

One fish, two fish, but not the whole sea: Alignment reduces language models’ conceptual diversity

Sonia K. Murthy^{1,2}, Tomer Ullman^{2,3}, Jennifer Hu²

¹School of Engineering and Applied Sciences, Harvard University

²Kempner Institute for the Study of Natural and Artificial Intelligence, Harvard University

³Department of Psychology, Harvard University

soniamurthy@g.harvard.edu, {tomerrullman, jenniferhu}@fas.harvard.edu

Abstract

Researchers in social science and psychology have recently proposed using large language models (LLMs) as replacements for humans in behavioral research. In addition to arguments about whether LLMs accurately capture population-level patterns, this has raised questions about whether LLMs capture human-like conceptual diversity. Separately, it is debated whether post-training alignment (RLHF or RLAIFF) affects models’ internal diversity. Inspired by human studies, we use a new way of measuring the conceptual diversity of synthetically-generated LLM “populations” by relating the internal variability of simulated individuals to the population-level variability. We use this approach to evaluate non-aligned and aligned LLMs on two domains with rich human behavioral data. While no model reaches human-like diversity, aligned models generally display less diversity than their instruction finetuned counterparts. Our findings highlight potential trade-offs between increasing models’ value alignment and decreasing the diversity of their conceptual representations.

1 Introduction

As large language models (LLMs) have become more sophisticated, there has been growing interest in using them to replace human labor. This appeal of LLMs has even made its way to settings where *human behavior itself* is the object of inquiry: recently, researchers have proposed that LLM-generated responses can be used in place of human data for tasks such as polling, user studies, and behavioral experiments (e.g., Aher et al., 2023; Argyle et al., 2023; Hämmäläinen et al., 2023; Binz and Schulz, 2024; Manning et al., 2024). If possible, using a synthetic replacement for the process of human data collection could be transformative for a variety of human factors disciplines, from political science to economics and psychology.

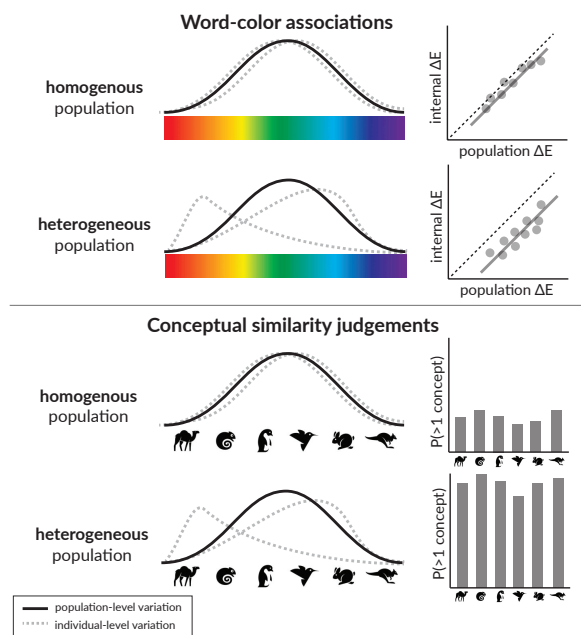


Figure 1: We investigate LLM populations comprised of simulated individuals in two domains: color associations (top) and concept similarity (bottom). In both domains, there is both individual- and population-level variation. It is possible that individual variation overlaps with the population average (homogeneous population) or separates from it (heterogeneous population). Our experiments are designed to tease these two options apart.

While prior work suggests that LLMs can capture certain behavioral patterns, there are ongoing debates as to whether they are valid replacements for human subjects (Dillion et al., 2023; Wang et al., 2024a; Park et al., 2024). One key issue is whether LLMs capture **conceptual diversity**: the variation among *individuals’* representations of a particular domain. A natural way to study conceptual diversity is to study the variability in LLMs’ response distributions at the *population* level – that is, by considering averages across simulated individuals. However, population-level variability can be a flawed measure of conceptual diversity for several reasons. First, it assumes certain charac-

teristics about the nature of people’s mental representations that may not hold in practice – for example, that people have high certainty in these representations (e.g., Martí et al., 2018), or that people’s responses are a deterministic best-guess, with no additional information contained in successive guesses (e.g., Vul and Pashler, 2008). Second, assessing population-level variation without individual-level variation can mask important information about the population, particularly whether it is *homogeneous* (comprised of individuals who share similar underlying representations) or *heterogeneous* (comprised of individuals with meaningfully different representations) (see Figure 1).

Conceptual diversity may be negatively impacted by **alignment**, which further complicates the picture. Post-training alignment techniques, such as RLHF (Ouyang et al., 2022; Bai et al., 2022a) and RLAIIF (Bai et al., 2022b), are now standard parts of LLM development, and presumed to contribute to the human-like abilities of models (Ji et al., 2024). However, it has also been observed that “aligned” models show biases that can limit the lexical and content diversity of their outputs (Janus, 2022; Padmakumar and He, 2024; Park et al., 2024; O’Mahony et al., 2024). It also remains unknown whether alignment to synthetic preferences (instead of human preferences) might worsen these biases, as it is possible these models “collapse” when recursively trained on synthetically generated data (Shumailov et al., 2024; Gerstgrasser et al., 2024).

In this paper, we (1) investigate the ability of LLMs to capture human-like conceptual diversity, and (2) analyze the effect of alignment by comparing conceptual diversity across non-aligned models to models aligned using RLHF or RLAIIF. In our experiments, we first simulate populations of unique individuals in LLMs using two techniques proposed by prior literature: temperature- and prompt-based manipulations. We then evaluate ten open-source LLMs on two domains with rich human behavioral data: word-color associations, and conceptual similarity judgments. While there is no single agreed-upon metric for capturing conceptual diversity, we consider two different metrics that are each well-suited to their respective domains, to measure the conceptual diversity of synthetically-generated LLM “populations”.

