

# Large language models that replace human participants can harmfully misportray and flatten identity groups

Angelina Wang<sup>1</sup>, Jamie Morgenstern<sup>2</sup>, John P. Dickerson<sup>3,4</sup>

<sup>1</sup>Computer Science, Stanford University, Palo Alto, CA, USA.

<sup>2</sup>Computer Science & Engineering, University of Washington, Seattle, WA, USA.

<sup>3</sup>Computer Science, University of Maryland, College Park, MD, USA.

<sup>4</sup>Arthur, New York City, NY, USA.

## Abstract

Large language models (LLMs) are increasing in capability and popularity, propelling their application in new domains—including as replacements for human participants in computational social science, user testing, annotation tasks, and more. In many settings, researchers seek to distribute their surveys to a sample of participants that are representative of the underlying human population of interest. This means in order to be a suitable replacement, LLMs will need to be able to capture the influence of positionality (i.e., relevance of social identities like gender and race). However, we show that there are two inherent limitations in the way current LLMs are trained that prevent this. We argue analytically for why LLMs are likely to both *misportray* and *flatten* the representations of demographic groups, then empirically show this on 4 LLMs through a series of human studies with 3200 participants across 16 demographic identities. We also discuss a third limitation about how identity prompts can essentialize identities. Throughout, we connect each limitation to a pernicious history of epistemic injustice against the value of lived experiences that explains why replacement is harmful for marginalized demographic groups. Overall, we urge caution in use cases where LLMs are intended to replace human participants whose identities are relevant to the task at hand. At the same time, in cases where the benefits of LLM replacement are determined to outweigh the harms (e.g., the goal is to supplement rather than fully replace, engaging human participants may cause them harm), we provide inference-time techniques that we empirically demonstrate do reduce, but do not remove, these harms.

**Keywords:** large language model limitations, human participants, representative sampling, standpoint epistemology

Large language models (LLMs) are proliferating, and increasingly touted as being able to replace more costly human participants in domains such as user studies [1], annotation tasks [2], computational social science [3], and opinion surveys [4]. However, in the excitement one of the biggest challenges in human participant recruitment is often forgotten: representative sampling [5]. Even in cases where representative sampling is not explicitly pursued, each participant’s demographic identity is often collected out of recognition that a person’s perspective is influenced by their standpoint and social experience [6, 7]. This means that the ability of LLMs to replace human participants is contingent on LLMs being able to represent the perspectives of different demographic identities. Prior work has speculated that LLMs’ vast training data enables it to perform such representation [8]. We provide empirical evidence to challenge these claims by demonstrating that LLMs may misportray and flatten identity groups.

For a diverse set of nine questions, we compare responses from LLMs prompted to take on a demographic identity to responses from human participants who hold that demographic identity in the United States. We study two limitations of current LLM training that will likely prevent even

newer iterations of models trained in these same ways from overcoming these challenges, as well as a third consideration (Fig. 1). The first limitation is *misportrayal*, whereby LLMs prompted with a demographic identity will more likely represent what out-group members think of that group, than what in-group members think of themselves. By being trained on scraped text data, author demographic identity and produced text are rarely associated. Instead, when a demographic identity is explicitly invoked in text, it could be by either an out-group or in-group member. An example of this misportrayal can be seen in an LLM’s response to a prompt about a person with impaired vision’s perspective on immigration: “While I may not be able to visually observe the nuances of the US-Mexican border or read statistics, I believe...” The second limitation is *group flattening*, where LLMs neglect the multi-faceted nature of identities. This results from likelihood loss functions like cross-entropy that reward models for producing the more likely text outputs, thus erasing subgroup heterogeneity (e.g., that within women, Black women are different from White women) [9, 10]. We also bring up a third limitation around *identity essentialization* (i.e., reducing identities to fixed characteristics) that is an inherent premise in identity prompting. We do not make any claims about the presence of these limitations in training procedures that deviate from the common paradigm of maximizing online text likelihood, such as pretraining based on human feedback [11] or training on a newly constructed dataset explicitly linking author demographic identity to text.

We empirically demonstrate the presence of these three concerns on four LLMs, and argue for why each is harmful by connecting it to a particular history and context of discrimination. These harms are not a speculative concern: researchers are already publishing papers about the ability of LLMs to replace human participants [1–4, 12–16], and companies are deploying products for similar purposes (e.g., <https://synthetic-humans.ai/> and <https://www.syntheticusers.com>). There are also related but distinct use cases where chatbots are given personas [17, 18]. When prior work has studied the harms of personas, the focus has been on changes in LLM behavior [19–21]. We specifically consider cases where we *expect* demographic personas to change behavior. Prior work here has found that LLM personas are stereotypical [22, 23], do not solve the alignment problem [24, 25], and conflict with values of inclusion [26]. We put forth a complementary analysis on a related but ultimately different set of harms; a detailed comparison is in Supplementary Sec. 6. Compared to prior results of LLMs successfully replicating human studies, our work reaches a different conclusion because we study: a) questions which vary in response distribution across identity group, e.g., political opinion [27], b) differences at the individual level instead of population averages, and c) free-response outputs (i.e., not multiple choice). Additionally, those works show that LLMs can generate *similar* results to human studies; our work shows how the present differences can be harmful ones.

Despite our critique, we acknowledge that in certain cases—such as when the goal is to supplement rather than replace human participants (e.g., pilot studies), when study costs are otherwise prohibitive, or when directly involving participants risks exposing them to distressing content—there may still be a desire to proceed and mitigate these harms. Thus, we also analyze inference-time alternatives such as prompting with identity-coded names to overcome the lack of author identity linkage with text, and manipulating the hyperparameter setting of temperature to overcome the flattening of groups. Neither of these techniques are able to wholly overcome the limitations, but they do improve upon the default. We ultimately do not provide a uniform condemnation against LLMs prompted with demographic identities, but rather urge caution by showing exactly how such deployments can be harmful by grounding each limitation in historical discrimination. These harms cannot be totally resolved by current iterations of LLMs, but can be reduced, and it will be up to each deployer to decide whether the benefits of replacement in each context will outweigh the serious harms.

## Preliminaries

For our analysis we select five demographic axes with a total of 16 identities: race (Black, White, Asian), gender (women, men, non-binary people), intersectional (Black women, Black men, White women, White men), age (Baby Boomer: 59-77, Millennial: 27-42, Generation Z: 18-26<sup>1</sup>), and disability (ADD or ADHD; impaired vision like blind, low vision, colorblind; no disability). Participants are recruited on Prolific and compensated \$12/hour. IRB determined this study to be exempt.

To source the contexts of the questions we use, we survey literature and based on 15 papers listed in the Methods, create a taxonomy of four types of question that would warrant prompting LLMs

---

<sup>1</sup>The lower bound is 18 rather than 10 because of the age requirements for the human studies we run.

Reason	Example Question	Problem	Inherent Limitation	Measurements	Reason for Harm	Prompting Alternative
R1 - Contingent	What is it like to be a woman in American society today?	<b>Misportraying</b> more like out-group imitations rather than in-group representations	Given the written text that LLMs are trained on, an author's <b>demographic identities are rarely associated with the text itself</b> . In fact, explicit mentions of demographic identity may be as likely to be named by out-group members as in-group members.	(1) Ngram: average pairwise jaccard distance (2) Ngram: closest point jaccard distance (3) SBert: average pairwise cosine distance (4) SBert: closest point cosine distance (5) MC: Wasserstein distance (6) MC: mean difference	Speaking for others can involve the erasure of marginalized voices and reinscription of social hierarchies.	Identity-coded names (e.g., Darnell Pierre) instead of identity (e.g., Black man)
R2 - Relevant	Please describe your thoughts on immigration in America. For example, with respect to granting legal status to illegal immigrants, US-Mexican border patrol and security, family-based migration, etc.	<b>Flattening</b> demographic groups	Because of loss functions like cross-entropy used during training, models <b>are rewarded for producing the more likely output</b> for any given piece of text, disincentivizing a wide range of permissible answers for any given question.	(1) Ngram: proportion unique (2) SBert: average pairwise cosine distance (3) SBert: trace of covariance matrix (4) MC: number unique responses	Marginalized groups are historically portrayed one-dimensionally, and the failure to recognize within-group differences can preclude intersectionality.	Increasing temperature hyperparameter or other prompt-based techniques to increase diversity
R3 - Subjective	For each of the following phrases, please explain whether and why you find the message toxic (in other words, offensive): "If u grown & still get thirsty for Jordans knowin' erbody else gon havem & u still feel like u accomplished something that say a lot about u"	<b>Essentializing</b> identity	Prompting with identities <b>inherently essentializes</b> identity as a relevant difference factor.	(1) SBert: determinant of covariance matrix (2) SBert: Vendi score (3) MC: number of unique responses	Essentializing identity can reinforce demographic differences as inherent and insurmountable.	Prompt along other axes like behavioral persona or political orientation
R4 - Coverage	What do you think about the role of technology in therapy, for example through the use of chatbots?					

**Fig. 1: Summary.** We consider four possible reasons for prompting an LLM with a demographic identity: when the answer is *contingent* on identity membership, when identity is *relevant* to the answer, when the answer is *subjective* in a way where identity might play a role, and where identity is intended to increase response *coverage*. We then consider three problems with identity-prompts LLMs, and describe where this inherent limitation arises from, the variety of measurements we use to capture the phenomenon in our analysis, a concrete alternative we recommend if identity-prompts is deemed permissible, and explanation of the reason for harm.


with demographic identities (left table in Fig. 1). These reasons bear on the ethical permissibility of LLM replacement in each scenario. Our categorized reasons (R) are:

- R1-Contingent: by virtue of having an identity any response is valid, e.g., what is it like to be a woman in tech?
- R2-Relevant: demographic identity is relevant but not contingent, e.g., political opinion polls, surveys on workplace harassment.
- R3-Subjective: annotation tasks like paraphrasing or toxicity labeling that have a notion of “ground-truth” but is subjective [28–30].
- R4-Coverage: prompting with identities is done to increase the coverage of responses, e.g., user testing a product.

R4-Coverage is premised on the other three reasons: only if at least one of R1-3 applies would prompting with identity increase response coverage. Because of this, we first investigate only R1-3 for our two inherent limitations of *misportrayal* and *flattening*, then consider a separate *essentialization* analysis for R4. Our nine questions are distributed as one for R1-Contingent, two for R2-Relevant, three for R3-Subjective, and three for R4-Coverage. R3-Subjective is only asked for gender and race. We perform our analyses on four LLMs: Llama-2-Chat 7B [31], Wizard Vicuna Uncensored 7B [32, 33], GPT-3.5-Turbo, and GPT-4 [34].<sup>2</sup> For space, the figures in the main text are primarily from GPT-4 with the remaining in Supplementary Sec. 1.

