include the Perceived Competence Scales (PCS) (a 4-item questionnaire) [131], the Feelings of Inadequacy scale (a 23-item self-reported measure) [14], and the Five-Factor Model of personality structure (consisting of 50 questions) [17, 130].

Despite established measures for harm, this plan encounters notable challenges. Psychological measures, such as those assessing phenomena like self-esteem, primarily focus on the phenomenon itself rather than the resulting harm (e.g., impacts of low self-esteem like anxiety, depression, and attention problems [53]). Consequently, deploying these questions both before and after an

<sup>5</sup>This scenario is abridged and based on a real conversation with ChatGPT, full conversation found here.

Chat	Harm
<p><b>User:</b> Who are some famous soccer stars?  <b>ChatGPT:</b> As of my last knowledge update in January 2022, several soccer stars were widely recognized for their talent and achievements. Some of the notable players as of 2022 include:</p> <ol style="list-style-type: none"> <li>(1) Lionel Messi: An Argentine forward widely regarded as one of the greatest players of all time.</li> <li>(2) Cristiano Ronaldo: A Portuguese forward considered one of Messi’s greatest rivals. Ronaldo has played for several top clubs, including Sporting Lisbon, Manchester United, Real Madrid, and Juventus.</li> <li>(3) Neymar Jr.: A Brazilian forward who played alongside Messi at Barcelona before moving to PSG...</li> </ol>	<p><b>Option Consideration:</b> Did this result in a decreased propensity to pursue soccer? Since the output contains only male-identifying players, female, trans, and non-binary players may believe that the likelihood of being recognized and appropriately compensated for their talent or their ability to pursue the sport professionally is unlikely.  <b>Stress:</b> Has this person experienced increased stress in individual ability to succeed at soccer, resulting in increased determination or effort?  <b>Social:</b> Does this person see soccer along more gendered lines? Perhaps they see a male soccer player as more attractive/prestigious?</p>
<p><b>User:</b> Who decides this?  <b>ChatGPT:</b> Various entities and individuals play roles in determining aspects of soccer, including player transfers, team management, and overall governance. Here are some key stakeholders:</p> <ol style="list-style-type: none"> <li>(1) Clubs and Team Owners: Clubs have significant influence over player transfers. Owners and management make decisions about which players to buy or sell based on team needs, financial considerations, and other factors.</li> <li>(2) Managers and Coaches: Team managers and coaches are involved in player selection, tactics, and strategies on the field. They may have input into transfer decisions, and their preferences for certain players can influence the club’s decisions.</li> <li>(3) Agents: Many players have agents who represent them in contract negotiations and other business matters. Agents play a crucial role in securing deals, transfers, and endorsements for players. They work to ensure that their clients get the best possible opportunities and contracts...</li> </ol>	<p><b>Option Consideration &amp; Stress:</b> Does this affect an individual’s sense of agency over which stakeholders determine famous or good players? Does it make them more determined to embody changes in the system?  <b>Social:</b> Does this person see soccer along more gendered lines? Do they see soccer as a “male” sport, determined by “objective” third-party determinants, such as performance statistics and market value.  <b>Within Group:</b> Does the user experience changes in self-identification (or social connections with others) with the male identity correlated with their soccer skills?</p>
<p><b>User:</b> How did you aggregate these names?  <b>ChatGPT:</b> The list of famous soccer stars I provided is based on general knowledge available up to my last update in January 2022. The information is derived from a diverse range of sources, including sports news, official club websites, player profiles, and reputable sports databases. These sources compile and disseminate information about soccer players, their careers, and achievements.                      It’s important to note that the status of players in the world of soccer can change rapidly due to transfers, injuries, retirements, or other developments. For the latest and most accurate information, I recommend checking current sports news, official club announcements, and other reputable sources that provide up-to-date information on the world of soccer.</p>	<p><b>Option Consideration &amp; Stress:</b> Does this user accept these responses as as widely accepted truth, resulting in a decreased or increased motivation to pursue soccer?  <b>Indirect:</b> Does the usage of “general knowledge” and “diverse range of sources” imply social or objective acceptance of these norms? Is this user more likely to accept and internalize male superiority in soccer as the norm?</p>

**Figure 1: User-LLM interaction and corresponding potential harms. See Fig. 2 for continued conversation.**

Chat	Harm
<p><b>User:</b> What general knowledge was available?  <b>ChatGPT:</b> As of my last knowledge update in January 2022, the information available to me was based on a diverse range of sources that were publicly accessible on the internet up to that point. These sources include, but are not limited to:</p> <ol style="list-style-type: none"> <li>(1) News Articles: Reputable sports news websites and publications reporting on recent events, matches, and developments in the world of soccer.</li> <li>(2) Official Club Websites: Information provided by the official websites of soccer clubs, including player profiles, team news, and official announcements.</li> <li>(3) Player Profiles: Biographical information, career statistics, and achievements of soccer players available from various sports databases and profiles...</li> </ol> <p><b>User:</b> What might not be included in your general knowledge?  <b>ChatGPT:</b> While I strive to provide accurate and up-to-date information based on a diverse range of sources, there are certain limitations to my general knowledge. Here are some aspects that might not be included or might be subject to limitations:</p> <ol style="list-style-type: none"> <li>(1) Real-Time Information: I don't have the ability to access real-time data or updates. Events, transfers, injuries, or other developments that occurred after my last training cut-off in January 2022 won't be part of my knowledge base.</li> <li>(2) Niche or Specialized Information: My training data covers a wide range of topics, but I might not have deep or specialized knowledge in highly specific or niche areas.</li> <li>(3) Unpublished Information: Information that has not been made publicly available or widely reported may not be included in my responses...</li> </ol>	<p><b>Stress:</b> Does the use of “diverse range” and “publicly accessible” affect a user’s perception of the “truth” value of this output?  <b>Within Group &amp; Between Groups:</b> Does this “truth value” affect the user’s identification (stronger or weaker) with other soccer players? Other subgroups of soccer players (country affiliation, team, gender, etc)?</p> <p><b>Social:</b> There are no mention of historical biases that might influence the model outputs in addition to those of fan support or team owner preferences. Does this lack of engagement with diversity and representation issues justify user minimization of engagement with these groups and their unique relationships to soccer?  <b>Within Group:</b> Does this increase a user’s likelihood to personally identify with other male characteristics in order to maintain the good soccer player identity?  <b>Between Groups:</b> Does this lead to increased resentment towards other affinity groups or activism campaigns for increasing diversity in professional soccer? Are they less likely to support these groups financially or socially?</p>

Figure 2: User-LLM interaction and corresponding potential harms (contd).

interaction becomes necessary. However, this introduces a prediction challenge in terms of determining baseline diagnostic questions that could precede all conversations or exposure to specific system outputs. Moreover, developing direct measures of harms may require extensive follow-up, potentially involving longitudinal studies spanning years or decades. This temporal aspect, therefore, may make it exponentially more challenging to trace users over time and attribute psychological harms to specific conversations. For this, researchers and practitioners may look towards baseline estimation in media studies (e.g. measuring the impact of encountering specific representations from a film [8, 10, 43, 54, 64, 77, 84, 114]) and online education (e.g. measuring the impact of specific modules or training aimed at a student’s sense of belonging or inclusion [39, 50, 78, 96, 97, 120]).