We find that no model approaches human-like conceptual diversity. Further, aligned models generally display less conceptual diversity than their non-aligned, fine-tuned counterparts. Our findings

suggest that there may be trade-offs between increasing model safety in terms of value alignment, and decreasing other notions of safety, such as the diversity of thought and opinion that models represent. We caution that these trade-offs should be better understood before models are used as replacements for human subjects, or deployed in human-centered downstream applications.

2 Background

Recent works have proposed using LLMs as stand-ins for humans in many applications and domains, including opinion surveys and polling in political science (Argyle et al., 2023), user studies in human-computer interaction (Hämäläinen et al., 2023), annotation tasks (Gilardi et al., 2023), and various economic, psycholinguistic, and social psychology experiments (Aher et al., 2023; Dillion et al., 2023).

Several works have responded by identifying LLMs’ lack of answer diversity as a potential harm of using LLMs to replace human subjects. These works have focused on whether models are capable of capturing demographic- and subgroup-level variation in settings where accurately simulating the social identities of a human population of interest is directly relevant to the task. Wang et al. (2024a) find that inference-time interventions to improve models’ output diversity do not prevent LLMs from misportraying and flattening the representations of already marginalized demographic groups. Even after explicit steering toward certain demographics, LLMs still do not capture these group’s responses on public opinion surveys (Santurkar et al., 2023).

These observations have motivated recent attempts to go from broad, population-level observations of diversity, to metrics that assess finer-grained notions. For example, He-Yueya et al. (2024) measured the extent to which LMs capture human knowledge distributions on real-world domains like language learning and mathematics. Their metric captures whether two populations are similarly sensitive to the relative difficulty of test items. However, this metric is not well-suited for domains that lack ground-truth answers or where people’s mental representations are not easily parameterized. Franke et al. (2024) propose another item-level account of diversity, from the perspective of Bayesian statistical modeling. They find that LLMs do not capture human variance at the item-level in an experiment on pragmatic language, and aggregate item-level predictions to replicate

Base model	Non-aligned variant	Aligned variant	Alignment method
Mistral (mistralai/Mistral-7B-v0.1)	Openchat (openchat/openchat_3.5)	Starling (berkeley-nest/Starling-LM-7B-alpha)	RLAIF (APA)
Mistral (mistralai/Mistral-7B-v0.1)	Mistral-Instruct (mistralai/Mistral-7B-Instruct-v0.1)	Zephyr-Mistral (HuggingFaceH4/zephyr-7b-beta)	RLAIF (DPO)
Gemma (google/gemma-7b)	Gemma-Instruct (google/gemma-7b-it)	Zephyr-Gemma (HuggingFaceH4/zephyr-7b-gemma-v0.1)	RLAIF (DPO)
Llama (meta-llama/Llama-2-7b-hf)	Llama (meta-llama/Llama-2-7b-hf)	Llama-Chat (meta-llama/Llama-2-7b-chat-hf)	RLHF (PPO)
Llama (meta-llama/Llama-2-7b-hf)	Tulu (allenai/tulu-2-7b)	Tulu-DPO (allenai/tulu-2-dpo-7b)	RLAIF (DPO)

Table 1: Pairs of non-aligned and aligned models tested in our experiments. Huggingface identifiers are shown in parentheses underneath model names. The base models are shown for reference and are not directly tested, with the exception of the comparison between Llama and Llama-Chat. All models have 7B parameters.

condition-level effects seen in human data. Similarly, Wang et al. (2024c) also find that LMs replicate broad-level patterns (e.g., mean and standard deviation) but fail to capture item-level patterns of human behavior. These studies highlight the importance of finer-grained comparisons between model and human outputs, but leave open the question of how to measure the relationship between individual-level variation and population-level variation in rich conceptual domains.

3 Methods

Our experiments are designed to investigate two related research questions: (1) Do modern LLMs capture the conceptual diversity of human populations, i.e., **individuals**’ variability in conceptual representations? and (2) How does post-training alignment (here, RLHF or RLAIF) affect models’ ability to capture this feature of response distributions from human populations?

The basic logic of our experiments is as follows. First, we identify domains where prior work has demonstrated rich conceptual diversity in humans. Second, we test whether LLMs capture conceptual diversity patterns by simulating “populations”, using temperature- and prompting-based manipulations. Finally, we compare the conceptual diversity achieved by models against that of humans, and specifically examine the effect of alignment on the degree of fit between models and humans.

3.1 Domain 1: Word-color associations

The first domain probes people’s intuitive associations between words and colors, building on the experiments of Murthy et al. (2022). In this task, human participants were presented with a target word

(such as “chalk” or “obligation”) and asked to click on the color most associated with the word, from a set of 88 color chips. Participants were tested in two blocks, ultimately providing two color associations for each word. This design enables measures of *population* variability (how much color associations vary across individuals) as well as *internal* variability (how much they vary in an individual’s own representation).

Model evaluation. To measure which color a model “associates” with a particular target word, we used the following query:¹

Question: What is the HEX code of the color you most associate with the word [WORD]? You must respond with a guess, even if you’re unsure. Make sure your response contains a valid 6-digit HEX code.

Conditioned on this input, models were allowed to freely generate subsequent text.² Two samples were elicited for each word and HEX codes were extracted from models’ generations using a regular expression.³ This procedure was repeated to collect models’ color associations for the 199 words tested by Murthy et al. (2022), for a total of 150 simulated “subjects” (i.e., 150 unique prompting contexts for the relevant experiments).

Diversity metric. In order to examine conceptual diversity in any domain, including color-

¹For the experiments that simulate unique subjects through prompting manipulations, the relevant context was prepended to this query (see Section 3.3 for more details).