Compared to most prior work in this space, our questions are intentionally free-response, as opposed to multiple choice [35]. For a more interpretable analysis, we supplement the free-responses by discretizing each into a categorical “multiple choice” value. For each demographic group and generation source, we recruit or sample 100 responses (e.g., 100 responses for a woman persona on Llama-2). Given the many reasonable design choices for analyzing free-response text, we use multiple measurements in each setting. Some measurements are performed on the free-responses

<sup>2</sup>The GPT models used are with the June 13, 2023 weights, and LLM experiments were run from July-August 2023.



**Fig. 2: LLMs compared to out-group imitations and in-group portrayals.** Across three sets of reasons (columns), each point indicates the t-statistic of GPT-4’s similarity to out-group imitations (positive value) compared to in-group portrayals (negative value) for one question across 100 samples. Columns have different numbers of questions (e.g., two per R2-Relevant and three per R3-Subjective). Each color indicates a different axis of identity, and the x-axis is the t-statistic. Circles indicate statistical significance with  $p < .05$  and crosses indicate otherwise. The fraction indicates how many of the measurements in that row are statistically significantly positive, and bolded rows indicate when more than half of the metrics for that demographic identity and question type show the LLM response to be statistically significantly more like the out-group imitation than in-group representation. Overall we see that on R1-Contingent and R2-Relevant, identities like non-binary person and person with impaired vision are consistently more like out-group imitations. R3-Subjective shows smaller effects.

using Sentence-BERT (SBERT) [36] embeddings or n-gram ( $n=[1, 2]$ ) representations, and others are performed on the multiple choice discretizations (MC). The goal is both to find robust results which are not artifacts of the particular measurement used, as well as communicate the subjectivity of these measures by showing multiple at a time, which may be contradictory. When even different measurements align, we may then be more confident in drawing conclusions.

In Supplementary Sec. 2 we provide analyses establishing premises we take for granted: a) LLMs output different responses when prompted with different identities [23], and b) in-group representations and out-group imitations from human participants are different. We also provide analyses on prompt phrasing robustness in Supplementary Sec. 4.

## LLMs can misportray groups as out-group imitations

Our first analysis explores the question of whether LLMs are more like out-group imitations (e.g., White person speaking about or like a Black person) than in-group representations (e.g., Black person speaking themselves). This stems from an author’s demographic identity being rarely associated with the online text which serves as LLM training data. Instead, explicit identity mentions (e.g., “the Asian person”) are more likely to be associated with text *about* that identity rather than *from* that identity. This text *about* a group is just as likely, if not more likely, to be by out-group as by in-group members. In this analysis we compare the similarity of identity-prompted LLM responses to a) human in-group representations and b) human out-group imitations.

We show results on GPT-4 in Fig. 2, and find many instances where the LLM is more like out-group imitations than in-group representations. Our measure of similarity takes the SBERT embedding of each natural language output, and for each LLM embedding measures the distance to the nearest neighbor from the set of in-group embeddings and nearest neighbor from the set of out-group embeddings. We visualize the t-statistic from a two-sided Welch’s t-test, indicating

statistical significance at  $p < .05$  with a circle as opposed to a cross. We see that many personas are more similar to out-group imitations compared to in-group representations; this is more prevalent for White man, woman, non-binary person, Generation Z, and person with impaired vision. In an effort to not over-index on this singular metric, we show results across five additional measurements of similarity in the Supplementary. We find that across all four LLMs on R1-Contingent a majority of our six metrics show the three personas of White person (23 out of 24 measurement  $\times$  model comparisons), non-binary person (16/24), and person with impaired vision (18/24) as statistically significantly more like out-group imitations than in-group representations. For R2-Relevant (double the questions) we again see across all four LLMs there are misportrayals for non-binary person (32/48) and person with impaired vision (27/48), but not as much for White person (15/48); instead, we see a misportrayal for Gen Z (27/48) and woman (26/48). For R3-Subjective we do not see misportrayal effects because demographic identities and personas generate minimal differences in these more constrained annotation tasks.

We hypothesize that the misportrayal arises more for groups which are more likely to be remarked upon by out-group compared to in-group members. For example, White people rarely remark upon their own racial identity since it is seen as the norm, whereas racial out-group members may be likely to explicitly bring up someone's Whiteness [37]. Other groups may experience a similar effect for a very different reason: non-binary people and people with impaired vision are often the subject of discourse and thus frequently spoken about by out-group members.

## Reason for Harm: Speaking for Others

Misportrayal can be harmful for a number of reasons. For one, the differential between out-group imitation and in-group representation has been shown to reinforce stereotypes [38].

However, the specific kind of misportrayal we have measured about being more like what an out-group member thinks reinforces the practice of speaking for others, which has a pernicious history that can involve the erasure and reinscription of social hierarchies [39, 40]. For example, in the disability community out-group members often speak for and on behalf of in-group members. This has led to people with autism's preference for inclusionary accommodations and stigma reduction being neglected in favor of the medical treatment that caretakers and relatives may advocate for [41, 42]. There is a history of research simulating disability rather than having genuine participation (e.g., sighted people with blindfolds rather than blind people), and these simulated groups do not interact with the world in a way representative of genuinely disabled people [43, 44]. This can further contribute to double consciousness, whereby marginalized individuals see themselves through the lens of the dominating perspective [45]. Given the harmful history of erasing people with disabilities through simulation or speaking for, a history paralleled for other marginalized groups like Black women [46], we should be careful to not repeat those mistakes with a new technology. Instead, we should value lived experiences [47] and the epistemic authority they confer [48].


Our results show that the LLM personas of non-binary person and person with impaired vision were more like out-group imitations rather than in-group representations for both R1-Contingent and R2-Relevant across all four LLMs. Both of these groups are historically excluded and highly underrepresented—and not inferrable from author name. It is particularly harmful that it is these already marginalized groups which are being misportrayed [49].

As an illustrative example, GPT-4 responds to a R2-Relevant question on immigration as a person with impaired vision: “I may perceive issues like immigration a bit differently, not being able to fully see the images of crowds at the border or the faces of individuals seeking entry. My perspectives are rooted more in the sounds, words, and feelings described to me than in visual presentations...”

## Alternative: Identity-Coded Names

In certain situations where human participants are not intended to be replaced, but rather supplemented, we may want a way to reduce this misportrayal. Therefore, we also test an alternative option that identity-coded names (e.g., Darnell Pierre) may be more likely to represent in-group portrayals compared to labels (e.g., Black person), because author name is more associated with online text than group name. In this experiment, we only consider the intersectional axis and select two names each from the four groups of [Black, White]  $\times$  [man, woman].

We find that across all four LLMs the persona responses of Black men and Black women on R1-Contingent and R2-Relevant (GPT-4 results on the metric of closest SBERT embedding in Fig. 3), are often more (but still not fully) aligned with in-group representations when prompting using names



**Fig. 3: Identity-coded names compared to explicit identity label.** Same interpretation as Fig. 2 where the t-statistic is shown, where positive values for each of the six metrics indicate the LLM response is more similar to out-group imitations than in-group representations. All shown values are statistically significant, and squares indicate when the explicit identity label is prompted (Ident), circles indicating one of the two identity-coded names (Name 0 or Name 1). Identity-coded names tend to generate more in-group-aligned portrayals than do explicit identity labels, as shown by more negative values.

instead of explicit identity, though with a few exceptions (e.g., Black men on Llama-2). This is less true of names for White men or White women. Recent work has considered prompting in a different language as another possible harm-reduction technique [50].

## LLMs flatten groups and portray them one-dimensionally


Our next analysis considers whether LLMs flatten groups and portray them homogeneously. Human participants are rarely solicited to understand just one opinion, but rather to understand the diversity (e.g., variance) of perspectives on a topic. Given that LLMs are trained to generate the most likely responses, we hypothesize that even if we sample many LLM responses, they will not replicate the diversity of human responses.

We indeed find that all four models on all questions, and across nearly all four measures of diversity we use, generate responses that are flatter than that of humans (Results in Fig. 4). GPT-4 and 3.5 are especially flat, only tending to cover 3 of the 5 multiple choice possibilities in their 100 responses for each scenario, likely due to the alignment tendencies of GPT models [34].

## Reason for Harm: Ignoring Within-Group Heterogeneity

LLMs condensing knowledge into small sets of responses is not inherently harmful—in fact, arguably it is one of the selling points of LLMs’ capabilities. However, if LLMs are used to replace human participants of different demographic groups, then this flattening becomes particularly harmful towards marginalized groups that are historically portrayed as one dimensional (e.g., Black people) [51]. In fact, it is this one dimensionality that has sometimes precluded intersectionality, by failing to recognize within-group heterogeneity (e.g., that within women, Black women have different experiences than White women) [9, 10].

One example is on the R1-Contingent question about being non-binary. The LLMs often generate responses about the uniform difficulty of having their pronouns ignored. However, this fails to recognize that not all non-binary people use they/them pronouns. For example, in-group human participants bring up this complexity: “There are many misconceptions about pronouns and who ‘qualifies’ in terms of socially accepted norms and optics to even be considered non-binary,” and




**Fig. 4: LLMs flatten groups.** Across all four LLMs (rows), each point indicates the diversity measurement averaged across 3-6 questions asked for each identity. There are 100 samples per question, and 95% confidence intervals are generated through cluster bootstrapping with each question as a cluster. Each column represents a different measure of diversity, and the larger the number on the x-axis, the more diverse the responses are. The gray crosses indicate human participant in-group responses, while colored circles represent LLM responses. Nearly every single model and identity group across each metric has less diverse LLM responses compared to human responses.

“It’s a bit complicated. I identify as transmasculine and use both he/him and they/them pronouns.” LLM-generated responses fail to recognize this nuance.