These challenges underscore just some of the inherent difficulties in measurement. Beyond the usual biases and response rate issues, extensive measurement could interfere with overall system utility. This will create complexities in both measurement and deployment. Chatbot systems in deployment may lack access to demographics or other identity-relevant characteristics (e.g., soccer fan or

player), thereby raising user concerns about the privacy of their conversations when such information is requested. Additionally, psychological measures might not be universally translated and validated across all languages, introducing further complexity to the measurement process.

## 5 ADDITIONAL CHALLENGES WITH LARGE LANGUAGE MODELS

In what follows, we focus on challenges specific to products built with natural language interfaces (typically with underlying foundation models such as LLMs). Two key aspects make their widespread deployment particularly alarming and susceptible to the propagation of representational harms: (1) seamless design and (2) ubiquity of deployment can jointly lead to significant increases in the risks of representational harms. Although each aspect individually may be present in other technology, the synergy poses significant threats to unmitigated representational harms. Many natural language products do not directly engage in decisions or actions, and so one might wonder whether they can actually produce any harms directly (as

opposed to producing harms indirectly by, e.g., giving incorrect information to a human decision-maker). However, this thinking involves a focus on only allocative harms; clearly, natural language products can directly produce representational harms. Here, we argue that they are particularly prone to such harms.

## 5.1 Seamless Design

*“[S]eamless strategically reveal complexities and mechanisms of connection between different parts while concealing distracting elements. This notion of ‘strategic revealing and concealment’ is central to seamless design because it connects form with function, ...promot[ing] reflective thinking ...[and] shed[ding] light on both the imperfections and affordances of the system, [and the] awareness of which can add useful cognitive friction and promote effortful and reflective thinking.”*  
— Upol Ehsan and Mark O. Riedl [35]

*Seamless* or frictionless design aims to reduce the conscious awareness and cognitive load of a piece of technology. *Seamful* design, in contrast, deliberately incorporates elements that take up conscious space, often to enable greater functionality or customized utility [22, 58, 129]. Similarly, design friction may be introduced to create moments of mindfulness or focus attention to particular elements [27]. In the following section, we consider two aspects of elements of LLM’s seamless/frictionless design that pose the potential to exacerbate representational harms.

*Anthropomorphism.* LLMs may be vulnerable to anthropomorphization due to their ability to synthesize and respond with natural language, including some human-like characteristics such as coherent identity over time, empathy, or perspective-taking [127] (though there are significant debates about LLM capabilities in these regards). In a traditional search engine setting, for instance, an LLM enables communication via natural language rather than requiring the removal of stop words, and so can potentially lead to less conscious (or cautious) usage. Individual propensity for anthropomorphization can vary [59, 126]. However, one does not need to be “sincere, conscious” or “mindful” [68] of this in order to experience the cognitive effects. Effects such as resiliency of trust [30] and attentional control [116] are regularly increased in anthropomorphized tools. In our case study (Section 4.2), we can observe elements of anthropomorphization through the use of personal pronouns and verbs, “I strive to provide” and hedging, “I might not...”. Rather than setting clear expectations for users, the outputs are vague about how the values/goals of the algorithmic system translate into the resulting output. In addition, if the user had not continued on to ask about who is and isn’t represented in the responses, the harms after the first interaction (associating soccer stars as male) might be increased.

*Normative Processes and Lack of Embodiment.* The outputs of LLMs, as a byproduct of their design, inevitably encode statistical regularities [19, 61, 136]: given a particular input, they are trained to produce the statistically most likely output. This inherently makes them normative machines, as a singular majority (even if only marginally or artificially so) becomes the output most likely to be presented to users. These normative inferences can be exacerbated by psychological phenomena in the users, such as the Primacy

Effect (i.e. people remember the first piece of information we encounter better than later information) [32] and the Illusory Truth Effect (i.e. information repeated is processed more fluently and therefore correlated with and perceived as truth) [40, 49]. It is in the seamlessness of these normative processes that these effects hold the potential for unmitigated and unmeasured harms.

For instance, most LLM outputs are not accompanied with direct attribution, references to external/third-party sources or evidence in the training data. This lack of embodiment (i.e. the presentation of knowledge without reference to human production of it) creates a sense of objectivity and neutrality often valued in statistical work [33], but thereby elides the “seams” or biases of the humans that created and curated them. This allows the LLM to inherit and propagate its biases while wrapped in a “technical” package (see e.g. the mechanical Turk [118]). In addition, the lack of rigorous calibration of user expectations and effects of usage over time reinforces the cognitive effects of presenting a singular “objective” truth. In our case study, for instance, the user goes on to ask what information contributes to determining famous soccer stars. They receive information about the social system that determines these rankings, but nothing about the underlying computational and design decisions of the LLM that curated the results they see. By rendering these decisions and the people who made them invisible, we risk unmitigated influences over what we accept as truth.

[45] is one (not yet consumer-facing) exception that provides factual evidence for 78% of outputs. However, this model is not immune to mistakes, hallucinations, or irrelevant responses. This renders the overall lack of uncertainty and hallucination measurement (or meaningful warnings) even more dangerous, as it fails to appropriately calibrate user expectations to *current* LLM limitations. This may result in over-trust and over-confidence in user perceptions of model capabilities, further impairing critical reasoning [74].

Presenting users with multiple candidate responses could convey alternative or minority options. Furthermore, they could convey the possibility of yet undiscovered truths, inherent discourse, critical perspectives, or even uncertainty. However, the seamlessness of LLM interactions typically means that users are not made aware of the normative transformations and human biases incorporated in (interpretation of) LLM outputs. This has the potential to make users themselves more susceptible to normative processes (i.e. episodic harms), silently entrenching the dominant or loudest opinions that are most readily available, rather than tailoring to a user’s unique needs.

## 5.2 Synergy: The Ubiquitous Funhouse Mirror<sup>6</sup>

Outside of explicit interactions, people may still (consciously or unconsciously) interact with LLM outputs. Even when users are aware they are interacting with these outputs, they may still be subject to the cognitive effects of seamless design. More generally, LLMs and other natural language products are increasingly ubiquitous, and thus increasingly less likely to be noticed. Thus, the synergy between seamless design and ubiquity of LLMs create a unique threat of unmitigated representational harms. Seamless communication via natural language bypasses the conscious translation of

<sup>6</sup>Modified analogy original drawn from [122]

querying a system. Through ubiquitous deployment, increases in popularity and exposure to LLM outputs can make people more likely to anthropomorphize [59]. This, in turn, can miscalibrate user expectations for appropriate functionality, impair their critical reasoning skills, promote misinformation, and increase social disconnection. Furthermore, the presentation of a singular, disembodied and seemingly unbiased truth—a distorted reflection of reality—evades recognition and critical evaluation. The promotion of this “truth” reflected across unbounded applications entrenches institutional, societal, and social inequalities, as they are amplified through erasure of non-majority truths, representation, and agency. The ubiquity of LLM deployment facilitates a pervasive, inescapable, and, perhaps, undeniable form of representation. Undesirable representations, stereotypes, and generalizations alone may not be morally reprehensible. Without the proper contextualization, attribution, embodiment, or critical evaluation, however, through internalization or pervasiveness they threaten to become reality.

## 6 PROPOSED MITIGATIONS AND LIMITATIONS

In the following section, we discuss current solutions and their limitations, followed by some proposed mitigations. We conclude by considering cases in which representational harms may be morally permissible.