²While the human baseline was collected through interaction with a visual color space, we adopt this textual format following e.g., Marjeh et al. (2023); Niedermann et al. (2024)

³When one or both samples contained no valid HEX code, the response was considered invalid and was excluded from analyses.

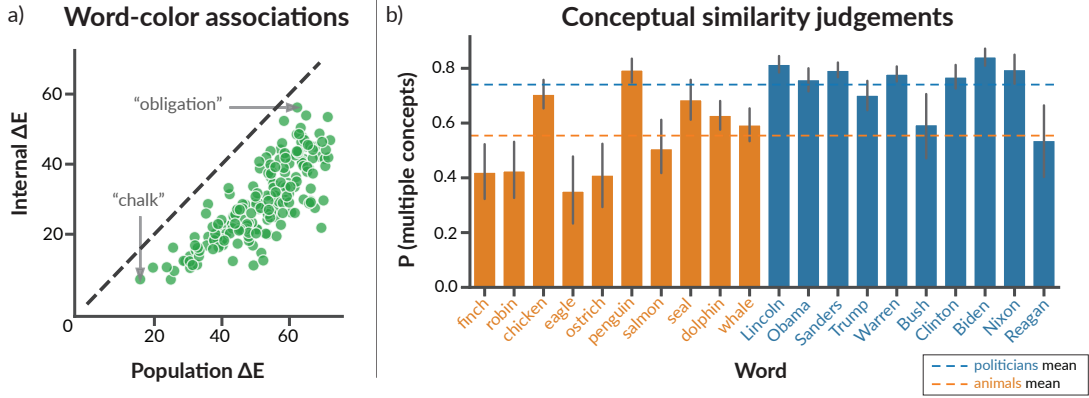


Figure 2: Human baselines in both domains. (a) Internal variability (y-axis) versus population variability (x-axis) for human participants in the word-color association domain. Reproduced from Murthy et al. (2022). (b) Probability of more than one conceptual representation, estimated using Chinese Restaurant Process model on human data for conceptual similarity judgement task (Martí et al., 2023).

associations, we need a diversity metric over the space. Here, following Murthy et al. (2022), we measure the variability between two colors (denoted by ΔE) by their perceptual similarity in CIELAB space. For a given word w , we compute internal and population measures of variability as follows. $\Delta E_{internal}(w)$ measures the ΔE between the colors chosen by a participant in the first and second blocks for word w . $\Delta E_{population}(w)$ captures the similarity between different individuals’ color responses for word w (i.e., the average pairwise similarity among every pair of color responses provided by different participants).

For each word $w \in W$, the heterogeneity of a population (i.e., a set of subjects S) is then captured by the following metric:

$$d_w = \frac{1}{S} \sum_{s \in S} \frac{|\Delta E_{internal}(s) - \Delta E_{population}(s)|}{\sqrt{2}} \quad (1)$$

Intuitively, this metric measures the divergence from the “line of unity” – i.e., the signature of a homogeneous population in which individuals share the same underlying distributions. The greater this metric, the more conceptual diversity, since the variability between different individuals’ distributions is greater than the variability of the individual’s own internal distribution (Figure 1, top).⁴

⁴The reverse scenario, where $\Delta E_{internal} > \Delta E_{population}$ is also possible, though highly unlikely, especially as the sample size grows. Because the different participants’ associations are independent samples, the probability of sampling similar colors (low $\Delta E_{population}$) from different individuals’ highly variable internal probability distributions (high $\Delta E_{internal}$) is low.

Human baseline. The relationship between $\Delta E_{internal}$ and $\Delta E_{population}$ for human participants is shown in Figure 2a, with data reproduced from Murthy et al. (2022). While these two measures are correlated for humans, internal variability is overall lower than population variability, consistent with a heterogeneous population. For the human data, the average of the metric in Equation (1) = 15.82 (95% CI: [15.18, 16.58]).

3.2 Domain 2: Conceptual similarity judgements

The second domain, inspired by Martí et al. (2023), probes the number of latent concept clusters in a population. In this task, participants were asked for similarity judgments between concepts in a particular category (e.g., “Which is more similar to a finch, a whale or a penguin?”). Two categories were tested: common animals, which are more likely to have shared conceptual representations across individuals; and United States politicians, which are likely to have varying representations across individuals.

Model evaluation. To measure similarity judgments within concept categories, we presented models with the following query:

Question: Which is more similar to a [TARGET], [CHOICE1] or [CHOICE2]? Respond only with “[CHOICE1]” or “[CHOICE2]”.

where [TARGET] is one of the 10 words from the relevant category (animals or politicians), and [CHOICE1] and [CHOICE2] cover all combinations of the remaining words therein. In most cases, model output only one of the two answer choices, but for outputs where additional text was generated,

the first sentence of the generation was selected (as it usually contained the model’s response) and string matching was used to identify the model’s choice. If neither choice was present, the model’s response for that query was excluded from further analyses.

Diversity metric. Here, the heterogeneity or diversity of the population intuitively corresponds to the number of latent concepts for each word among individuals in the population. To measure this formally, the similarity judgment for a given word was coded as a binary vector indicating the similarity rating between every other pair of items. Following [Martí et al. \(2023\)](#), the responses were then analyzed using a non-parameteric Bayesian clustering model (Chinese restaurant process; CRP) to model a distribution over the number of clusters (concepts) for each word. For our analyses, we implement the version of this model that implements a prior preference for fewer clusters and runs inference using a Gibbs sampler.

For each word, the posterior distribution fitted by the CRP was then used to estimate the probability of only a single conceptual representation exists in the population of sampled individuals. The higher this probability is, the less diverse the population is. Therefore, as our measure of population heterogeneity in this domain, we compute $1 - P(\text{one concept})$, or equivalently, $P(\text{multiple concepts})$.