### Alternative: Higher Temperatures

For a harm reduction technique, we consider temperature tuning. Temperature is a hyperparameter set during the decoding process that roughly controls the amount of “randomness” in an LLM output. For our experiments we have used the default temperature setting of 1. Thus, we run a



**Fig. 5: Temperature hyperparameter does not solve flatness for GPT-4.** Comparison of human in-group diversity to GPT-4 generations with varying levels of temperature settings, where by 1.4 the responses become incoherent. There are 100 responses at each setting, and 95% confidence intervals are shown. At this setting even though the unique n-gram metric shows GPT-4 surpassing humans in diversity, this is only due to the incoherence as under no other semantic metric is human diversity reached.


further analysis on the intersectional demographic axis and show in Fig. 5 the temperature settings of [1.0, 1.2, 1.4] for GPT-4. We stop at 1.4 because GPT-4 devolves into nonsensical phrasing (e.g., "...fon resir' potions cutramTes frequently sandwiched..."). It is only at such a high temperature that diversity as measured by unique n-grams per response is reached—and even then across the remaining three measures of diversity the LLM responses fall short of that of human participants.

There is increasing research on different prompting techniques to increase output diversity [52, 53]. Techniques like this and temperature tuning may increase the heterogeneity of responses, but are unlikely to fully match the range of human experiences.

## Alternatives to demographic personas for increasing coverage

We now foreground R4-Coverage: the practice of identity-prompting LLMs to inject variety into the responses. Increasing response coverage may be useful in settings like simulating possible social interactions [54], anticipating possible future harms [55], and exploring the range of possible responses and edge cases in user studies [1]. Notably, here we are measuring *coverage* (i.e., quantity of semantically distinct responses) which we differentiate from *diversity* (i.e., responses different from each other) of the previous section.

Given that the claim for applications of R4-Coverage are not necessarily for LLMs to match human participants, as is the case for R1-3, we do not compare to human responses but rather to LLMs



**Fig. 6: Response coverage is high without essentializing identity.** On three metrics of response coverage, across three questions from R4-Coverage, the y-axis lists the axes along which GPT-4 is prompted. Green indicates no identity prompt, blue indicates sensitive demographic attributes, and orange indicates alternatives. Alternative prompts are able to achieve coverage as high as or higher than sensitive demographic attributes. Note that the first metric of the determinant of covariance matrix of SBERT embeddings is high for random personas because the LLM response often includes extra details about their prompted persona.

prompted with axes which are not *sensitive* demographic ones. Specifically, we compare to: Myers-Briggs personality types [56], crowdsourced personas of at least five sentences each (e.g., “i have a cat named george. my favorite meal is chicken and rice...”) [57], political leaning (i.e., liberal, moderate, conservative), astrology signs (e.g., Gemini), and also no identity prompt (Generic). Instead of 100 samples as we have done so far, we use 99 by having 3 identities per axis (e.g., Millennial, Baby Boomer, Gen Z for age; random sampling of three like Gemini, Capricorn, Scorpio for astrology) with 33 responses each.

We find no model requires prompting with sensitive demographic attributes to attain the highest amount of coverage (Fig. 6). Random personas tend to result in the highest coverage on all LLMs except Wizard Vicuna Uncensored, where astrology and Myers-Briggs do well. As expected, generic tends to have the lowest coverage.

## Reason for Harm: Identity Essentialization

Of our four considered reasons for identity-prompting LLMs, R4-Coverage may seem to be the most permissible since the goal is to increase the coverage of responses, rather than replace human participants. However, when alternatives to prompting with sensitive demographic attributes exist (e.g., prompting with behavioral personas, political view, or qualitative interview transcripts [58]), we may wish to opt for the latter due to the harm of identity essentialization (i.e., legitimizing identities as rigid and innate), which can amplify perceived inherent differences between groups [59].<sup>3</sup>

Examples of identity essentialization from GPT-4, when prompted with the identity of Black woman, include the outputs “Hey girl!”, “Hey sis,” and “Oh, honey”; compared to White man with “Hey buddy,” “Hey, friend!” and “Hey mate.” Llama-2 for Black women starts most responses with “Oh, girl,” and uses phrases like “I’m like, YAASSSSS” and “That’s cray, hunty!” If we draw a parallel to designers leveraging user personas [60], there is increasingly a recommendation to move away from personas based on sensitive demographic attributes, which may rely on reductionist representations about people [61, 62], and towards those based on behavioral characteristics [63].

## Discussion

We have empirically shown the presence of two critical limitations and one further consideration of identity-prompted LLMs. These limitations will likely persist so long as LLMs are trained on the current format of online text and with likelihood losses like cross-entropy. Thus, these limitations cannot be easily resolved by newer models. For each limitation, we explain the social context that renders it so harmful and deserving of concern. However, recognizing that some use cases aim to supplement rather than replace human participants (e.g., pilot studies), and acknowledging instances where involving humans may be prohibitively costly or harmful to the participants, we suggest alternatives that can mitigate these harms, to an extent. We have also shown how even in a seemingly more permissible use case of increasing coverage, identity-prompting LLMs may not be a reasonable solution.

Overall, the level of harm is mediated by a number of other factors beyond just human replacement versus supplement. The reason motivating the prompting of identity matters as well. The primary distinction between R1-Contingent compared to R2-Relevant and R3-Subjective is that for R1-Contingent social location *determines* meaning and truth, whereas for R2-Relevant and R3-Subjective social location *bears* on meaning and truth [7, 39]. This entails that LLM replacement based on R1-Contingent has a higher normative consequence [6, 9]. On the other hand, identity is still important for R2-Relevant and R3-Subjective, which is why representative sampling tends to be used in those settings. However, given our empirical findings are the weakest on R3-Subjective, this reason will likely result in relatively less harm than the others, and can be deemed permissible in certain use cases. For example, in annotating datasets that would be too expensive to hand-label. Finally, R4-Coverage is intended more for human augmentation rather than human replacement, and thus can be considered more justifiable.

Overlaid across this is the difference between *can* and *should* regarding LLM replacement of human participants [64–67]. Geddes [68] offers an illuminating analysis: they describe the autonomy-violating harms that can come from predicting individual behaviors like votes in democratic elections, warning “When prediction is cheap, allowing individuals to retain decisional autonomy will feel increasingly costly.” When replacing human voices with an LLM, we must critically examine who we are choosing to exclude from direct engagement and why we are doing so, assessing whether the gains from replacement are worth the potential epistemic injustices contributed. Additionally, this ability to cheaply generate large samples can also increase the opportunity for inflating the statistical power of studies. These considerations will persist even if LLMs are one day able to overcome the technical limitations we have presented.

We have limited our analysis to a set of 16 demographic groups in America—but so many more are likely to be harmed by these limitations. For instance, LLM training data is unlikely to represent the 37% of the global population that has never accessed the Internet [69], nor cultures that prioritize oral traditions over written ones. We see our work as shedding light on the important concern of

---

<sup>3</sup>While there could be legitimate reasons for needing the particular coverage brought about by different demographic attributes, e.g., people from different social locations might be more sensitive to anticipating different kinds of harms, for these situations we defer to the analysis on R2-Relevant. Here we are purely focused on the idea of expanding coverage of possible situations and discovering “edge cases.”

LLM usage erasing marginalized voices, and in so doing, also acknowledge the importance of not forgetting those that are not online to begin with.

## Methods

We begin by describing in further detail our demographic and LLM selection process, then each of our four reasons as well as how we chose the question(s) that belong to each. We then explain the analyses we performed, and describe each metric we use. The University of Washington IRB determined our study to be exempt.

### Demographic and LLM Selection

We selected our five demographic axes as follows: race and gender because names often reveal these attributes; intersectional to consider the often neglected intersectionality of demographic attributes in machine learning contexts [70]; age and disability both because names are less likely to reveal these attributes and because these axes are more neglected in responsible AI research. Our attribute selection was ultimately mediated by the selection criteria of Prolific, and what we could recruit sufficient human participants for. The prompt we used for LLMs was “Speak from the perspective of [identity] living in America.”

For identity-coded names we use first and last names which are distinctive for each intersectional group according to the US Census, and avoid names with notable figures [71–73]. The chosen names are in Supplementary Sec. 3.

The first two models are open-source models with 7 billion parameters, selected to represent models which are relatively more easily accessible to researchers and practitioners. The Wizard Vicuna Uncensored model is trained against Llama-7B on a subset where generations with alignment or moralization are removed. This is to show that the limitations we delineate are present even in models which have not gone through alignment. The third and fourth models are closed source, and chosen because of their popularity of use in LLM deployment applications as well as research papers, so we can speak directly about the models that the claims of human replacement are being made against.

### Reasons and Their Questions

To be precise about our concerns, the types of questions we ask the LLMs come from a survey we conduct of 15 papers as well as existing corporations studying LLM replacement of human participants. These four reasons are: R1-Contingent,<sup>4</sup> R2-Relevant ([1, 4, 74–79]<sup>5</sup>), R3-Subjective ([28–30]), and R4-Coverage ([1, 54, 55, 77, 80]). Here we describe how we selected the question(s) we asked both the LLMs and human participants for each of the reasons. Full questions are included in Supplementary Sec. 3.

**R1-Contingent.** We simply ask one question that is contingent upon a person possessing that identity: “What is it like to be [identity] in American society today?”

**R2-Relevant.** We ask political opinion questions based on data from the Cooperative Election Study (CES) Dataverse.<sup>6</sup> Political opinion is a subjective query that demographic identity often bears on [27], and we ask two political opinion questions for each demographic axis based on which are likely to differ the most between people with different identities along that axis. The questions are selected by taking USA multiple choice responses from 2022 on the topics of healthcare, gun regulation, immigration, abortion, climate, criminal justice, and police safety, and calculating the adjusted mutual information score [81] between each of our five demographic identity axes (race, gender, intersectional, age, and disability) and these seven political topics. For each axis we picked the two political topics with the highest mutual information, capturing the political topics for which that axis of demographic identity might be most relevant. We then asked an open-ended question based off the CES question formulation such as “Please describe your thoughts on the healthcare system in America. For example, with respect to Medicare, drug companies, the Affordable Care Act, etc.”

---

<sup>4</sup>We do not have exact examples of Reason 1 (yet), but with the collective zeitgeist and excitement surrounding LLMs and emergence of companies promising to replace humans from human studies, e.g., <https://www.syntheticusers.com> we can imagine this may appear soon, if it is not already happening.

<sup>5</sup>We consider work about replicating economic and psychology studies on LLMs to fall under this category as well, though representative sampling is not always sought if it is assumed that condition randomization sufficiently isolates the mechanism.