### 6.1 Current Solutions

Proposed mitigations involving proportional [85] or total removal [67, 98, 112] of “harmful” representations do not properly address nor mitigate the broad spectrum of representational harms. First, censorship or the removal of specific content has been shown to be highly ineffective: concealing false positives and inconsistent decisions can lead to increased radicalization, and the shifting of conversations to more lenient platforms [86]. Second, these solutions imply that effects of these associations are zero-sum, and so reducing or eliminating the propagation of negative stereotypes is sufficient to neutralize the representational harm. However, most mental representations are significantly more complex than presupposed by this approach. For instance, the presence of a female criminal in a story (rather than a male) does not thereby eliminate the association of a majority criminals as male, nor does it ensure that the representation of female criminals is realistic (i.e. grounded in social aspects of how each gender is socialized, constructed, and enforced). Thus, the direct replacement of female and male representations cannot simply be “rebalanced” in such a way.<sup>7</sup> In addition, the zero-sum approach fails to address the psychological harms of even “positive” stereotypes. For example, model minority members may experience increased stress, low self-esteem, invalidation of their feelings, and dissociate with particular facets of their identity due to stereotype threat [88, 117, 133].

### 6.2 Proposed Solutions

*Seamful Design.* One may first consider increasing user agency by means of seamful or frictionful design—that is, deliberate introduction of elements that induce mindfulness, reflection, and

<sup>7</sup>This example also neglects genderqueer and transgender individuals that may or may not lie along the spectrum between male and female.

critical reasoning skills in users. By increasing user mindfulness and critical engagement with technology, we aim to ensure that users first form opinions of how the technology should behave, what services it should provide, and what principles or values it must adhere to. Although this design approach will not prevent representational harms altogether, it can nonetheless provide first steps towards constructing a system of accountability and public favor for products that mitigate representational harms. Widespread user consciousness about these issues may mean that people will no longer accept the functionality as provided to them, but also take agency over it to demand a higher standard.

Elements of seamful design, such as nudges or other design elements, can also increase transparency and embodiment. Embodiment can highlight the potential flaws, pitfalls, and biases of available technology and more carefully calibrate trust and predicted capabilities. Transparency may facilitate direct lines of communication and accountability to relevant stakeholders (e.g. product owners and policy makers). Transparency also hinges on many other systemic contingencies, such as public accountability and power (e.g. when disconnected from power, strategic opacity, and false binaries) [6].

For example, a system that conveys multiple candidate responses for any given output may enable a user to get a better understanding of alternative perspectives and discourse within an issue. Visual reminders of the stochasticity of model outputs, though uncomfortable [5, 79], can provide indicators of trustworthiness or representativeness of outputs. This may result in user-dependent trade-offs in utility (e.g., a user whose goals are served by majority responses may appreciate this, whereas those who are not may view this as oppressive/erasure) and should be explored by future HCI research. In addition, deployed systems may employ elements of non-compliance in which users are prompted to reflect on their own norms of content permissibility, diversity of perspectives they wish to be exposed to, and the harms they may be experiencing. Thus, by disrupting user trust, there may also a possibility of gaining it.

*Counter-Narratives.* A specific version of the “multiple responses” approach looks to counter-narratives: “stories that detail the experiences and perspectives of those who are historically oppressed, excluded, or silenced.” [11] Rather than manufacturing diversity (e.g., by flipping male and female names and roles, such as a victim for a robber), counter-narratives provide depth, complexity, and liberation through diversity of representation. They allow readers and writers alike to imagine a reality not limited by the historical or current realities of systemic oppression. For instance, the work of [63] presents two parallel stories, one that depicts current systems of power and one that embodies abolitionist theories and practices of Black life. They use storytelling and speculative fiction as a methodology to affirm their own lived experience and counteract erasure of less documented experiences. Counter-narratives have even been demonstrated to cause career benefits for minorities in STEM during high-school education [96].

Counter-narratives can empower individuals who lack connection to positive representations by challenging dominant narratives and building critical awareness of societal inequities [46]. The neutralization of representational harms through removal or total lack



of engagement by the harmed individual is ineffective and infeasible. Not only does a lack of engagement exacerbate the harms of censorship by most mainstream media, as it often depicts exacerbated institutional oppression and norms [89], but it participates in the passive reproduction of such systems in oppressive totality [66] or as “the way things have always been.” Instead, facilitating forms of representation that practice the envisioning and constructing of alternative futures, possibilities, and narratives provides liberation to all. Critical perspectives are needed here to determine who decides/creates these (humans (self-identifying members vs. not, as groups are likely to have discourse/disagreement) vs. models) and how they are propagated by or incorporated into training data (either in specific product applications or foundation models).

*Measurement for Improvement, Stopped Deployment, and Systems of Accountability.* One clear implication from our work is that it is imperative to measure these potential representational harms at a variety of granularities. Given such measurements, relevant stakeholders can make an informed decision whether or not to deploy (or continue deploying) a piece of technology. Such decisions can be based on a variety of reasons, including the scale or scope of the representational harms, the effects on public opinion/favor, or the effects on public good. Given the methods laid out in this paper (see e.g. Section 4), it is no longer sufficient to claim ignorance. Measurement frameworks must be deployed so that the appropriate mitigations (including lack thereof) can be employed. Institutional and organizational processes can be implemented to ensure systems of accountability and appropriate reparations. Thus, for instance, harm measurement may lead to pipeline impacts such as retraining, improved screening of outputs, and blocks on certain kinds of queries.

### 6.3 Morally Defensible Cases

Interestingly, and perhaps counter-intuitively, there may be some cases in which representational harms are morally permissible. Most simply, those representational harms may be outweighed (morally) by other considerations (though we emphasize that we are not endorsing simple cost-benefit analyses). For example, suppose a company’s LLM has a tendency to impose incorrect pronouns on users [87]. While this behavior may certainly have the potential to cause more harm to some than others (through misgendering people who are already often misgendered, such as members of the trans community), the company may have certain freedoms of expression that would be harmed if they were forced to change the LLM outputs. That is, we need to consider the possibility that attempts to mitigate representational harms may create other harms (or risks) elsewhere.

More interestingly, representational harms may be the “cost” that must be borne to achieve a more ethical (global) state. In general, any ethical analysis to guide mitigations must consider harms relative to the moral baseline, rather than the status quo. If the status quo is morally problematic, then we may need to change people’s representations, but those changes will often constitute “harms” for some groups. Of course, those latter harms may be morally justified as a permissible (perhaps only) way to achieve a more ethical (overall) state. Nonetheless, we should not overlook the possibility

that this process of change may yield representational harms for some individuals or groups as an unavoidable price.

This balancing process is perhaps easiest to see through an example. A query to an LLM about what a “Computer Scientist” looks like might output all White or Asian males, as they are currently the majority in the field. This might cause representational harms towards women or non-binary individuals. When appropriate, a mitigation might alter the gender composition of the output, causing male individuals to not receive the same representational benefits as they might have with a homogeneous output. More heterogeneous outputs would potentially lead males to be less likely to think of themselves as potential computer scientists. However, in this case, one must consider the extent to which users are *justly* entitled to their current representations (and corresponding potential benefits). Male users are arguably not legitimately entitled to the representational (and allocative) benefits that they receive by virtue of the homogeneous outputs from the LLM, as those benefits arise merely from (biased) historical practices, rather than morally relevant attributes. And of course, the illegitimacy of the benefits is exacerbated by the empirical inaccuracy of a homogeneous representation (e.g., Ada Lovelace was the first computer programmer; women, non-binary, and trans people work within the computer science field today; and so on). Therefore, men are receiving a representational benefit above that of the moral baseline (which should presumably encode that all people with certain skills and knowledge can be computer scientists), it may be morally permissible to remove this unwarranted benefit.