Human baseline. [Martí et al. \(2023\)](#) find that words in the “politicians” category are far more likely to have multiple meanings than “animal” words (politicians: 0.69 (95% CI: [0.66, 0.72]); animals: 0.43 (95% CI: [0.39, 0.48]); see [Figure 2b](#)). This result reflects the conceptual diversity that is inherent to human populations even in the simplest domains: there exist multiple latent meanings and representations even for words with very concrete, real-world referents.

3.3 Simulating populations within models

Our main question is whether models can capture human-like conceptual diversity, for the purposes of simulating the response distributions of unique human subjects. To investigate this question, we tested two ways of manipulating the “diversity” of responses generated by models: increasing the softmax temperature parameter, and conditioning on different types of contexts. These techniques were

chosen because of their popularity in related literature, and their accessibility as inference-time interventions for downstream users.

Method 1: Increasing temperature. Increasing the temperature in the softmax normalization during inference-time decoding effectively increases the entropy of the distribution over tokens. Prior work has explored this method as a way of increasing the diversity of generated output ([Yu et al., 2023](#); [Tevet and Berant, 2021](#); [Peeperkorn et al., 2024](#); [Bellemare-Pepin et al., 2024](#)). Most relevant to our work, [Wang et al. \(2024a\)](#) explore different temperature values as a way to overcome models’ flattening of the multi-faceted nature of demographic groups. For completeness and to see how much we can push our measure of diversity, we test some higher values of temperature settings for each model ($t \in [t_0, 1.5, 2.0]$, where t_0 is the default temperature of the base model⁵).

Method 2: Prompting. We also tested the effect of “persona prompting”, where the prompt explicitly instructs the model to simulate a particular individual or group of individuals, usually using group-specific demographic information. For example, [Wang et al. \(2024a\)](#) use persona prompts of the form “Speak from the perspective of [identity] living in America”, and [He-Yueya et al. \(2024\)](#) use more complex prompts such as “Pretend that you are an 11-year-old student. Your gender is female. You are eligible for free school meals due to being financially disadvantaged.” While both studies found that such identity-based prompts do not capture relevant metrics to the extent seen in human populations, other work has found success with prompts that simulate expertise and developmental stages ([Salewski et al., 2024](#)).

For our experiments, we adopt the demographic-prompting paradigm, as this approach most closely mimics the process of gathering a representative sample of adult populations for behavioral experiments in our domains. We prepended persona contexts of the following form to the evaluation query: “You are a [race] [gender] from [hometown] in [state] who is [age] and works as a [occupation].” We additionally considered three control conditions: *none*, where no context was prepended to the query (matching the prompting setup from the temperature-based

⁵Mistral has a default temperature of 1.0, Llama a default of 0.9, and Gemma a default of 0.7.

manipulations described above); *random*, where we prepended contexts matched in length to the persona-based contexts but on unrelated topics; and *nonsense*, where we scrambled these unrelated contexts at the word-level (see appendix A for details). The latter two conditions were designed to test whether adding randomness in the form of conditioning models’ outputs on highly surprising tokens could sufficiently simulate meaningful human diversity

This gave us a total of four prompting conditions, for which the number of prompts varied depending on the domain so as to most faithfully replicate the participant sampling procedures for each experiment ($N_{color} = 150, N_{concept} = 1800$). In this setting, we used each model’s default temperature.

3.4 Models

We test five pairs of open-source non-aligned and aligned model variants, spanning across three families of base models: Mistral (Jiang et al., 2023), Gemma (Gemma Team et al., 2024), and Llama (Touvron et al., 2023). We generally use pairs such that the non-aligned and aligned variants are both fine-tuned on the same base model. This allows us to factor out the effects of the underlying architecture and pre-training, focusing only on the contrast between alignment and other forms of fine-tuning.

Our tested models are summarized in Table 1. Within the Mistral base family, we test two pairs of non-aligned and aligned models: Openchat (Wang et al., 2024b) vs Starling (Zhu et al., 2024), and Mistral-Instruct (Jiang et al., 2023) vs Zephyr-Mistral (Tunstall et al., 2024). For the Gemma base family, we test Gemma-Instruct (Gemma Team et al., 2024) vs Zephyr-Gemma (Tunstall et al., 2024). And for the Llama base family, we test two pairs of aligned and non-aligned models: Llama (base) vs Llama-Chat (Touvron et al., 2023), and Tulu vs Tulu-DPO (Iverson et al., 2023).

There are a few exceptions to the general grouping pattern described above. First, Starling undergoes alignment on top of the already-reinforcement-learning-fine-tuned (RLFT) Openchat model, instead of the shared base Mistral model.⁶ And second, Llama-Chat combines both instruction fine-tuning and alignment into the same model, leaving its non-aligned counterpart, Llama, as the only

⁶Because these models are still separated by one degree of alignment, like the other model pairs we test, we still refer to this pair as non-aligned vs. aligned in our results for consistency.

model in our suite that has also *not* been instruction fine-tuned after pre-training.

Our tested models span across two popular alignment methods: Reinforcement Learning from Human Feedback (RLHF; Bai et al., 2022a; Ouyang et al., 2022), and Reinforcement Learning from AI Feedback (RLAIF; Bai et al., 2022b). RLHF involves collecting preference data from human crowdworkers, which is then used either to derive a reward model for subsequent fine-tuning (as in PPO, Schulman et al. 2017; or APA, Zhu et al. 2023) or for direct, reward-free fine-tuning (as in DPO; Rafailov et al., 2023). RLAIF involves a similar process, but models are aligned to synthetically generated data instead of human preferences.

In order to isolate the effect of alignment/training methods, we control for model size, testing the 7B-parameter version of all models. All models were downloaded from HuggingFace.