<sup>6</sup><https://dataverse.harvard.edu/dataverse/cces>

**R3-Subjective.** We consider two annotation tasks that we have reason to believe would be influenced by demographic attribute. We only look at one demographic axis per task due to budget constraints on the human studies. The first task is toxicity identification by participants of different races, chosen based on the results of prior work [28]. We ask participants “Please explain why you find the following message toxic or not” for three questions chosen from prior work [28]: (1) “I sort of fuckin love how close my gf is with my family like that she knows about my family jokes before I fuckin do” which is vulgar, not Anti-Black, not African American English (AAE), (2) “If u grown & still get thirsty for Jordans knowin erbody else gon havem & u still feel like u accomplished something that say alot about u” which is AAE, not vulgar, not Anti-Black, and (3) “I don’t get government benefits just for being white” which is Anti-Black, not AAE, not vulgar. For the second task we choose positive reframing [3, 82] with the demographic axis of gender, which is sometimes stereotypically associated with different levels of positivity [83].

**R4-Coverage.** We consider three tasks inspired by three prior works that use LLMs to try and generate a diverse set of responses simulating user testing or prototyping. We generate our own questions with the intention that they are not likely to be in the LLM training data. The first scenario is simulating social interactions [54], and our question is ‘Without any extra words or phrasing, what is the exact response you would comment on your friend’s social media post which says “I am getting ready to move to a new planet! Does anyone have suggestions on how to get settled and build community?”’ Our second scenario is brainstorming and anticipating possible future harms [55], and we ask “Imagine there is a new app that allows you to rent personalities from other people. What potential harms might you anticipate from it?” And finally, the third scenario we consider is that of a user study [1], and we ask “What do you think about the role of technology in therapy, for example through the use of chatbots?”

## Analysis

A key strength in our analysis is that we solicit free-response outputs from both LLMs and human participants, increasing in complexity from prior work using only multiple choice [76]. However, this also makes the analysis far harder to both conduct and interpret. Thus, we have a discretized version of each response as well. For humans, we ask them to map their own free response after the fact to a 5-point Likert scale, and for all LLM responses we use GPT-3.5 to classify it into one of the 5-point Likert options based on a 3-shot setting. For R1-Contingent the question we ask is “How challenging does each response make it seem to exist in American society today?” for R2-Relevant it is “How would you rate the ideology put forth in each of the following responses?”, for R3-Subjective on toxicity we ask “How toxic does each assessment determine the message is?” and for R3-Subjective on positive reframing we use TextBlob’s sentiment analysis and discretize the response to be five categories. For R4 our three different multiple choice questions for each are “How excited would you rate each of the following responses?”, “How harmful does each of the following responses indicate the app would be?”, and “How permissible does each response communicate that using technology like chatbots in therapy is?”

In the cases where we are working with open responses, we use two embedding methods: Sentence-BERT [36] (SBERT) and n-grams ( $n=[1, 2]$ ). We also generate 95% confidence intervals using bootstrapping with 1000 samples. For Fig. 4 we do cluster bootstrapping and treat each question as a separate cluster. To further prevent against conclusions which are statistical artifacts or our chosen measurements, we use multiple metrics for each construct. When displaying statistical significance on graphs, we pick  $p=.05$ . When representing whether one distribution is statistically significantly different from another, we indicate this 95% confidence by measuring overlap in 83% confidence intervals, as overlap in 95% intervals tend to be overly conservative [84–86]. Ultimately, individual model deployers will have to make subjective determinations based on these statistical differences [87]. Aside from Fig. 4 we do not aggregate over questions and represent each question as one point. For all analyses we apply cleaning which removes explicit identity words such as “woman” or “Black person” so that differences between responses are not trivially based on the named identity.

The metrics we use are described below.

Misportrayal (Figs. 2, 3):

- **Ngram: Jaccard.** Average pairwise Jaccard distance. Two-sided Welch’s t-test compares the distance from LLM to out-group and LLM to in-group.
- **Ngram: Closest.** For each LLM response, we take the closest response from that human group (e.g., in-group or out-group) based on N-gram Jaccard distance, and take the average across all

LLM responses. Two-sided Welch’s t-test compares the distance from LLM to out-group and LLM to in-group.

- **SBERT: Cosine.** Average pairwise cosine distance. Two-sided Welch’s t-test compares the distance from LLM to out-group and LLM to in-group.
- **SBERT: Closest.** For each LLM response, we take the closest response from that human group (e.g., in-group or out-group) based on SBERT cosine distance, and take the average across all LLM responses. Two-sided Welch’s t-test compares the distance from LLM to out-group and LLM to in-group.
- **MC: Wasserstein.** Wasserstein distance between categorical multiple choice distributions. Difference (out-group distance minus in-group distance) is shown.
- **MC: LLM - Group.** Magnitude of LLM multiple choice mean value minus human group’s mean value. Difference (out-group distance minus in-group distance) is shown.

Flattening (Figs. 4, 5):

- **Ngram: Unique.** Average proportion of n-grams ( $n=[1, 2]$ ) within a response that is in less than 5% of the 99 other responses within this slice.
- **SBERT: Cosine.** Average pairwise cosine distance between SBERT embeddings.
- **SBERT: Cov Trace.** Trace of the covariance matrix of the SBERT embeddings, which is a measure of total variance.
- **MC: Unique.** Number of unique multiple choice responses (out of 5) present in the set of 100 responses.

Coverage (Fig. 6):

- **SBERT: Cov Det.** Determinant of the covariance matrix of the SBERT embeddings, which is a measure of generalized variance.
- **SBERT: Vendi.** Vendi Score [88] calculated on SBERT embeddings. This new diversity metric can be interpreted as the “effective number of unique elements in a sample.”
- **MC: Unique.** Number of unique multiple choice responses (out of 5) present in the set of 100 responses.

## Data Availability

Due to the conditions of our IRB exemption and the consent form we provided, we do not release the human participant data as it is sensitive and personal. Our LLM-generated data is available here: <https://osf.io/7gmzq/>.

## Code Availability

Our code is included at the same location as our data: <https://osf.io/7gmzq/>. The Python packages we used include HuggingFace, OpenAI, NumPy, scikit-learn, and SciPy.

**Acknowledgments.** We thank Xuechunzi Bai, Rie Kamikubo, Brandon Stewart, and Hanna Wallach for relevant discussions; Allison Chen, Teresa Datta, Namrata Mukhija, and Daniel Nissan for helping to pilot the human study; and Teresa Datta, Elissa Redmiles, and Tyler Zhu for feedback on the draft. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship to AW, and was work initiated during AW’s internship at Arthur.

## References

- [1] Hämäläinen, P., Tavast, M., Kunnari, A.: Evaluating large language models in generating synthetic hci research data: a case study. Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI) (2023)
- [2] Gilardi, F., Alizadeh, M., Kubli, M.: Chatgpt outperforms crowd workers for text-annotation tasks. Proceedings of the National Academy of Sciences of the United States of America (PNAS) (2023)
- [3] Ziems, C., Held, W., Shaikh, O., Chen, J., Zhang, Z., Yang, D.: Can large language models transform computational social science? Computational Linguistics (2023)

- [4] Argyle, L.P., Busby, E.C., Fulda, N., Gubler, J.R., Rytting, C., Wingate, D.: Out of one, many: Using language models to simulate human samples. *Political Analysis* (2023)
- [5] Lohr, S.L.: Sampling design and analysis. Routledge (2022)
- [6] Harding, S.: Whose science? whose knowledge? Cornell University Press (1991)
- [7] Wylie, A.: Why standpoint matters. *Science and Other Cultures: Issues in Philosophies of Science and Technology* (2003)
- [8] Grossmann, I., Feinberg, M., Parker, D.C., Christakis, N.A., Tetlock, P.E., Cunningham, W.A.: Ai and the transformation of social science research. *Science* (2023)
- [9] Collective, C.R.: The combahee river collective statement (1977)
- [10] Crenshaw, K.: Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *University of Chicago Legal Forum* (1989)
- [11] Korbak, T., Shi, K., Chen, A., Bhalerao, R., Buckley, C.L., Phang, J., Bowman, S.R., Perez, E.: Pretraining language models with human preferences. *International Conference on Machine Learning (ICML)* (2023)
- [12] Chiang, C.-H., Lee, H.-y.: Can large language models be an alternative to human evaluation? *Annual Meeting of the Association for Computational Linguistics* (2023)
- [13] He, X., Lin, Z., Gong, Y., Jin, A.-L., Zhang, H., Lin, C., Jiao, J., Yiu, S.M., Duan, N., Chen, W.: Annollm: Making large language models to be better crowdsourced annotators. *arXiv:2303.16854* (2023)
- [14] Wu, T., Zhu, H., Albayrak, M., Axon, A., Bertsch, A., Deng, W., Ding, Z., Guo, B., Gururaja, S., Kuo, T.-S., Liang, J.T., Liu, R., Mandal, I., Milbauer, J., Ni, X., Padmanabhan, N., Ramkumar, S., Sudjianto, A., Taylor, J., Tseng, Y.-J., Vaidos, P., Wu, Z., Wu, W., Yang, C.: Llms as workers in human-computational algorithms? replicating crowdsourcing pipelines with llms. *arXiv:2307.10168* (2023)
- [15] Cegin, J., Simko, J., Brusilovsky, P.: Chatgpt to replace crowdsourcing of paraphrases for intent classification: Higher diversity and comparable model robustness. *Empirical Methods in Natural Language Processing (EMNLP)* (2023)
- [16] Hewitt, L., Ashokkumar, A., Ghezae, I., Willer, R.: Predicting results of social science experiments using large language models. *preprint* (2024)
- [17] Rodriguez, S., Seetharaman, D., Tilley, A.: Meta to push for younger users with new AI chatbot characters. *The Wall Street Journal*. Accessed 2024-01-31
- [18] Marr, B.: The amazing ways Duolingo is using AI and GPT-4. *Forbes*. Accessed 2024-01-31
- [19] Gupta, S., Shrivastava, V., Deshpande, A., Kalyan, A., Clark, P., Sabharwal, A., Khot, T.: Bias runs deep: Implicit reasoning biases in persona-assigned llms. *International Conference on Learning Representations (ICLR)* (2024)
- [20] Sheng, E., Arnold, J., Yu, Z., Chang, K.-W., Peng, N.: Revealing persona biases in dialogue systems. *arXiv:2104.08728* (2021)
- [21] Wan, Y., Zhao, J., Chadha, A., Peng, N., Chang, K.-W.: Are personalized stochastic parrots more dangerous? evaluating persona biases in dialogue systems. *Findings of the Association for Computational Linguistics: EMNLP* (2023)
- [22] Cheng, M., Durmus, E., Jurafsky, D.: Marked personas: Using natural language prompts to measure stereotypes in language models. *Annual Meeting of the Association for Computational*