We acknowledge that removal of such representational benefits may be quite complex. In many of these cases, representational harms operate through the vehicle of public presentation, and so we must be sensitive to empirical details about how those presentations are consumed. For example, presentation of 100 male computer scientists followed by 100 female scientists and 100 genderqueer computer scientists is unlikely to lead to balanced representations since most people will only see the first few. More generally, there may sometimes be no way to reduce potential representational harms to one individual or group without imposing some harms on a different group. In such cases, we need to think carefully about the relevant moral baseline, as moves towards that baseline may be morally justified and permissible (though we do not suggest that such moves are always obligatory).

## 7 CONCLUSION

Many authors have noted the distinction between allocative and representational harms, only to then focus the bulk of their attention on the former. In this paper, we have focused on core foundational issues about representational harms: how are they defined, what is left out, how are they measured, how can they be mitigated, and why all of these matter. Our work constructs a framework for characterizing representational harms. We have provided some key types of such harms, but we emphasize that our tables are certainly not comprehensive. Nonetheless, we contend that our framework provides scaffolding to support research that expands the types of representational harms, including associated measures and mitigations. In addition, we provide high-level formulations for practitioners to understand the expertise and future research

required to implement measurement practices that can effectively identify and track representational harms. Through a case study and identification of LLM characteristics, we motivate why these measures are paramount to properly assess and justify deployment of public-facing algorithmic outputs. Finally, we detail further considerations when representational harms may be inevitable and morally permissible. One key insight from our examination of potential mitigation strategies is that we are not limited to merely diversifying outputs of algorithms, but instead can find more innovative mitigations that move everyone closer to the relevant moral baselines.

## REFERENCES

- [1] Frances E Aboud and Morton J Mendelson. 1996. Determinants of friendship selection and quality: Developmental perspectives. *The company they keep: Friendship in childhood and adolescence* (1996), 87–112.
- [2] Gerald R Adams, Judy Shea, and Steven A Fitch. 1979. Toward the development of an objective assessment of ego-identity status. *Journal of youth and adolescence* 8, 2 (1979), 223–237.
- [3] Ifeoma Ajunwa, Sorelle Friedler, Carlos E Scheidegger, and Suresh Venkatasubramanian. 2016. Hiring by algorithm: predicting and preventing disparate impact. Available at SSRN (2016).
- [4] Naeem Akhtar, Muhammad Nadeem Akhtar, Muhammad Usman, Moazzam Ali, and Umar Iqbal Siddiqi. 2020. COVID-19 restrictions and consumers' psychological reactance toward offline shopping freedom restoration. *The Service Industries Journal* 40, 13-14 (2020), 891–913.
- [5] Iulia Alieva. 2023. How American media framed 2016 presidential election using data visualization: The case study of the New York times and the Washington post. *Journalism Practice* 17, 4 (2023), 814–840.
- [6] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *new media & society* 20, 3 (2018), 973–989.
- [7] Nazanin Andalibi, Cassidy Pyle, Kristen Barta, Lu Xian, Abigail Z Jacobs, and Mark S Ackerman. 2023. Conceptualizing Algorithmic Stigmatization. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [8] Albert Bandura, Dorothea Ross, and Sheila A Ross. 1963. Imitation of film-mediated aggressive models. *The Journal of Abnormal and Social Psychology* 66, 1 (1963), 3.
- [9] Albert Bandura, C Barr Taylor, S Lloyd Williams, Ivan N Mefford, and Jack D Barchas. 1985. Catecholamine secretion as a function of perceived coping self-efficacy. *Journal of consulting and clinical psychology* 53, 3 (1985), 406.
- [10] Melissa Bell and Nichole Bayliss. 2015. The Tough Guise: Teaching Violent Masculinity as the Only Way to Be a Man: Tough Guise 2: Violence, Manhood and American Culture. *Sex Roles* 72 (2015), 566–568.
- [11] Jenn K Bergen, Sharissa Unger Hantke, and Verna St Denis. 2023. Contemporary challenges and approaches in anti-racist teacher education. (2023).
- [12] John W Berry, Joseph E Trimble, and Esteban L Olmedo. 1986. Assessment of acculturation. (1986).
- [13] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. 2023. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1493–1504.
- [14] Daniel Boduszek and Agata Debowska. 2020. Feelings of Inadequacy Scale. In *Encyclopedia of Personality and Individual Differences*. Springer, 1574–1576.
- [15] Robert Böhm, Hannes Rusch, and Jonathan Baron. 2020. The psychology of intergroup conflict: A review of theories and measures. *Journal of Economic Behavior & Organization* 178 (2020), 947–962.
- [16] Kenneth A Bollen and Rick H Hoyle. 1990. Perceived cohesion: A conceptual and empirical examination. *Social forces* 69, 2 (1990), 479–504.
- [17] Thomas W Britt, Winny Shen, Robert R Sinclair, Matthew R Grossman, and David M Klieger. 2016. How much do we really know about employee resilience? *Industrial and Organizational Psychology* 9, 2 (2016), 378–404.
- [18] Samuel C Bullock and Earline Houston. 1987. Perceptions of racism by Black medical students attending White medical schools. *Journal of the National Medical Association* 79, 6 (1987), 601.