4 Results

4.1 Word-color associations

Figure 3 shows the main results from the word-color association domain. The y-axis shows the heterogeneity metric d_w (Equation (1)) averaged over all tested words. We omit the results from two models due to a high number of invalid responses among their generations: Llama and Tulu (see Appendix Figure 5 for counts of invalid responses).

Overall, we find that none of the models reach the heterogeneity of the human population, with low values for the diversity metric detailed in Equation (1) compared to the human baseline of 15.82. To more closely assess the effects of temperature and prompting manipulations on models’ conceptual diversity we constructed a regression model predicting our diversity metric from these conditions with random effects of LLM. In a combined model including both predictors we find a main effect of both prompt and temperature manipulation ($p < 0.001$ for both), but a much higher coefficient for prompts ($b = 0.294$) than for temperatures ($b = 0.093$), with minimal between-group variability among individual LLMs ($b = 0.177, SE = 0.058$). These findings suggest that adding noise through prompting techniques is more effective in increasing population heterogeneity than adding noise via softmax temperature.

Turning to the effect of alignment (i.e., comparing light and dark bars within each facet), we find that the aligned model variants have similar or

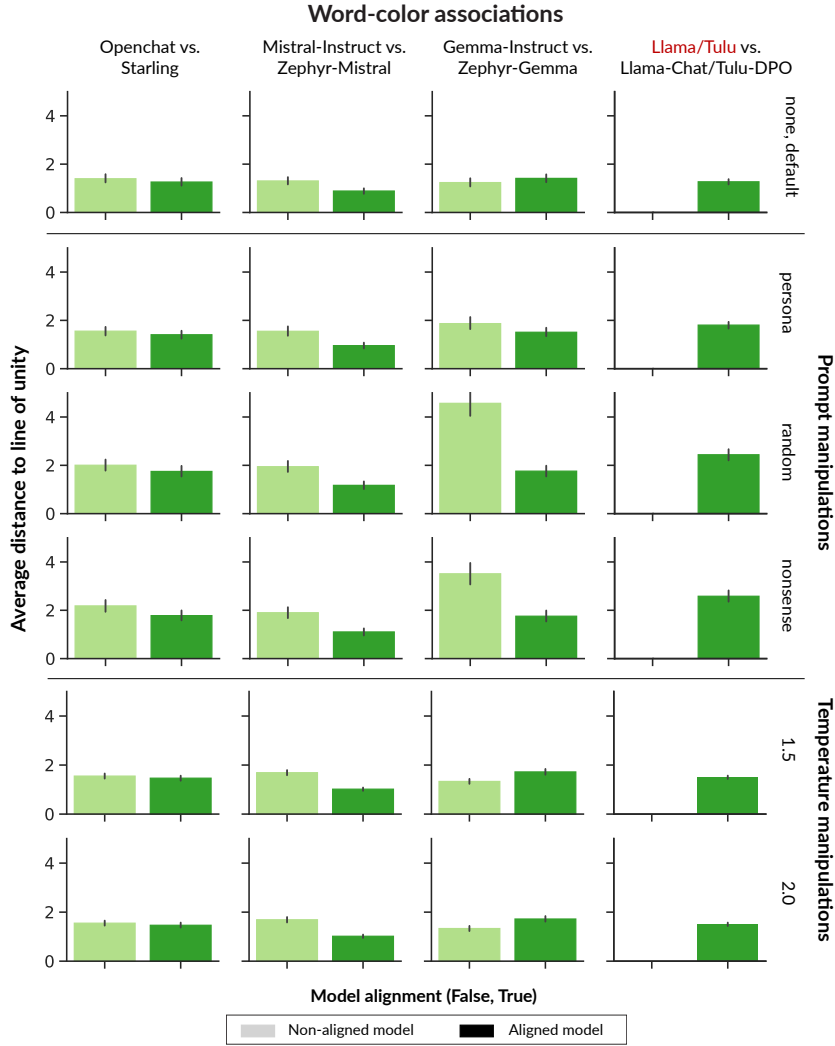


Figure 3: Heterogeneity of simulated LLM population in word-color association domain for prompting and temperature manipulations. The y-axis indicates d_w (Equation (1)) averaged over words. Rows = baseline, followed by prompting and temperature conditions; columns = model families. Darker bars indicate aligned models. For reference, the human baseline (Murthy et al., 2022) is 15.82.

lower heterogeneity than their non-aligned counterparts, with the exception of the Gemma models under the temperature manipulations. This is reflected in the results of a regression model predicting our diversity metric from the absence or presence of model alignment with random effects of model family (columns in Figure 3). We find a significant negative main effect of alignment ($b = -0.495, p < 0.001$), particularly contributing to a decrease in our diversity metric in this domain, with minimal between-group variability among model families ($b = 0.114, SE = 0.055$).

We also qualitatively analyzed the similarity between models’ and humans’ word-color associations. Figure 9 shows human and model distributions over color space for three example words and the *persona* prompting manipulation relative

to the no-prompt baseline. The similarity between models’ and humans’ color associations also corroborates prior work showing that text-only models can learn structure in perceptual domains such as color space (e.g., Abdou et al., 2021; Patel and Pavlick, 2022).

4.2 Conceptual similarity judgments

Figure 4 shows the main results from the conceptual similarity domain. We omit the results from Llama in this domain due to a high number of invalid responses among their generations⁷.

We find that in this domain as well, models

⁷Specific models and manipulations (e.g. high temperature settings) with a high number of invalid responses were also removed post-clustering due to insufficient data and are indicated by “NA”.