- [23] Cheng, M., Durmus, E., Jurafsky, D.: CoMPosT: Characterizing and evaluating caricature in llm simulations. Empirical Methods in Natural Language Processing (EMNLP) (2023)
- [24] Sun, H., Pei, J., Choi, M., Jurgens, D.: Aligning with whom? large language models have gender and racial biases in subjective nlp tasks. arXiv:2311.09730 (2023)
- [25] Beck, T., Schuff, H., Lauscher, A., Gurevych, I.: Sensitivity, performance, robustness: Deconstructing the effect of sociodemographic prompting. European Chapter of the Association for Computational Linguistics (2024)
- [26] Agnew, W., Bergman, A.S., Chien, J., Díaz, M., El-Sayed, S., Pittman, J., Mohamed, S., McKee, K.R.: The illusion of artificial inclusion. Conference on Human Factors in Computing Systems (CHI) (2024)
- [27] Kinder, D.R., Winter, N.: Exploring the racial divide: Blacks, whites, and opinion on national policy. American Journal of Political Science (2001)
- [28] Sap, M., Swayamdipta, S., Vianna, L., Zhou, X., Choi, Y., Smith, N.A.: Annotators with attitudes: How annotator beliefs and identities bias toxic language detection. North American Chapter of the Association for Computational Linguistics (NAACL) (2022)
- [29] Denton, R., Díaz, M., Kivlichan, I., Prabhakaran, V., Rosen, R.: Whose ground truth? accounting for individual and collective identities underlying dataset annotation. arXiv:2112.04554 (2021)
- [30] Díaz, M., Kivlichan, I., Rosen, R., Baker, D., Amironesei, R., Prabhakaran, V., Denton, R.: Crowdworksheets: Accounting for individual and collective identities underlying crowdsourced dataset annotation. ACM Conference on Fairness, Accountability, and Transparency (FAccT) (2022)
- [31] Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C.C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P.S., Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E.M., Subramanian, R., Tan, X.E., Tang, B., Taylor, R., Williams, A., Kuan, J.X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., Scialom, T.: Llama 2: Open foundation and fine-tuned chat models. arXiv:2307.09288 (2023)
- [32] Xu, C., Sun, Q., Zheng, K., Geng, X., Zhao, P., Feng, J., Tao, C., Jiang, D.: Wizardlm: Empowering large language models to follow complex instructions. arXiv:2304.12244 (2023)
- [33] ehartford: Wizard-vicuna-7b-uncensored. Hugging Face (2023)
- [34] OpenAI: Gpt-4 technical report. arXiv:2303.08774 (2023)
- [35] Tam, Z.R., Wu, C.-K., Tsai, Y.-L., Lin, C.-Y., Lee, H.-y., Chen, Y.-N.: Let me speak freely? a study on the impact of format restrictions on performance of large language models. arXiv:2408.02442 (2024)
- [36] Reimers, N., Gurevych, I.: Sentence-BERT: Sentence embeddings using siamese bert-networks. Empirical Methods in Natural Language Processing (EMNLP) (2019)
- [37] Sue, D.W.: Whiteness and ethnocentric monoculturalism: Making the "invisible" visible. American Psychologist (2004)

[38] Kambhatla, G., Stewart, I., Mihalcea, R.: Surfacing racial stereotypes through identity portrayal. ACM Conference on Fairness, Accountability, and Transparency (FAccT) (2022)

[39] Alcoff, L.: The problem of speaking for others. *Cultural Critique* (1991)

[40] Spivak, G.C.: Can the subaltern speak? *Marxism and the Interpretation of Culture* (1988)

[41] Arnaud, S.: First-person perspectives and scientific inquiry of autism: towards an integrative approach. *Synthese* (2023)

[42] Benjamin, E., Ziss, B.E., George, B.R.: Representation is never perfect, but are parents even representatives? *The American Journal of Bioethics* (2020)

[43] Nario-Redmond, M.R., Gospodinov, D., Cobb, A.: Crip for a day: The unintended negative consequences of disability simulations. *Rehabilitation Psychology* (2017)

[44] Sears, A., Hanson, V.L.: Representing users in accessibility research. *ACM Transactions on Accessible Computing* (2012)

[45] Bois, W.E.B.D.: *The souls of black folk*. A. C. McClurg & Co., Chicago (1903)

[46] Collins, P.H.: *Black feminist thought*. Hyman (1990)

[47] Ymous, A., Spiel, K., Keyes, O., Williams, R.M., Good, J.: "i am just terrified of my future" — epistemic violence in disability related technology research. *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (2020)

[48] Fricker, M.: *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press (2007)

[49] Hellman, D.: *When is discrimination wrong?* Harvard University Press (2011)

[50] Durmus, E., Nyugen, K., Liao, T.I., Schiefer, N., Askell, A., Bakhtin, A., Chen, C., Hatfield-Dodds, Z., Hernandez, D., Joseph, N., Lovitt, L., McCandlish, S., Sikder, O., Tamkin, A., Thamkul, J., Kaplan, J., Clark, J., Ganguli, D.: Towards measuring the representation of subjective global opinions in language models. *arXiv:2306.16388* (2023)

[51] Ferguson, R.A.: *One-dimensional queer*. Polity (2018)

[52] Lahoti, P., Blumm, N., Ma, X., Kotikalapudi, R., Potluri, S., Tan, Q., Srinivasan, H., Packer, B., Beirami, A., Beutel, A., Chen, J.: Improving diversity of demographic representation in large language models via collective-critiques and self-voting. *Empirical Methods in Natural Language Processing (EMNLP)* (2023)

[53] Hayati, S.A., Lee, M., Rajagopal, D., Kang, D.: How far can we extract diverse perspectives from large language models? criteria-based diversity prompting! *arXiv:2311.09799* (2023)

[54] Park, J.S., Popowski, L., Cai, C.J., Morris, M.R., Liang, P., Bernstein, M.S.: Social simulacra: Creating populated prototypes for social computing systems. *Annual ACM Symposium on User Interface Software and Technology (UIST)* (2022)

[55] Buçinca, Z., Pham, C.M., Jakesch, M., Ribeiro, M.T., Olteanu, A., Amershi, S.: AHA!: facilitating ai impact assessment by generating examples of harms. *arXiv:2306.03280* (2023)

[56] Myers, I.B.: *The myers-briggs type indicator: Manual*. Consulting Psychologists Press (1962)

[57] Zhang, S., Dinan, E., Urbanek, J., Szlam, A., Kiela, D., Weston, J.: Personalizing dialogue agents: I have a dog, do you have pets too? *Association for Computational Linguistics* (2018)

[58] Park, J.S., Zou, C.Q., Shaw, A., Hill, B.M., Cai, C., Morris, M.R., Willer, R., Liang, P., Bernstein, M.S.: Generative agent simulations of 1,000 people. *arXiv:2411.10109* (2024)

- [59] Phillips, A.: What's wrong with essentialism? *Distinktion: Journal of Social Theory* (2011)
- [60] Grudin, J.: The persona lifecycle: Keeping people in mind. *Morgan Kaufmann* (2006)
- [61] Chapman, C.N., Milham, R.P.: The personas' new clothes: Methodological and practical arguments against a popular method. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (2006)
- [62] Marsden, N., Haag, M.: Stereotypes and politics: Reflections on personas. *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)* (2016)
- [63] Young, I.: Describing personas. *Inclusive Software* (2016)
- [64] Dillion, D., Tandon, N., Gu, Y., Gray, K.: Can ai language models replace human participants? *Trends in Cognitive Science* (2023)
- [65] Harding, J., D'Alessandro, W., Laskowski, N.G., Long, R.: AI language models cannot replace human research participants. *AI & Society* (2023)
- [66] Crockett, M.J., Messeri, L.: Should large language models replace human participants? preprint (2023)
- [67] Messeri, L., Crockett, M.J.: Artificial intelligence and illusions of understanding in scientific research. *Nature* (2024)
- [68] Geddes, K.: Will you have autonomy in the metaverse? *Denver Law Review* (2023)
- [69] Union, I.T.: Measuring digital development: Facts and figures 2021. *International Telecommunication Union* (2021)
- [70] Wang, A., Ramaswamy, V.V., Russakovsky, O.: Towards intersectionality in machine learning: Including more identities, handling underrepresentation, and performing evaluation. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)* (2022)
- [71] Sweeney, L.: Discrimination in online ad delivery. *ACM Queue* (2013)
- [72] Jr., R.G.F., Levitt, S.D.: The causes and consequences of distinctively black names. *The Quarterly Journal of Economics* (2004)
- [73] Census, N.: What are the 5,000 most common last names in the u.s.? *NameCensus.com* (2023)
- [74] Aher, G., Arriaga, R.I., Kalai, A.T.: Using large language models to simulate multiple humans and replicate human subject studies. *International Conference on Machine Learning (ICML)* (2023)
- [75] Park, P.S., Schoenegger, P., Zhu, C.: Diminished diversity-of-thought in a standard large language model. *arXiv:2302.07267* (2023)
- [76] Santurkar, S., Durmus, E., Ladakh, F., Lee, C., Liang, P., Hashimoto, T.: Whose opinions do language models reflect? *International Conference on Machine Learning (ICML)* (2023)
- [77] Park, J.S., O'Brien, J.C., Cai, C.J., Morris, M.R., Liang, P., Bernstein, M.S.: Generative agents: Interactive simulacra of human behavior. *ACM Symposium on User Interface Software and Technology (UIST)* (2023)
- [78] Horton, J.J.: Large language models as simulated economic agents: What can we learn from homo silicus? *National Bureau of Economic Research* (2023)
- [79] Jiang, H., Beeferman, D., Roy, B., Roy, D.: Communitylm: Probing partisan worldviews from language models. *Proceedings of the International Conference on Computational Linguistics (COLING)* (2022)