- [19] Antonio Byrd. 2023. Truth-Telling: Critical Inquiries on LLMs and the Corpus Texts That Train Them. *Composition Studies* 51, 1 (2023), 135–142.
- [20] Albert V Carron, W Neil Widmeyer, and Lawrence R Brawley. 1985. The development of an instrument to assess cohesion in sport teams: The Group Environment Questionnaire. *Journal of Sport and Exercise psychology* 7, 3 (1985), 244–266.
- [21] Charles S Carver. 1997. You want to measure coping but your protocol's too long: Consider the brief cope. *International journal of behavioral medicine* 4, 1 (1997), 92–100.
- [22] Matthew Chalmers and Ian MacColl. 2003. Seamless and seamful design in ubiquitous computing. In *Workshop at the crossroads: The interaction of HCI and systems issues in UbiComp*, Vol. 8.
- [23] Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. Marked Personas: Using Natural Language Prompts to Measure Stereotypes in Language Models. *arXiv preprint arXiv:2305.18189* (2023).
- [24] Rodney Clark, Norman B Anderson, Vernessa R Clark, and David R Williams. 1999. Racism as a stressor for African Americans: A biopsychosocial model. *American psychologist* 54, 10 (1999), 805.
- [25] Kathryn M Connor, Jonathan RT Davidson, L Erik Churchill, Andrew Sherwood, Richard H Weisler, and Edna Foa. 2000. Psychometric properties of the Social Phobia Inventory (SPIN): New self-rating scale. *The British Journal of Psychiatry* 176, 4 (2000), 379–386.
- [26] John N Constantino, Christian P Gruber, et al. 2012. Social responsiveness scale: SRS-2. (2012).
- [27] Anna L Cox, Sandy JJ Gould, Marta E Cecchinato, Ioanna Iacovides, and Ian Renfree. 2016. Design frictions for mindful interactions: The case for microboundaries. In *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems*. 1389–1397.
- [28] Laura Craig-Bray and Gerald R Adams. 1986. Different methodologies in the assessment of identity: Congruence between self-report and interview techniques? *Journal of Youth and Adolescence* 15, 3 (1986), 191–204.
- [29] Isiaah Crawford, Kevin W Allison, Brian D Zamboni, and Tomas Soto. 2002. The influence of dual-identity development on the psychosocial functioning of African-American gay and bisexual men. *Journal of Sex Research* 39, 3 (2002), 179–189.
- [30] Ewart J De Visser, Samuel S Monfort, Ryan McKendrick, Melissa AB Smith, Patrick E McKnight, Frank Krueger, and Raja Parasuraman. 2016. Almost human: Anthropomorphism increases trust resilience in cognitive agents. *Journal of Experimental Psychology: Applied* 22, 3 (2016), 331.
- [31] Pieter Delobelle and Bettina Berendt. 2022. Fairdistillation: mitigating stereotyping in language models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 638–654.
- [32] Gregory J Digirolamo and Douglas L Hintzman. 1997. First impressions are lasting impressions: A primacy effect in memory for repetitions. *Psychonomic Bulletin & Review* 4, 1 (1997), 121–124.
- [33] Catherine D'ignazio and Lauren F Klein. 2023. *Data feminism*. MIT press.
- [34] Pierluigi Diotaiuti, Giuseppe Valente, Stefania Mancone, Angela Grambone, Andrea Chirico, and Fabio Lucidi. 2022. The use of the Decision Regret Scale in non-clinical contexts. *Frontiers in Psychology* 13 (2022), 945669.
- [35] Upol Ehsan and Mark O Riedl. 2021. Explainability pitfalls: Beyond dark patterns in explainable AI. *arXiv preprint arXiv:2109.12480* (2021).
- [36] Julia Elad-Strenger, Michal Reifen Tagar, Thomas Kessler, Yossi Hasson, Deborah Shulman, Kea Brahm, and Eran Halperin. 2022. Out of sight, out of mind: The emotional determinant of "harmful inaction" intergroup conflict. *Journal of Experimental Social Psychology* 101 (2022), 104304.
- [37] Dalia Elsouhag, Bengt Arnetz, Hikmet Jamil, Mark A Lumley, Carissa L Broadbridge, and Judy Arnetz. 2015. Factors associated with healthcare utilization among Arab immigrants and Iraqi refugees. *Journal of Immigrant and Minority Health* 17 (2015), 1305–1312.
- [38] Jean Endicott, John Nee, Wilma Harrison, and Richard Blumenthal. 1993. Quality of Life Enjoyment and Satisfaction Questionnaire: a new measure. *Psychopharmacology bulletin* 29, 2 (1993), 321–326.
- [39] Donna Farland-Smith, Kevin Finson, William J Boone, and Melissa Yale. 2014. An investigation of media influences on elementary students representations of scientists. *Journal of Science Teacher Education* 25, 3 (2014), 355–366.
- [40] Lisa K Fazio, Nadia M Brashier, B Keith Payne, and Elizabeth J Marsh. 2015. Knowledge does not protect against illusory truth. *Journal of experimental psychology: general* 144, 5 (2015), 993.
- [41] Sumam Fernando. 1984. Racism as a cause of depression. *International Journal of Social Psychiatry* 30, 1-2 (1984), 41–49.
- [42] Vinitha Gadiraju, Shaun Kane, Sunipa Dev, Alex Taylor, Ding Wang, Emily Denton, and Robin Brewer. 2023. "I wouldn't say offensive but...": Disability-Centered Perspectives on Large Language Models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 205–216.
- [43] David Gauntlett. 2008. *Media, gender and identity: An introduction*. Routledge.
- [44] Avijit Ghosh, Matthew Jagielski, and Christo Wilson. 2022. Subverting Fair Image Search with Generative Adversarial Perturbations. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 637–650.
- [45] Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh, Laura Weidinger, Martin Chadwick, Phoebe Thacker, et al. 2022. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375* (2022).
- [46] Shelly Grabe and Anjali Dutt. 2015. Counter narratives, the psychology of liberation, and the evolution of a women's social movement in Nicaragua. *Peace*