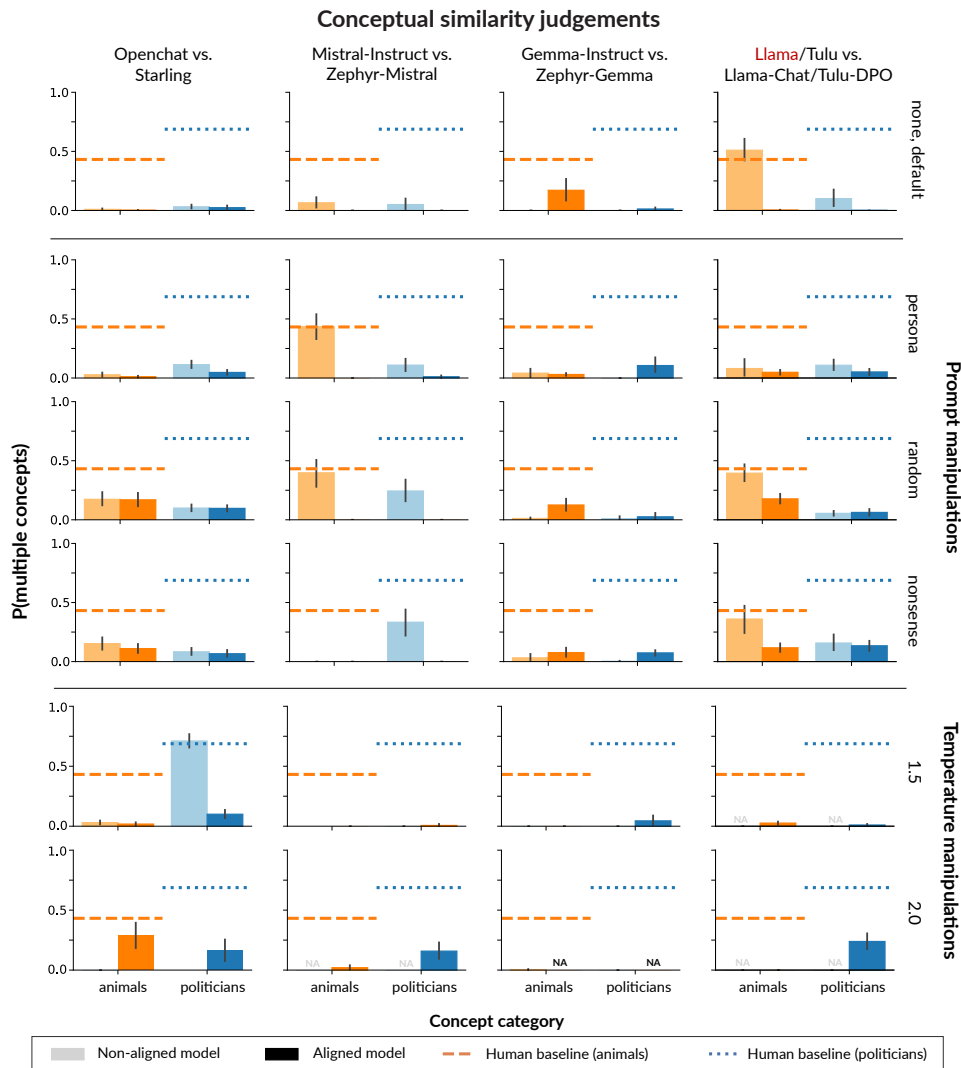


Figure 4: Heterogeneity of simulated LLM population in conceptual similarity domain. The y-axis indicates the probability of more than one conceptual representation, estimated using Chinese Restaurant Process model. Rows = baseline, followed by prompting and temperature conditions; columns = model families. Darker bars indicate aligned models. Human baselines for each conceptual category (Martí et al., 2023) are shown as horizontal lines.

largely fall short of the heterogeneity of the human population (horizontal lines), with only a few non-aligned models reaching the human baseline. Across models and concept words, the CRP estimates a very low probability that multiple conceptual representations exist among the population. Qualitatively, in Figure 4, we see that with the exception of the Gemma-Instruct and Zephyr-Gemma pair (and in some of the temperature manipulation settings), this effect is more pronounced among the models that have been aligned. The results of a regression model show a significant main effect of alignment ($b = -0.079, p < 0.001$) on the probability of multiple concepts existing among an LLM-simulated population, with minimal variability across model families ($b = 0.002, SE =$

0.009).

The results of a combined regression model predicting our diversity metric in this domain from both predictors of prompt and temperature manipulations (with random effects of model family) showed small, though significant effects of both manipulations ($b = 0.019, p < 0.001$) and ($b = -0.018, p < 0.001$), respectively. This result corroborates that of the word-color domain, where prompting manipulations seem to have a greater impact on increasing models' conceptual diversity than temperature manipulations, though the effects are small.

Interestingly, the models also fail to capture another signature pattern in the human data: increased diversity for politician words than animal words

(cf. Figure 2b). The heterogeneity measure generally stays the same or even slightly decreases between animals and politicians (with a few exceptions). It is reasonable to expect that text-only models would be able to learn conceptual similarity relationships based on the contexts in which the terms appear, and yet, we do not find greater diversity for politician words. We speculate that this could be related to political concepts being more important “targets” for alignment, as they may be associated with more extreme or emotionally charged language.

5 Discussion

In both domains, the base Llama and Tulu models generated a high number of invalid responses (text generations that did not mention a valid color/target word). We hypothesize that for Llama, as the only model in our suite that had not been instruction fine-tuned, this indicates a fundamental difficulty comprehending the task instructions. However, for Tulu, (which had been instruction fine-tuned, but not aligned) this could indicate weaker understanding of subjective, non-fact-based domains like the ones we test. Further, our findings that models aligned through both RLHF and RLAIIF display less conceptual diversity suggests that the preference information captured by pairwise judgements in such feedback datasets could be too narrowly constraining the post-training latent space of models. Exploring alternative paradigms of value alignment that allow for more complete pictures of human values (cf. Zhi-Xuan et al.) and further studying models’ internals pre- and post-alignment could help overcome such issues.