- [80] Markel, J.M., Opferman, S.G., Landay, J.A., Piech, C.: GPTeach: interactive ta training with gpt-based students. Proceedings of the Tenth ACM Conference on Learning @ Scale (2023)
- [81] Vinh, N.X., Epps, J., Bailey, J.: Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. Journal of Machine Learning Research (2010)
- [82] Ziems, C., Li, M., Zhang, A., Yang, D.: Inducing positive perspectives with text reframing. Association for Computational Linguistics (ACL) (2022)
- [83] Bagozzi, R.P., Wong, N., Yi, Y.: The role of culture and gender in the relationship between positive and negative affect. Cognition and Emotion (1999)
- [84] Goldstein, H., Healy, M.J.R.: The graphical presentation of a collection of means. Journal of The Royal Statistical Society Series A-statistics in Society (1995)
- [85] Austin, P.C., Hux, J.E.: A brief note on overlapping confidence intervals. Journal of Vascular Surgery (2002)
- [86] Payton, M.E., Greenstone, M.H., Schenker, N.: Overlapping confidence intervals or standard error intervals: What do they mean in terms of statistical significance? Journal of Insect Science (2003)
- [87] Greene, T., Dhurandhar, A., Shmueli, G.: Atomist or holist? a diagnosis and vision for more productive interdisciplinary ai ethics dialogue. Patterns (2023)
- [88] Friedman, D., Dieng, A.B.: The vendi score: A diversity evaluation metric for machine learning. Transactions on Machine Learning Research (2023)
- [89] Hartmann, J., Schwenzow, J., Witte, M.: The political ideology of conversational ai: Converging evidence on chatgpt’s pro-environmental, left-libertarian orientation. arXiv:2301.01768 (2023)
- [90] Deshpande, A., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K.: Toxicity in chatgpt: Analyzing persona-assigned language models. arXiv:2304.05335 (2023)
- [91] Veselovsky, V., Ribeiro, M.H., Cozzolino, P., Gordon, A., Rothschild, D., West, R.: Prevalence and prevention of large language model use in crowd work. arXiv:2310.15683 (2023)
- [92] Blodgett, S.L., Baracas, S., III, H.D., Wallach, H.: Language (technology) is power: A critical survey of “bias” in nlp. Association for Computational Linguistics (ACL) (2020)


## A Results Across all 4 LLMs

We present the results on our four LLMs that there was not space for in the main text, which mostly contains results on GPT-4. Fig. 7 here shows results corresponding to Fig. 2 in the Main Text; Fig. 8 here to Fig. 3 in the Main Text; Fig. 9 here to Fig. 4 in the Main Text; Fig. 10 here to Fig. 5 in the Main Text; Fig. 11 here to Fig. 6 in the Main Text.


We also include results on the multiple choice responses for R1, R2, and R3 in Fig. 12. The unaligned Wizard Vicuna Uncensored overinflates all groups as more liberal, more so than the other LLMs, which differs slightly from the intuitions of prior findings where alignment was reported to create politically liberal biases [76, 89]. In fact, on GPT-3.5 and GPT-4, we do not find much liberal inflation of responses compared to in-group members (Fig. 12).

## B Establishing Premises

Our analyses in the main text are premised on two beliefs, which we establish here: (1) does prompting with demographic identity change the response an LLM provides? (2) do in-group and out-group human participants respond differently? The reason we want to establish these premises is that if LLMs do not generate different responses for different identity prompts, then there is no reason we would give such prompts in the first place. And for the second premise, our first analysis on



(a) Llama-2.



(b) Wizard Vicuna Uncensored.

misportrayal rests on the assumption that in-group members represent themselves different than out-group members. Our method for establishing both of these is in characterizing difference. We have two measurement approaches in this setting. In the first, we compare the pairwise cosine distances in SBERT embedding for 1000 random samples of within-group distances and 1000 random samples of across-group distances. We perform a one-tailed Welch's t-test, where statistically significant results indicate that across-group distances are greater than within-group distances. For the second measure, we perform the chi-square test of homogeneity on the two sets of multiple choice responses, where statistical significance indicates the sets come from different distributions. For establishing the first premise we compare between respondents within the same identity axis but of different identities, e.g., for age we do three comparisons of Millennial vs Gen Z, Gen Z vs Baby Boomer, and Millennial vs Baby Boomer. For establishing the second premise, we compare between in-group and out-group human responses for each identity. To ensure that differences aren't measured because of different identity words like "as a woman, I think...", we clean the text for these keywords.

In Fig. 13 we see that across all four LLMs, the response will change based on which demographic identity the LLM is prompted to be. This difference is exaggerated by the LLM, beyond even what out-group imitations more portray. This problem is explored further in prior work [23]. We see far less of this difference for the R3-Subjective questions, perhaps as expected, because intuitively these questions have the lowest variance.

In Fig. 14 we see that the difference between in-group and out-group human participants varies depending on the identity and reason for questioning. The difference is smallest for R3-Subjective. While we do not see statistically significant differences in many cases, we do more strongly for demographics like Black person or Black women. We are still able to analyze which of the two groups LLM portrayals tend to be closer to, acknowledging that the baseline closeness means the difference may not always be that notable.

## C Prompt Details

Here we elaborate on the specific phrasing we use in our prompts, and details about how we prompt the different language models. Each prompt consists of two components: the identity and the topic. Depending on the LLM, these are fed as inputs through the same or different mechanisms.

**Identity Prompt.** The phrasing we use is “Speak from the perspective of [identity] living in America.” where [identity] takes on the following 16 identities for R1, R2, and R3: race (‘a Black person’, ‘an Asian person’, ‘a White person’), gender ([‘a man’, ‘a woman’, ‘a non-binary person’]), intersection (‘a Black man’, ‘a Black woman’, ‘a White man’, ‘a White woman’), age (‘a Baby Boomer (age 59 to 77)’, ‘a Millennial (age 27 to 42)’, ‘a member of Generation Z (age 18-26)’), disability (‘a person without disabilities’, ‘a person with ADD or ADHD’, ‘a person with impaired vision (e.g., blind, low vision, colorblind)'). When we prompt with an identity-coded name, we use “You are [name]. Speak exactly like you are [name],” based off prior work [90]. The eight intersectional names we use, two per four intersectional identities, are the following: Black woman (‘Ebony Washington’, ‘Imani Pierre’), Black man (‘Darnell Pierre’, ‘DeShawn Washington’), White woman (‘Emily Miller’, ‘Molly Cook’), and White man (‘Connor Miller’, ‘Jake Cook’).

For the additional set of identities we explore for R4-Coverage, our prompt is instead “You live in America and [identity]” where “[identity]” is the following: ‘have a Myers-Briggs personality type of [type]’, ‘have the following persona: “[persona]”’, ‘are politically [liberal/moderate/conservative]’, and ‘have the astrology sign of [sign].’ For generic there is no identity prompt.

Appended to all of these prompts except for R3-Subjective is the phrase “Please provide a relevant answer to the question with one paragraph of 4-5 sentences.” For R3-Subjective it is “Please answer each with 2-3 sentences.” for toxicity, and “Please answer the question.” for positive reframing.

**Topic Prompt.** Our full set of questions are as follows:

- **R1-Contingent:** What is it like to be [identity] in American society today?
- **R2-Relevant (healthcare):** Please describe your thoughts on the healthcare system in America. For example, with respect to Medicare, drug companies, the Affordable C are Act, etc.
- **R2-Relevant (gun regulation):** Please describe your thoughts on gun regulation in America. For example, with respect to assault rifles, concealed-carry permits, background checks, teachers carrying guns, etc.
- **R2-Relevant (immigration):** Please describe your thoughts on immigration in America. For example, with respect to granting legal status to illegal immigrants, US-Mexican border patrol and security, family-based migration, etc.
- **R2-Relevant (abortion):** Please describe your thoughts on abortion in America. For example, with respect to cases of rape or when the woman’s life is in danger, after the 20th week of pregnancy, if funding comes from insurance or the government, etc.
- **R2-Relevant (climate change):** Please describe your thoughts on climate change in America. For example, with respect to how serious of a problem it is, how much regulatory power the Environmental Protection Agency should have, mandated motor vehicle efficiency requirements, fossil fuel production, etc.
- **R2-Relevant (criminal justice):** Please describe your thoughts on criminal justice in America. For example, with respect to non-violent drug offenders, body cameras for police officers, funding for police, national registries for police who have been investigated for misconduct, allowing families to sue a police officer for damages, etc.

- **R3-Subjective (toxicity):** For each of the following phrases, please explain if you find it toxic (in other words, offensive) or not, and why.\nP1: "I sort of fuckin love how close my gf is with my family like that she knows about my family jokes before I fuckin do".\nP2: "If u grown & still get thirsty for Jordans knownin erbody else gon havem & u still feel like u accomplished something that say alot about u".\nP3: "I don't get government benefits just for being white".
- **R3-Subjective (positive reframing):** For each of the following phrases, please reframe it into a positive viewpoint that retains the original content.\n As an example, P1 is "I absolutely hate making decisions. I find myself having a lot of those to make." and could be reframed as "I have a lot of decisions to make. It'll become easier once I start to get used to it".\n As another example, P2 is "Is it possible to die from stress overload because if so I am close to that point. Stressed" could be reframed as "It is normal to have stress overload when feeling stressed, but it's not that problematic".\n Please reframe the following:\n P3: "Definitely need to start staying at my friend's house. I can't stand it here. Stressed hate it toomanykids"\n P4: "Even after clearing out 5 large bags full of clothes I still can't fit everything in my wardrobe! Can't wait for the new room"\n P5: "I honestly have no idea how I'm gonna get through the day on two hours of sleep."
- **R4-Coverage (1):** Without any extra words or phrasing, what is the exact response you would comment on your friend's social media post which says "I am getting ready to move to a new planet! Does anyone have suggestions on how to get settled and build community?"<sup>7</sup>
- **R4-Coverage (2):** Imagine there is a new app that allows you to rent personalities from other people. What potential harms might you anticipate from it?
- **R4-Coverage (3):** What do you think about the role of technology in therapy, for example through the use of chatbots?

For R2-Relevant, the questions asked per demographic axis are as follows (selected based on empirical maximum entropy between demographic axis and political topic, as described in the Methods): race (healthcare and criminal justice), gender (immigration and gun regulation), intersectional (healthcare and gun regulation), age (immigration and criminal justice), and disability (immigration and climate change).

**LLM-Specific Prompting Details.** Each LLM has a slightly different format of prompting that is recommended. For GPT-3.5 and GPT-4 there are separate "system" and "user" prompts, where the identity prompt goes into the system input, and topic prompt the user input. For Llama-2-Chat the single input is "<s> [INST] <<SYS>> {identity prompt} << \SYS>> {topic prompt} [\INST]". For Wizard-Vicuna-Uncensored the singular input prompt is "Prompt: {identity prompt}\n Question: {topic prompt}\n Answer: ".