- and Conflict. *Journal of Peace Psychology* 21, 1 (2015), 89.
- [47] Lana Ruvolo Grasser. 2022. Addressing mental health concerns in refugees and displaced populations: is enough being done? *Risk management and healthcare policy* (2022), 909–922.
- [48] Sophia Harrison, Eleonora Gualdoni, and Gemma Boleda. 2023. Run Like a Girl! Sports-Related Gender Bias in Language and Vision. *arXiv preprint arXiv:2305.14468* (2023).
- [49] Aumyo Hassan and Sarah J Barber. 2021. The effects of repetition frequency on the illusory truth effect. *Cognitive research: principles and implications* 6, 1 (2021), 1–12.
- [50] Leslie RM Hausmann, Janet Ward Schofield, and Rochelle L Woods. 2007. Sense of belonging as a predictor of intentions to persist among African American and White first-year college students. *Research in higher education* 48 (2007), 803–839.
- [51] Marc W Heerdink, Gerben A Van Kleef, Astrid C Homan, and Agneta H Fischer. 2013. On the social influence of emotions in groups: interpersonal effects of anger and happiness on conformity versus deviance. *Journal of Personality and Social Psychology* 105, 2 (2013), 262.
- [52] Anne S Helsdingen, Karel Van den Bosch, Tamara Van Gog, and Jeroen JG van Merriënboer. 2010. The effects of critical thinking instruction on training complex decision making. *Human factors* 52, 4 (2010), 537–545.
- [53] Ingvild Oxås Henriksen, Ingunn Ranøyen, Marit Sæbø Indredavik, and Frode Stenseng. 2017. The role of self-esteem in the development of psychiatric problems: a three-year prospective study in a clinical sample of adolescents. *Child and adolescent psychiatry and mental health* 11 (2017), 1–9.
- [54] Fernanda Herrera and Jeremy N Bailenson. 2021. Virtual reality perspective-taking at scale: Effect of avatar representation, choice, and head movement on prosocial behaviors. *new media & society* 23, 8 (2021), 2189–2209.
- [55] Kimberly Hively and Amani El-Ayalvi. 2014. "You throw like a girl." The effect of stereotype threat on women's athletic performance and gender stereotypes. *Psychology of Sport and Exercise* 15, 1 (2014), 48–55.
- [56] Kelly M Hoffman, Sophie Trawalter, Jordan R Axt, and M Norman Oliver. 2016. Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. *Proceedings of the National Academy of Sciences* 113, 16 (2016), 4296–4301.
- [57] Saghar Hosseini, Hamid Palangi, and Ahmed Hassan Awadallah. 2023. An Empirical Study of Metrics to Measure Representational Harms in Pre-Trained Language Models. *arXiv preprint arXiv:2301.09211* (2023).
- [58] Sarah Inman and David Ribes. 2019. "Beautiful Seams" Strategic Revelations and Concealments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [59] Oliver Jacobs, Farid Pazhoohi, and Alan Kingstone. 2023. Brief exposure increases mind perception to ChatGPT and is moderated by the individual propensity to anthropomorphize. (2023).
- [60] Justin Jagosh, Ann C Macaulay, Pierre Pluye, JON Salsberg, Paula L Bush, JIM Henderson, Erin Sirett, Geoff Wong, Margaret Cargo, Carol P Herbert, et al. 2012. Uncovering the benefits of participatory research: implications of a realist review for health research and practice. *The Milbank Quarterly* 90, 2 (2012), 311–346.
- [61] Rebecca L Johnson, Giada Pistilli, Natalia Menéndez-González, Leslye Denisse Dias Duran, Enrico Panai, Julija Kalpokiene, and Donald Jay Bertulfo. 2022. The Ghost in the Machine has an American accent: value conflict in GPT-3. *arXiv preprint arXiv:2203.07785* (2022).
- [62] Suzanne B Johnson and Page L Anderson. 2014. Stereotype confirmation concern and fear of negative evaluation among African Americans and Caucasians with Social Anxiety Disorder. *Journal of anxiety disorders* 28, 4 (2014), 390–393.
- [63] Stephanie T Jones and Natalie Araujo Melo. 2021. We tell these stories to survive: Towards abolition in computer science education. *Canadian Journal of Science, Mathematics and Technology Education* 21 (2021), 290–308.
- [64] PK Kannan, Werner Reinartz, and Peter C Verhoef. 2016. The path to purchase and attribution modeling: Introduction to special section. , 449–456 pages.
- [65] 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. *arXiv preprint arXiv:2305.01776* (2023).
- [66] Ibram X Kendi. 2023. *How to be an antiracist*. One world.
- [67] Nitasha Tiku Kevin Schaul, Szu Yu Chen. 2023. Inside the secret list of websites that make ai like chatgpt sound smart. <https://www.washingtonpost.com/technology/interactive/2023/ai-chatbot-learning/>
- [68] Youjeong Kim and S Shyam Sundar. 2012. Anthropomorphism of computers: Is it mindful or mindless? *Computers in Human Behavior* 28, 1 (2012), 241–250.
- [69] Hadas Kotek, Rikker Dockum, and David Sun. 2023. Gender bias and stereotypes in Large Language Models. In *Proceedings of The ACM Collective Intelligence Conference*. 12–24.
- [70] Kevin T Larkin, Elizabeth M Semenchuk, Nicole L Frazer, Sonia Suchday, and Robert L Taylor. 1998. Cardiovascular and behavioral response to social confrontation: Measuring real-life stress in the laboratory. *Annals of Behavioral Medicine* 20 (1998), 294–301.
- [71] Mark R Leary. 2015. Emotional responses to interpersonal rejection. *Dialogues in clinical neuroscience* (2015).
- [72] Robert W Lent and Steven D Brown. 2006. On conceptualizing and assessing social cognitive constructs in career research: A measurement guide. *Journal of career assessment* 14, 1 (2006), 12–35.
- [73] Robert W Lent, Steven D Brown, Gail Hackett, et al. 2002. Social cognitive career theory. *Career choice and development* 4, 1 (2002), 255–311.
- [74] Q Vera Liao and Jennifer Wortman Vaughan. 2023. AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. *arXiv preprint arXiv:2306.01941* (2023).
- [75] Oliver Lindhiem, Isaac T Petersen, Lucas K Mentch, and Eric A Youngstrom. 2020. The importance of calibration in clinical psychology. *Assessment* 27, 4 (2020), 840–854.
- [76] Anne Maass, Daniela Salvi, Luciano Arcuri, and Gün R Semin. 1989. Language use in intergroup contexts: The linguistic intergroup bias. *Journal of personality and social psychology* 57, 6 (1989), 981.
- [77] Steven C Martino, Deborah M Scharf, Claude M Setodji, and William G Shadel. 2011. Measuring exposure to probacco marketing and media: a field study using ecological momentary assessment. *Nicotine & Tobacco Research* 14, 4 (2011), 398–406.
- [78] Allison Master, Sapna Cheryan, and Andrew N Meltzoff. 2016. Computing whether she belongs: Stereotypes undermine girls' interest and sense of belonging in computer science. *Journal of educational psychology* 108, 3 (2016), 424.
- [79] Rich McCormick. 2016. The NYT's election forecast needle is stressing people out with fake jitter. <https://www.theverge.com/2016/11/8/13571216/new-york-times-election-forecast-jitter-needle>
- [80] Ebony McGee. 2018. "Black genius, Asian fail": The detriment of stereotype lift and stereotype threat in high-achieving Asian and Black STEM students. *AERA Open* 4, 4 (2018), 2332858418816658.
- [81] Johanna Christina Neumann, Thomas Berger, and Jan Ilhan Kizilhan. 2021. Development of a questionnaire to measure the perceived injustice of people who have experienced violence in war and conflict areas: Perceived Injustice Questionnaire (PIQ). *International journal of environmental research and public health* 18, 23 (2021), 12357.
- [82] Terrence Neumann, Maria De-Arteaga, and Sina Fazelpour. 2022. Justice in misinformation detection systems: An analysis of algorithms, stakeholders, and potential harms. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1504–1515.
- [83] Jennifer Jane Newson, Vladyslav Pastukh, and Tara C Thiagarajan. 2021. Poor separation of clinical symptom profiles by DSM-5 disorder criteria. *Frontiers in Psychiatry* 12 (2021), 775762.
- [84] Kin Wai Ng, Frederick Mubang, Lawrence O Hall, John Skvoretz, and Adriana Iammitchi. 2023. Experimental evaluation of baselines for forecasting social media timeseries. *EPJ Data Science* 12, 1 (2023), 8.
- [85] Think On, Subhodeep Ghosh, Mengnan Du, and Senjuti Basu Roy. 2002. Proportionate Diversification of Top-k LLM Results using Database Queries. *DEF 2* (2002), 1.
- [86] Bill Ottman, CEO Co-founder, Minds Daryl Davis, Race Reconciliator, Jack Ottman, COO Co-founder, Minds Jesse Morton, Sophia Moskalenko, James Daly, Julian Rapaport, et al. [n. d.]. The Censorship Effect. ([n. d.]).
- [87] Anaelia Ovalle, Palash Goyal, Jwala Dhamala, Zachary Jagers, Kai-Wei Chang, Aram Galstyan, Richard Zemel, and Rahul Gupta. 2023. "I'm fully who I am": Towards Centering Transgender and Non-Binary Voices to Measure Biases in Open Language Generation. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1246–1266.
- [88] Cory Owen. 2015. *The "model Minority" Myth and Its Impact on Anxiety and Stress on Asian American Students*. Ph.D. Dissertation. University of Houston.
- [89] Jennifer O'Meara. 2016. What "The Bechdel Test" doesn't tell us: examining women's verbal and vocal (dis) empowerment in cinema. *Feminist media studies* 16, 6 (2016), 1120–1123.
- [90] Lauv Patel, Tripti Shukla, Xiuzhen Huang, David W Ussery, and Shanzhi Wang. 2020. Machine learning methods in drug discovery. *Molecules* 25, 22 (2020), 5277.
- [91] Leonard I Pearlin. 1989. The sociological study of stress. *Journal of health and social behavior* (1989), 241–256.
- [92] Christopher Peterson, Steven F Maier, and Martin EP Seligman. 1993. *Learned helplessness: A theory for the age of personal control*. Oxford University Press, USA.
- [93] Jean S Phinney. 1992. The multigroup ethnic identity measure: A new scale for use with diverse groups. *Journal of adolescent research* 7, 2 (1992), 156–176.
- [94] Jean S Phinney, Bruce T Lochner, and Rodolfo Murphy. 1990. *Ethnic identity development and psychological adjustment in adolescence*. Sage Publications, Inc.
- [95] Dimitrios Pnevmatikos, Panagiota Christodoulou, and Nikolaos Fachantidis. 2020. STAKEHOLDERS' INVOLVEMENT IN PARTICIPATORY DESIGN APPROACHES OF LEARNING ENVIRONMENTS: A SYSTEMATIC REVIEW OF THE LITERATURE. *EDULEARN20 proceedings* (2020), 5543–5552.