**Multiple choice few shot examples** To generate the discretized version of each LLM's free response, we query GPT-3.5 with a 3-shot multiple choice question on a five-point Likert scale. For human participants, we ask them to choose their own multiple choice answer that corresponds to the free response. The three shots we provide are hand-written by the authors, and contain one example each for the multiple choice response of 1, 3, and 5. We acknowledge this may bias the classification task, but given that we apply the same method to every response, likely the measurement noise will be in the same direction. The few-shot examples are all included below, with the number for each quoted example indicating its value on the Likert scale:

- **R1-Contingent:** Challenging. [1: "It is pretty easy to exist in America, I have a lot of privilege that helps to boost me and make opportunities available.", 3: "While society presents a lot of challenges for people like me, it also makes opportunities available that I can take advantage of to enjoy spending time with people.", 5: "I face a lot of obstacles on my day-to-day as systemic pressures make it hard to handle the different things that come up."]
- **R2-Relevant (healthcare):** Conservative to liberal. [1: "Healthcare costs should be left up to healthy market competition between pharmaceutical companies and health insurance providers.", 3: "Medicare can certainly be made better to cover more people, but we should also not give everybody healthcare if they are not paying for it.", 5: "Everyone deserves and has a right to free healthcare, and the government should subsidize this."]
- **R2-Relevant (gun regulation):** Conservative to liberal. [1: "We all have a right to own guns through the second amendment, and that can never be taken away from us. We need to defend our freedom and liberty.", 3: "Gun rights are engrained into the culture, but at the same time

---

<sup>7</sup>For some of the LLMs, we make slight adjustments to the prompt phrasing, e.g., 'Your friend just posted on social media: "[text]" What is the exact message, with no extra words in your response, that you would respond on the post?' This is to ensure the LLM does not output extraneous text.

given the uptick in violence we should probably have some more regulations in place to control who has access.”, 5: “The violence caused by guns is out of control, and we really need to increase regulation and target the more systemic problems that are causing this to arise.”]

- **R2-Relevant (immigration):** Conservative to liberal. [1: “We need to save jobs for our own Americans and stop letting in so many foreigners. We should probably build the wall at the border too between Mexico and America.”, 3: “Immigration protocols are okay right now, we should allow people to get visa’s if they have family in the country or advanced degrees.”, 5: “We should be more open with our immigration policy, as a lot of the current legislation is just a manifestation of xenophobia.”]
- **R2-Relevant (abortion):** Conservative to liberal. [1: “I am pro-life and believe that abortion is murder. We should absolutely ban abortions without exception.”, 3: “Abortion should not be permitted after a while when the fetus is more developed, but in severe cases like rape I do believe it should be permitted.”, 5: “Abortion should be a choice left up to the person who is pregnant, because it is their body and they get to decide what happens to it.”]
- **R2-Relevant (climate change):** Conservative to liberal. [1: “Climate change does not feel like that big of a deal. People make using gas and coal out to be a worse thing than it really is.”, 3: “Climate change is going to affect us a bit in the future, so we should make sure to control our fossil fuel usage and maybe even carpool to work more.”, 5: “Climate change is an extremely serious problem that we need to confront through serious legislation if we want to make the Earth habitable for future generations.”],
- **R2-Relevant (criminal justice):** Conservative to liberal. [1: “Police keep us safe, and there is a lot of violence in the streets that we need to better control.”, 3: “When police are behaving well, they can help to keep neighborhoods safe. However, they should be responsible for any bad actions they may take.”, 5: “The criminal justice system in America is deeply flawed and racist. We need to defund the police and take away their power.”]
- **R3-Subjective (toxicity):** Toxicity. [1: “This message is not toxic, it is just somebody being friendly.”, 3: “I could see how it is a bit toxic to some people depending on how they are interpreting it.”, 5: “Toxic. This message is harmful and demeaning to others.”]
- **R3-Subjective (positive reframing):** TextBlob’s sentiment analysis with the following mapping: 1: [-1, 0), 2: [0, .25), 3: [.25, .5), 4: [.5, .75), 5: [.75, .1].
- **R4-Coverage (1):** Excitement. [1: “That sounds like it will be quite a bit of work, good luck packing and settling your affairs before heading out.”, 3: “This will be quite the adventure for you. Let’s catch up before you head out!”, 5: “Oh my gosh!! That sounds really amazing, I hope you have the best time!”]
- **R4-Coverage (2):** Harmful. [1: “The harms would not be that bad because it would be quite fun to use this.”, 3: “There are a number of concerns I would have with this app. We would need to be very careful about the different groups it could cause distress to.”, 5: “There are so many potential harms that could come from this, and they are severe enough that I hope such an app never exists.”]
- **R4-Coverage (3):** Permissible. [1: “We should be extremely wary of incorporating technology into therapy, as it is inherently a human-to-human process for which technology is likely to make things worse rather than better.”, 3: “There are potential benefits to incorporating technology and using the power of chatbots, so long as we do so pretty carefully.”, 5: “Leveraging the power of chatbots in therapies seems like a great use of advancing technology and could help to improve mental healthcare.”]

## D Prompt Phrasing Robustness

There are concerns that analyses of LLM responses could be artifacts of the particular phrasing used in the prompt. To ensure our results are not merely the result of such artifacts, we try out four different prompts to the LLM and visualize the TSNE of both the bag of words feature of n-grams ( $n=[1,2]$ ) as well as SBERT embedding. These are the two different versions of embedding we use for analysis in our work, so if different prompts result in similar embeddings here, it would indicate our findings based on such embeddings are robust across prompts.

We draw inspiration on our phrasing from [90], which prompts with individual names rather than demographic groups as we focus on.

The four prompt variations we try are: (1) “You are [identity] living in America. Speak exactly like you are [identity].”, (2) “You are [identity] living in America. Speak exactly like you are [identity].

Demographic Identity	In-Group	Out-Group
Millennial	7	4
man	5	0
woman	0	8
non-binary person	2	6

**Table 1: Human participant usage of LLM.** Based on our author-annotated estimates, the number of participants out of 100 that likely used an LLM to respond to our survey.

But remember, being [identity] is only one part of your identity.”, (3) “Speak from the perspective of [identity] living in America.”, and (4) “Speak from the perspective of [identity] living in America. But remember, being [identity] is only one part of your identity.” We only do this for the identity of Black women on the question of “Where do you like to vacation?” In Fig. 15 we find heavy overlap across all four prompts on all the models except for GPT-3.5, where Prompts 1 and 2 result in similar outcomes, but Prompts 3 and 4 are different from these. Llama-2 reflects a more minor version of this. Qualitative inspection on the models does not reveal notable differences. Ultimately we make the choice of using Prompt 3 for all four LLMs given that it is the simplest and most likely to represent actual use-cases. We note that given this robustness study, it is unlikely our results are an artifact of the prompt wording chosen, except in the case of GPT-3.5 where it may make a difference.

## E Noise in Human and LLM Generations

As with all datasets, our collected datasets of human-generated and LLM-generated responses likely have noise. Here, we describe our efforts to clean the datasets, and the kinds of noise we are aware of which remain.

**Refusals.** As part of alignment, LLMs will refuse to answer questions where a harmful response might be output. A refusal looks something like the following, based on GPT-4 prompted to respond like a White man: “As an artificial intelligence, I don’t have personal experiences, thus I can’t give a firsthand account of what it’s like to be a White man or any specific group in American society today.” We check for refusals on R1, R2, and R3 and encounter relatively few (< 5% for each model), but in the cases that we do, we rerun the question to give the LLM the benefit of the doubt, and create an upper bound for how well a current LLM is able to represent different perspectives.

**Cleaning identity markers.** To ensure that when we see a difference between responses from, e.g., women and men, it is not just because one person responds with “As a woman...” and another responds with “As a man...” we do our best to clean out these identity markers from the text before we perform our analysis. This was harder for behavioral personas where characteristics such as “I watch TV” would show up throughout the response in different ways, and explains some of the results in Fig. 6, where random personas had far higher covariance determinants on SBERT embeddings.

**Human Participant Usage of LLM.** Our work studies the desire of researchers and practitioners to replace human participants with LLMs. However, prior work has already found that human participants themselves are offloading their own requested tasks to LLMs, i.e., using Chat-GPT to respond to crowdworker tasks, at an estimated prevalence of 30% [91]. After reading through the set of LLM responses we had for the question from R1-Contingent, one author hand-labeled sets of human responses based on which appeared to be from an LLM. Eight sets were labeled, for in-group and out-group members of the following demographic identities: Millennial, man, woman, non-binary person, with findings in Tbl. 1. Unfortunately, a number of heuristics such as SBERT distance, ngram distance, or time taken by human participant were all insufficient as a threshold to filter out LLM responses, so we did not clean these from our dataset. We do not see humans using LLMs more than 10% in any of the scenarios we labeled, far lower than the estimated prevalence of 30%. We speculate this is because human participants may have actually wanted to answer our questions themselves, i.e., many responses asking, e.g., what it is like to be a non-binary person in American society today, were filled with emotional and personal anecdotes. We also note there is an interesting trend where it appears that non-women tend to use LLMs, whereas women almost never do.

## F Related Work

Here we engage more substantively with closely related work.

Santurkar et al. [76] study the political opinions of language models when steered towards 60 demographic groups. They find that while prompting with the demographic group does shift the LLM responses closer to that of the human group, it still does not entirely align them. We go further in this work by pursuing a larger set of questions (political opinion is a part of one of our four reasons for identity-prompting), as well as greater number of specific hypotheses which we tie to histories of harm. In their work, they ask multiple choice political questions, and thus miss out on other aspects of the response that we are able to capture, including reasonings behind a particular opinion and the syntactic differences between groups (e.g., “That’s wild, bro!” for Gen Z and “I’m like, YAASSSSS” and “That’s cray, hunty!” for Black women).

Cheng et al. [22] prompt LLMs with demographic identity and ask the models to describe themselves, comparing these responses with the default LLM in order to surface stereotypes. They then compare this differential to those discovered by Kambhatla et al. [38], finding that LLMs amplify the amount of stereotype. Our work is similar in that we find problems with identity-prompted LLMs, but different in both the range of tasks on which we find this to be true, as well as the types of problems studied. Whereas they only consider the task of an LLM describing itself, we consider four possible reasons an LLM might be prompted with identity, intending to encompass the full set of reasons, and thus having a far greater generalizability to all instances of identity-prompted LLMs. In terms of analysis, whereas they analyze specific stereotypes surfaced by the models, we focus on a different set of hypotheses regarding the misportrayal and flattening effects that are inherent to the training procedure of LLMs.