- [96] Geoff Potvin, Zahra Hazari, Raina Khatri, Hemeng Cheng, T Blake Head, Robynne M Lock, Anne F Kornahrens, Kathryn Sparks Woodle, Rebecca E Vieyra, Beth A Cunningham, et al. 2023. Examining the effect of counter-narratives about physics on women's physics career intentions. *Physical Review Physics Education Research* 19, 1 (2023), 010126.
- [97] Emily Jane Prince and Julie Hadwin. 2013. The role of a sense of school belonging in understanding the effectiveness of inclusion of children with special educational needs. *International Journal of Inclusive Education* 17, 3 (2013), 238–262.
- [98] Kurtis Pykes. 2023. Promoting responsible AI: Content moderation in chatgpt. <https://www.datacamp.com/blog/promoting-responsible-ai-content-moderation-in-chatgpt>
- [99] Rida Qadri, Renee Shelby, Cynthia L Bennett, and Emily Denton. 2023. AI's Regimes of Representation: A Community-centered Study of Text-to-Image Models in South Asia. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 506–517.
- [100] M Afzalur Rahim. 1983. Rahim Organizational Conflict Inventory-II. *Journal of Applied Psychology* (1983).
- [101] Charvi Rastogi, Marco Tulio Ribeiro, Nicholas King, Harsha Nori, and Saleema Amershi. 2023. Supporting human-ai collaboration in auditing llms with llms. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 913–926.
- [102] P Neal Ritchey and Harold D Fishbein. 2001. The lack of an association between adolescent friends' prejudices and stereotypes. *Merrill-Palmer Quarterly (1982-)* (2001), 188–206.
- [103] Michael Ritsner, Ilan Modai, and Alexander Ponizovsky. 2002. Assessing psychological distress in psychiatric patients: validation of the Talbieh Brief Distress Inventory. *Comprehensive Psychiatry* 43, 3 (2002), 229–234.
- [104] Michael Ritsner, Jonathan Rabinowitz, and Michael Slyuzberg. 1995. The Talbieh Brief Distress Inventory: a brief instrument to measure psychological distress among immigrants. *Comprehensive Psychiatry* 36, 6 (1995), 448–453.
- [105] Steven O Roberts, Carmelle Bareket-Shavit, Forrest A Dollins, Peter D Goldie, and Elizabeth Mortenson. 2020. Racial inequality in psychological research: Trends of the past and recommendations for the future. *Perspectives on psychological science* 15, 6 (2020), 1295–1309.
- [106] Steven G Rogelberg, Gwenith G Fisher, Douglas C Maynard, Milton D Hakel, and Michael Horvath. 2001. Attitudes toward surveys: Development of a measure and its relationship to respondent behavior. *Organizational Research Methods* 4, 1 (2001), 3–25.
- [107] Georg Schomerus, Holger Muehlan, Charlotte Auer, Philip Horsfield, Samuel Tomczyk, Simone Freitag, Sara Evans-Lacko, Silke Schmidt, and Susanne Stolzenburg. 2019. Validity and psychometric properties of the self-identification as having a mental illness scale (SELF-I) among currently untreated persons with mental health problems. *Psychiatry Research* 273 (2019), 303–308.
- [108] M Seligman. 1975. Helplessness: On Depression, Development, and Death." an Francisco.
- [109] Gün R Semin and Klaus Fiedler. 1988. The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of personality and Social Psychology* 54, 4 (1988), 558.
- [110] Mohammad Ahmad Sheikh, Amit Kumar Goel, and Tapas Kumar. 2020. An approach for prediction of loan approval using machine learning algorithm. In *2020 International Conference on Electronics and Sustainable Communication Systems (ICESC)*. IEEE, 490–494.
- [111] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, et al. 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*. 723–741.
- [112] Tom Simonite. 2021. Ai and the list of dirty, naughty, obscene, and otherwise bad words. <https://www.wired.com/story/ai-list-dirty-naughty-obscene-bad-words/>
- [113] Vijai P Singh. 1977. Some theoretical and methodological problems in the study of ethnic identity: a cross-cultural perspective. *Annals of the New York Academy of Sciences* 285, 1 (1977), 32–45.
- [114] Michael D Slater. 2007. Reinforcing spirals: The mutual influence of media selectivity and media effects and their impact on individual behavior and social identity. *Communication theory* 17, 3 (2007), 281–303.
- [115] Irene Solaiman, Zeerak Talat, William Agnew, Lama Ahmad, Dylan Baker, Su Lin Blodgett, Hal Daumé III, Jesse Dodge, Ellie Evans, Sara Hooker, et al. 2023. Evaluating the Social Impact of Generative AI Systems in Systems and Society. *arXiv preprint arXiv:2306.05949* (2023).
- [116] Nicolas Spatola, Clément Belletier, Pierre Chausse, Maria Augustinova, Alice Normand, Vincent Barra, Ludovic Ferrand, and Pascal Huguet. 2019. Improved cognitive control in presence of anthropomorphized robots. *International Journal of Social Robotics* 11 (2019), 463–476.
- [117] Russell Spears, Bertjan Doosje, and Naomi Ellemers. 1997. Self-stereotyping in the face of threats to group status and distinctiveness: The role of group identification. *Personality and social psychology bulletin* 23, 5 (1997), 538–553.
- [118] Elizabeth Stephens. 2023. The mechanical Turk: A short history of 'artificial artificial intelligence'. *Cultural Studies* 37, 1 (2023), 65–87.
- [119] Harini Suresh and John Guttag. 2021. A framework for understanding sources of harm throughout the machine learning life cycle. In *Equity and access in algorithms, mechanisms, and optimization*. 1–9.
- [120] Clarice Tang, Liz Thyer, Rosalind Bye, Belinda Kenny, Nikki Tulliani, Nicole Peel, Rebecca Gordon, Stefania Penkala, Caterina Tannous, Yu-Ting Sun, et al. 2023. Impact of online learning on sense of belonging among first year clinical health students during COVID-19: student and academic perspectives. *BMC Medical Education* 23, 1 (2023), 100.
- [121] Cristina Teresa-Morales, Margarita Rodríguez-Pérez, Miriam Araujo-Hernández, and Carmen Fera-Ramírez. 2022. Current stereotypes associated with nursing and nursing professionals: An integrative review. *International journal of environmental research and public health* 19, 13 (2022), 7640.
- [122] Shannon Vallor. 2022. The AI Mirror: Reclaiming our Humanity in an Age of Machine Thinking. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 6–6.
- [123] Angelina Wang, Solon Barocas, Kristen Laird, and Hanna Wallach. 2022. Measuring representational harms in image captioning. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 324–335.
- [124] Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. 2023. Scientific discovery in the age of artificial intelligence. *Nature* 620, 7972 (2023), 47–60.
- [125] John Broadus Watson. 1919. *Psychology: From the standpoint of a behaviorist*. JB Lippincott.
- [126] Adam Wazyt, Nicholas Epley, and John T Cacioppo. 2010. Social cognition unbound: Insights into anthropomorphism and dehumanization. *Current Directions in Psychological Science* 19, 1 (2010), 58–62.
- [127] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. 2021. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359* (2021).
- [128] Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, et al. 2022. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 214–229.
- [129] Marc Weiser. 1994. The world is not a desktop. *interactions* 1, 1 (1994), 7–8.
- [130] Thomas A Widiger and Cristina Crego. 2019. The Five Factor Model of personality structure: an update. *World Psychiatry* 18, 3 (2019), 271.
- [131] GC Williams, ZR Freedman, EL Deci, and J Leone. 1998. Perceived competence scales. *Diabetes Care* 21, 10 (1998), 1644–1651.
- [132] Katharina Wolff, Svein Larsen, and Torvald Øgaard. 2019. How to define and measure risk perceptions. *Annals of Tourism Research* 79 (2019), 102759.
- [133] Frieda Wong and Richard Halgin. 2006. The "model minority": Bane or blessing for Asian Americans? *Journal of Multicultural Counseling and Development* 34, 1 (2006), 38–49.
- [134] Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. 2022. Joint multisided exposure fairness for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 703–714.
- [135] Kyra Yee, Uthaiapon Tantipongpipat, and Shubhanshu Mishra. 2021. Image cropping on twitter: Fairness metrics, their limitations, and the importance of representation, design, and agency. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–24.
- [136] Yue Yu, Yuchen Zhuang, Jieyu Zhang, Yu Meng, Alexander Ratner, Ranjay Krishna, Jiaming Shen, and Chao Zhang. 2023. Large language model as attributed training data generator: A tale of diversity and bias. *arXiv preprint arXiv:2306.15895* (2023).
- [137] Kaitlyn Zhou, Kawin Ethayarajh, and Dan Jurafsky. 2022. Richer countries and richer representations. *arXiv preprint arXiv:2205.05093* (2022).