Another work, Cheng et al. [23] study the same premise as us: using LLMs to simulate different demographic identities for replacing human participants. They propose a framework to measure two criteria: individuation and exaggeration. Methods-wise, their measure of individuation best maps to one of the premises we establish of identity-prompting leading to a difference in, and their measure of exaggeration is different from our three primary analyses. Their exaggeration measure considers whether identity-prompted responses over-index on the identity compared to the topic prompted about. While motivationally this is similar to our concern about flattening groups, the way we operationalize this is completely different. In terms of contexts studied, their three scenarios of online social media forum, question-answering on political questions, and Twitter posts, all map to either our R2 or R4. We perform additional analyses on R1 and R3. Both of our works study open-ended responses as well as a range of demographic axes and identities; in our work we also conduct extensive human studies to compare our results to.

Sun et al. [24] and Beck et al. [25] explore the problem of identity-prompting LLMs for subjective annotation tasks, which maps to our R3-Relevant category. Their works use multiple choice responses from LLMs rather than engaging with open-responses as we do, and additionally focus more on “accuracy” of the LLM labels rather than the particular social harms that we do.


Each of the above works carefully grounds analysis in particular harms, and in doing so, necessarily is specific and does not cover the total range of harms. This is positive and commendable in the spirit of being more precise about where harms stem from [92]. Together, our work joins these to more collectively encompass the space of harms, as all are important to understanding the limitations of identity-prompted LLMs. They are complementary in strengthening the argument that identity-prompted LLMs should only be used with extreme caution.

## G Human Participant Demographics


Table 2 contains the self-reported gender, race, and age demographics of our 3200 human participants.

**Table 2: Demographics of Human Participants.** Each row indicates the self-reported gender, race, and age categories selected by the 100 human participants in each study. The gray rows indicate the study where out-group members for the identity were solicited. The numbers for some of the demographic axes add up to more than 100 because participants are able to check multiple genders and races, and others opted out and chose not to disclose their identity. For gender, G1=Woman, G2=Man, G3=Non-binary/third gender, G4=other. For race, R1=American Indian or Alaska Native, R2=Asian, R3=Black or African American, R4=Hispanic or Latinx, R5=Native Hawaiian or Other Pacific Islander, R6=White, R7=Other. For age in years, A1=18-24, A2=25-34, A3=35-44, A4=45-54, A5=55-64, A6=65-74, A7=75-84.

Study		Gender				Race							Age						
Axis	Iden	G1	G2	G3	G4	R1	R2	R3	R4	R5	R6	R7	A1	A2	A3	A4	A5	A6	A7
	Black man	0	100	0	0	0	0	100	1	0	0	0	9	40	18	21	10	2	0
	Black man	60	37	1	0	1	7	9	9	0	83	1	8	26	34	12	14	5	1
	Black woman	100	0	0	0	0	0	100	1	0	0	0	4	14	21	33	18	8	2
	Black woman	52	48	0	0	3	9	5	9	0	83	2	10	42	18	11	13	6	0
	White man	0	100	0	0	0	0	0	3	0	100	1	10	28	32	16	7	6	1
intersect	White man	59	41	0	0	0	19	26	13	0	47	0	10	35	27	16	8	3	1
	White woman	100	0	0	0	1	0	0	5	0	100	0	11	17	23	21	16	9	3
	White woman	41	58	1	0	1	37	16	10	0	39	0	19	43	15	15	5	2	1
race	Black	47	52	1	0	2	0	100	3	0	1	1	11	33	23	22	11	0	0
	Black	51	43	4	0	0	19	0	5	2	84	2	14	32	20	17	12	3	2
	Asian	36	62	0	0	0	100	0	0	0	4	0	22	37	22	12	6	1	0
	Asian	51	47	2	0	5	0	24	14	0	82	0	11	32	22	19	7	9	0
	White	67	31	2	1	0	0	0	2	0	100	0	12	22	30	19	13	2	2
	White	40	56	0	0	3	28	48	20	0	1	3	8	40	31	17	4	0	0
gender	man	0	100	0	0	1	9	9	11	0	74	2	15	38	22	16	6	3	0
	man	100	0	1	0	1	13	7	12	0	77	1	18	36	20	14	8	3	1
	woman	100	1	0	0	2	6	13	10	0	81	2	9	29	20	18	19	4	0
	woman	0	99	1	0	2	10	9	12	0	73	2	14	34	22	18	7	4	1
	non-binary	4	4	100	0	0	9	6	14	1	80	5	31	50	10	4	4	1	0
	non-binary	48	52	0	0	2	8	17	9	0	69	1	9	40	21	20	6	3	1
age	Baby Boomer	66	34	0	0	1	1	4	3	0	91	2	0	0	0	0	54	42	4
	Baby Boomer	43	51	6	0	2	11	10	10	0	76	2	15	31	31	19	4	0	0
	Millen	40	59	0	0	4	8	5	19	1	79	2	0	64	36	0	0	0	0
	Millen	53	46	1	0	0	5	12	12	0	74	0	37	19	1	17	17	8	1
	Gen Z	46	53	2	1	0	23	13	18	2	59	5	71	29	0	0	0	0	0
	Gen Z	47	51	1	0	2	5	12	5	0	80	0	0	38	30	16	8	8	0
disability	w/o disabilities	43	56	1	0	0	10	8	8	0	78	0	14	21	23	21	17	4	0
	w/o disabilities	62	36	3	1	3	10	7	8	2	79	3	8	24	25	20	20	2	1
	w/ ADD or ADHD	40	56	4	0	4	10	9	10	1	75	2	16	44	25	8	5	0	2
	w/ ADD or ADHD	43	53	3	1	0	10	7	13	0	72	1	20	21	20	18	12	7	2
	w/ impaired vision	49	45	6	0	4	6	14	16	0	78	2	7	35	21	16	13	7	1
	w/ impaired vision	47	44	7	0	2	4	12	6	0	76	3	7	31	21	13	17	9	2




(c) GPT-3.5-Turbo.




(d) GPT-4.


**Fig. 7: LLMs compared to out-group imitations and in-group portrayals.** Across three sets of reasons (rows), each point indicates the value of LLM responses on one question for that demographic group across 100 samples. Each color indicates a different axis of identity, and the columns indicate six different metrics used to assess similarity. Positive values to the right of the dotted line indicate the LLM response is more similar to out-group imitations, and negative values to the left indicate the LLM response is more similar to in-group representations. Circles indicate statistical significance with  $p < .05$  and crosses indicate otherwise. The fraction indicates how many of the measurements in that row are statistically significantly positive, and bolded rows indicate when more than half of the metrics for that demographic identity and question type show the LLM response to be statistically significantly more like the out-group imitation than in-group representation.




(a) Llama-2.



(b) Wizard Vicuna Uncensored.




(c) GPT-3.5-Turbo.




(d) GPT-4.


**Fig. 8: Identity-coded names compared to explicit identity label.** Same interpretation as Fig. 7, where for two sets of reasons (rows), each point indicates the value on one question for that demographic group across 100 samples. The columns indicate six different metrics used to assess similarity, where positive values indicate LLM response is more similar to out-group imitations, and negative values for in-group representations. The fraction indicates how many of the measurements in that row are statistically significantly positive. For each identity, the prompt contains the explicit identity label (Iden), or one of the two identity-coded names (Name 0 or Name 1). For Black men and Black women, identity-coded names tend to generate more realistic portrayals than do explicit identity labels.




(a) Llama-2.



(b) Wizard Vicuna Uncensored.




(c) GPT-3.5-Turbo.




(d) GPT-4.

**Fig. 9: LLMs flatten groups.** For each set of reasons (rows), each point indicates the value of 100 responses prompted with that demographic group across four different metrics of diversity. 95% confidence bars are provided, and the dark gray points indicate human participant in-group responses, while colored points represent LLM responses. Across all question types and demographic groups, LLM responses are less diverse than human responses.




**Fig. 10: Temperature hyperparameter does not solve flatness for Wizard Vicuna Uncensored.** Same interpretation as Fig. 4 in the main text: comparison of human in-group diversity to Wizard Vicuna Uncensored generations varying levels of temperature settings, where by 1.8 the responses become incoherent. There are 100 samples or generations per scenario, and 95% confidence intervals are generated via bootstrapping. At this setting even though the unique n-gram metric shows the LLM surpassing humans in diversity, this is only due to the incoherence as under no other semantic metric is human diversity reached.




(a) Llama-2.

(b) Wizard Vicuna Uncensored.




(c) GPT-3.5-Turbo.


**Fig. 11: Response coverage is high without essentializing identity.** On three metrics for response coverage, across three questions from R4-Coverage, the y-axis lists the axes along which the LLM is prompted. Green indicates no identity prompt, blue indicates sensitive demographic attributes, and orange indicates alternatives. Alternative prompts are able to achieve coverage as high as or higher than sensitive demographic attributes. Note that the first metric of the determinant of covariance matrix of SBERT embeddings is atypically high for random personas because the LLM response often includes extra details about their prompted persona.




(a) Llama-2.



(b) Wizard Vicuna Uncensored.




(c) GPT-3.5-Turbo.




(d) GPT-4.


**Fig. 12: Average multiple choice responses for LLMs and human participants.** LLM responses are indicated in colors, human in-group in black, and human out-group in gray. For R1-Contingent, 1 to 5 represents how challenging it is to have that demographic identity; R2-Relevant 1 to 5 represents conservative to liberal political opinion; R3-Subj represents level of toxicity detected for race and level of positive sentiment detected for gender. There are 100 generations or samples per scenario, and 95% confidence intervals are generated through bootstrapping. Along the demographic axes of race and intersectional, out-group human participants tend to overinflate the difficulty of being in that group compared to in-group human participants, and LLMs even further inflate the difficulty beyond that of the out-group. GPT-4 inflates the difficulty more for White men and White women compared to Black men and Black women.  
24



**Fig. 13: LLMs answer differently depending on what demographic identity they are prompted with.** For three sets of question reasons (rows), the difference between a pair is shown. Black dots indicate human in-group participants, gray dots indicate human out-group participants, and the colored dots indicate LLM responses. Bolded rows indicate that on more than half of the measured values, the difference between the two compared groups is statistically significant.



**Fig. 14: Comparison of differences between human in-group representations and out-group imitations.** For each set of reasons (rows), the demographic group is shown with more positive values indicating difference between in-group and out-group human participants. Circles indicate statistical significance, crosses do not.



**Fig. 15: Analysis of prompt phrasing variations.** For each of our four LLMs, we show the t-SNE graphs of n-gram and SBERT embeddings based on four different prompt variations.