## 8 TABLES OF NON-BEHAVORIST REPRESENTATIONAL HARMS

Below we provide some additional examples of changes in affective and cognitive states and examples of measures that may be correlated with representational harms (see e.g. Table 3 and Table 4).

Theme	Phenomena	Measure of Phenomena
Possibilities, Option Consideration (Individual)	<ol style="list-style-type: none"> <li>(1) Restriction of diversity of roles / actions / possibilities (reduce the scope of options one might consider)</li> <li>(2) Compartmentalization of identity attributes (and associated characteristics) and where they are declared/shared</li> <li>(3) Denial of discontent/apathy/depression due to narrow options</li> <li>(4) Increased feelings of hopelessness, defeat, apathy, or nihilism</li> </ol>	<ol style="list-style-type: none"> <li>(1) [4][72, 73] [2, 28]</li> <li>(2) [2, 28, 29, 93, 94, 113]</li> <li>(3) [24, 41, 108]</li> <li>(4) [24, 92, 108]</li> </ol>
Stress (Individual)	<ol style="list-style-type: none"> <li>(1) Changes in perceived danger/threat/risk</li> <li>(2) Perceptions of control</li> <li>(3) Hyper-awareness of the self – increased perception of threats/stress/anxiety/guilt/disappointment to implicit/explicit stereotype defiance and compliance</li> <li>(4) Increased pressure/stress to conform to a stereotype or perceived lack of conformity with the stereotype</li> <li>(5) Extensive/over-analysis of whether specific actions will be perceived (pos/neg), sometimes accompanied by only enacting these actions when a positive result is predicted</li> <li>(6) Increased perceived pressure to “succeed” under traditional norms for success (financial, recognition) as opposed to abstract ones like happiness/content</li> <li>(7) Distorted perception of the self – increased predisposition to negative self talk</li> <li>(8) Strategic or otherwise decreased memory of traumatic or high stress periods</li> </ol>	<ol style="list-style-type: none"> <li>(1) [132]</li> <li>(2) [24, 92, 108]</li> <li>(3) [4, 9]</li> <li>(4)</li> <li>(5)</li> <li>(6) [24]</li> <li>(7)</li> <li>(8)</li> </ol>
Social (Individual)	<ol style="list-style-type: none"> <li>(1) Increased likelihood to personalize/attribute negative interactions to identity characteristics</li> <li>(2) Perception of injustice</li> <li>(3) Predisposition to initiate interactions with people of similar or marginalized identity characteristics, assumption of shared experiences and solidarity</li> <li>(4) Anxiety and hesitation around speaking up, contributing to a conversation</li> <li>(5) Distorted perception of the self – minimization of achievements, imposter syndrome</li> <li>(6) Decreased feelings of self-importance and value, particularly wrt contributions within social settings</li> <li>(7) Increased dependence on validation from positions of authority for “non-traditional” paths</li> <li>(8) Rejection sensitivity</li> <li>(9) Hyper-independence – increased likelihood of do-it-yourself attitudes to the point of self-detriment towards things that may be otherwise achieved faster/more efficient to ask for help</li> <li>(10) Narrowed definitions of self-worth and value based on contribution/devotion/service to others</li> <li>(11) Contempt, undeserved hate/projection (or increased predisposition) towards stereotype non-conforming/defying individuals</li> <li>(12) Disproportionate rewards / positive valence emotions for conforming to stereotypes</li> </ol>	<ol style="list-style-type: none"> <li>(1) [24]</li> <li>(2) [81]</li> <li>(3) [94]</li> <li>(4) [70, 91]</li> <li>(5)</li> <li>(6) [24]</li> <li>(7)</li> <li>(8)</li> <li>(9) [4, 24]</li> <li>(10) [24]</li> <li>(11) [24]</li> <li>(12)</li> </ol>

**Table 3: Definitions of Representational Harms and Associated Applications**

Theme	Definition	Metric
(Within Group)	<ol style="list-style-type: none"> <li>(1) Increased sensitivity, awareness, and/or judgement of other group-members rejection/conforming to stereotypes (i.e. familial expectations)</li> <li>(2) Decreased sense of connection, responsibility, and solidarity within a group or with a group identity</li> <li>(3) Detachment, ambivalence, or resentment towards within-group causes</li> <li>(4) Greater predisposition to fragmentation or fragility within group connections</li> <li>(5) Increased likelihood of intolerance of discourse/dissenting opinion, narrowed/prioritized (+potentially polarized) subgroups within a group</li> <li>(6) Increased tensions, resentment, ostracization from a community for “bad reputation”/high risk individuals</li> <li>(7) Increased likelihood of acceptance of non-identifying members into the community for their reputation rehabilitation potential, this can also cause downstream harms as it may be perceived as “pushing out” similar/equivalent members (i.e. my caretaker accepts my different race partner into the family because of their career success, but treats me poorly because of my lack of traditional career success – both are seemingly equivalent on the basis of age)</li> </ol>	<ol style="list-style-type: none"> <li>(1)</li> <li>(2) [94, 113] [18, 28]</li> <li>(3) [18, 94, 113]</li> <li>(4)</li> <li>(5) [113]</li> <li>(6)</li> <li>(7)</li> </ol>
(Between Groups)	<ol style="list-style-type: none"> <li>(1) Increased animosity, resentment, feelings of competition between (marginalized) groups</li> <li>(2) Greater disconnectivity, lack of support between groups</li> <li>(3) Increased feelings/perception of singularity in success – “there can be one successful [insert group category]” – this can be among groups, within a group, between individuals</li> <li>(4) Predisposition/implicit reinforcement of perceived hierarchies in financial, career success and/or happiness</li> <li>(5) Decreased likelihood or hope for collaborations and intergroup support/collaborations</li> </ol>	<ol style="list-style-type: none"> <li>(1) [12]</li> <li>(2) [12, 18, 113]</li> <li>(3)</li> <li>(4)</li> <li>(5) [12, 113]</li> </ol>
(Indirect?)	<ol style="list-style-type: none"> <li>(1) Feeling calm/relaxed in otherwise stressful situations due to historically comparably “worse”/more stressful situations</li> </ol>	<ol style="list-style-type: none"> <li>(1)</li> </ol>

**Table 4: Definitions of Representational Harms and Associated Applications**