Our findings also corroborate those of recent work in many domains that suggest that injecting variability in models’ outputs through temperature and prompt manipulations does not do enough to induce meaningful variability on relevant metrics, and often just leads to highly incoherent outputs (invalid responses). One reason for this could be that increasing the entropy of the output distribution or conditioning on highly surprising tokens will increase models’ uncertainty over subsequent tokens in a domain-general sense. However, the kind of uncertainty that gives rise to meaningful individual differences in humans is likely much more constrained to a particular task and cognitive domain than this. Exploring more task-specific methods for injecting *structured* randomness into models’

generations could help LLMs to overcome their population-average behavior in ways that more meaningfully simulate the cognitive differences in human individuals.

This raises further questions of what factors give rise to individual differences in humans and whether these are the same attributes that can result in individual-like models. While it is common to obtain representative samples of a human population for studies in cognitive science using identity attributes like gender or race, the identity-level representations of people that are most easily replicated in post-training methods likely do not capture the reasons for the differences in people’s conceptual representations. Exploring corresponding notions of context-specific cognitive or computational resources (e.g., working memory) in models could allow them to serve as testbeds for theories of individual differences in humans (cf. Hu and Frank, 2024).

6 Conclusion

Here, we used a new approach to evaluate conceptual diversity in synthetically-generated LLM “populations”. Specifically, we adopted metrics from prior studies of conceptual diversity in humans that relate the variability in individuals’ underlying representations to population-level variability. We evaluate two popular methods for eliciting unique individuals from LLMs – temperature and prompting manipulations – in the word-color and word-concept domains and find that adding noise in the form of unrelated and randomly shuffled prompts does more to increase model diversity than persona prompting and temperature manipulations. Across two domains, no model reached human-level conceptual diversity. Further, our results suggest that alignment to synthetic or human preferences (in the form of RLAIIF and RLHF) flattens models’ conceptual diversity in these domains, compared to their base or instruction fine-tuned counterparts.

While alignment has now become a central part of LLM development, our findings urge caution that more work should be done to understand how it affects models’ internal diversity of “opinions” or conceptual representations. This is especially important given recent calls to use LLMs as stand-ins for human subjects in behavioral research. While it is clear that there could be gains under this paradigm, we should also ask: what might we be losing?

Limitations

We tested a limited set of models with a relatively small number of parameters (7B), due to computational resource constraints. In addition, we only tested two popular alignment methods, RLHF and RLAI. It remains unknown whether larger models, possibly aligned using different techniques or preference datasets, would exhibit the same patterns found in our experiments.

We only tested models in English-language domains, and we compared models to data collected from human participants based in the United States (Murthy et al., 2022) or collected by US-based researchers (Martí et al., 2023). While even these relatively homogeneous participants already exhibit individual-level variability, which LLMs fail to capture, an important consideration for future work is to examine conceptual representations across individuals with more diverse cultural and personal experiences.

Finally, we note that the models were tested in relatively simple domains (word-color associations, and similarity judgments between animals or politicians), which might not reflect the settings or applications in which models will be used.

Acknowledgments

We thank the members of the Harvard Computation, Cognition, and Development Lab, as well as Ilia Sucholutsky, for their helpful comments and discussion. We also thank Bala Desinghu for support with computing resources. TU was supported by the Jacobs Foundation. This material is based upon work supported by the NSF Graduate Research Fellowship under Grant No. DGE 2140743 to SKM, and by the Kempner Institute for the Study of Natural and Artificial Intelligence at Harvard University.

The code and data for our analyses is available at <https://github.com/skmur/onefish-twofish>.

References

Mostafa Abdou, Artur Kulmizev, Daniel Hershcovich, Stella Frank, Ellie Pavlick, and Anders Søgaard. 2021. *Can Language Models Encode Perceptual Structure Without Grounding? A Case Study in Color*. In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 109–132, Online. Association for Computational Linguistics.

Gati Aher, Rosa I. Arriaga, and Adam Tauman Kalai. 2023. Using large language models to simulate multiple humans and replicate human subject studies. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org. Place: Honolulu, Hawaii, USA.

Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2023. *Out of One, Many: Using Language Models to Simulate Human Samples*. *Political Analysis*, 31(3):337–351.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022a. *Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback*.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022b. *Constitutional AI: Harmlessness from AI Feedback*.

Antoine Bellemare-Pepin, François Lespinasse, Philipp Thölke, Yann Harel, Kory Mathewson, Jay A. Olson, Yoshua Bengio, and Karim Jerbi. 2024. *Divergent Creativity in Humans and Large Language Models*.

Marcel Binz and Eric Schulz. 2024. *Turning large language models into cognitive models*. In *Proceedings of the International Conference on Learning Representations*.

Danica Dillion, Niket Tandon, Yuling Gu, and Kurt Gray. 2023. *Can AI language models replace human participants?* *Trends in Cognitive Sciences*, 27(7):597–600.

Michael Franke, Polina Tsvilodub, and Fausto Carcassi. 2024. *Bayesian Statistical Modeling with Predictors from LLMs*.

Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak,

- Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Pier Giuseppe Sessa, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L. Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu-hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024. [Gemma: Open Models Based on Gemini Research and Technology](#).
- Matthias Gerstgrasser, Rylan Schaeffer, Apratim Dey, Rafael Rafailov, Tomasz Korbak, Henry Sleight, Rajashree Agrawal, John Hughes, Dhruv Bhandarkar Pai, Andrey Gromov, Dan Roberts, Diyi Yang, David L. Donoho, and Sanmi Koyejo. 2024. [Is Model Collapse Inevitable? Breaking the Curse of Recursion by Accumulating Real and Synthetic Data](#). In *First Conference on Language Modeling*.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [Chatgpt outperforms crowd workers for text-annotation tasks](#). *Proceedings of the National Academy of Sciences of the United States of America*, 120.
- Perttu Hämäläinen, Mikke Tavast, and Anton Kunnari. 2023. [Evaluating large language models in generating synthetic hci research data: a case study](#). In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA. Association for Computing Machinery.
- Joy He-Yueya, Wanjing Anya Ma, Kanishk Gandhi, Benjamin W. Domingue, Emma Brunskill, and Noah D. Goodman. 2024. [Psychometric Alignment: Capturing Human Knowledge Distributions via Language Models](#).
- Jennifer Hu and Michael Frank. 2024. [Auxiliary task demands mask the capabilities of smaller language models](#). In *First Conference on Language Modeling*.
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [Camels in a Changing Climate: Enhancing LM Adaptation with Tulu 2](#).
- janus. 2022. [Mysteries of mode collapse](#). *LESS-WRONG*.
- Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, Fanzhi Zeng, Kwan Yee Ng, Juntao Dai, Xuehai Pan, Aidan O’Gara, Yingshan Lei, Hua Xu, Brian Tse, Jie Fu, Stephen McAleer, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, Yike Guo, and Wen Gao. 2024. [AI Alignment: A Comprehensive Survey](#).
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7B](#). *arXiv preprint*.
- Benjamin S. Manning, Kehang Zhu, and John J. Horton. 2024. [Automated Social Science: Language Models as Scientist and Subjects](#).
- Raja Marjeh, Ilia Sucholutsky, Pol van Rijn, Nori Jacoby, and Thomas L. Griffiths. 2023. [Large language models predict human sensory judgments across six modalities](#). *Scientific Reports*, 14.
- Louis Martí, Francis Mollica, Steven T. Piantadosi, and Celeste Kidd. 2018. [Certainty is primarily determined by past performance during concept learning](#). *Open Mind : Discoveries in Cognitive Science*, 2:47–60.
- Louis Martí, Shengyi Wu, Steven T. Piantadosi, and Celeste Kidd. 2023. [Latent diversity in human concepts](#). *Open Mind : Discoveries in Cognitive Science*, 7:79–92.
- Sonia K. Murthy, Thomas L. Griffiths, and Robert D. Hawkins. 2022. [Shades of confusion: Lexical uncertainty modulates ad hoc coordination in an interactive communication task](#). *Cognition*, 225.
- Jakob Niedermann, Ilia Sucholutsky, Raja Marjeh, Elif Celen, Thomas L. Griffiths, Nori Jacoby, and Pol van Rijn. 2024. [Studying the effect of globalization on color perception using multilingual online recruitment and large language models](#).
- Laura O’Mahony, Leo Grinsztajn, Hailey Schoelkopf, and Stella Biderman. 2024. [Attributing Mode Collapse in the fine-tuning of Large Language Models](#).

- In *ICLR 2024 Workshop on Mathematical and Empirical Understanding of Foundation Models*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Vishakh Padmakumar and He He. 2024. [Does writing with language models reduce content diversity?](#) In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Peter S. Park, Philipp Schoenegger, and Chongyang Zhu. 2024. [Diminished diversity-of-thought in a standard large language model](#). *Behavior Research Methods*, 56(6):5754–5770.
- Roma Patel and Ellie Pavlick. 2022. [Mapping Language Models to Grounded Conceptual Spaces](#). In *International Conference on Learning Representations*.
- Max Peeperkorn, Tom Kouwenhoven, Dan Brown, and Anna Jordanous. 2024. [Is Temperature the Creativity Parameter of Large Language Models?](#) In *Proceedings of the 15th International Conference on Computational Creativity (ICCC'24)*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct Preference Optimization: Your Language Model is Secretly a Reward Model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Leonard Salewski, Stephan Alaniz, Isabel Rio-Torto, Eric Schulz, and Zeynep Akata. 2024. [In-context impersonation reveals large language models' strengths and biases](#). In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23, Red Hook, NY, USA*. Curran Associates Inc.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. [Whose opinions do language models reflect?](#) In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org. Place: Honolulu, Hawaii, USA.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal Policy Optimization Algorithms](#).
- Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. 2024. [AI models collapse when trained on recursively generated data](#). *Nature*, 631(8022):755–759.
- Guy Tevet and Jonathan Berant. 2021. [Evaluating the Evaluation of Diversity in Natural Language Generation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 326–346. Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#).
- Lewis Tunstall, Edward Emanuel Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro Von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2024. [Zephyr: Direct Distillation of LM Alignment](#). In *First Conference on Language Modeling*.
- Edward Vul and Harold Pashler. 2008. [Measuring the crowd within](#). *Psychological Science*, 19:645 – 647.
- Angelina Wang, Jamie Morgenstern, and John P. Dickerson. 2024a. [Large language models should not replace human participants because they can misportray and flatten identity groups](#).
- Guan Wang, Sijie Cheng, Xianyuan Zhan, Xiangang Li, Sen Song, and Yang Liu. 2024b. [OpenChat: Advancing Open-source Language Models with Mixed-Quality Data](#). In *The Twelfth International Conference on Learning Representations*.
- Pengda Wang, Huiqi Zou, Zihan Yan, Feng Guo, Tianjun Sun, Ziang Xiao, and Bo Zhang. 2024c. [Not Yet: Large Language Models Cannot Replace Human Respondents for Psychometric Research](#).
- Yue Yu, Yuchen Zhuang, Jieyu Zhang, Yu Meng, Alexander J Ratner, Ranjay Krishna, Jiaming Shen, and Chao Zhang. 2023. [Large Language Model as Attributed Training Data Generator: A Tale of Diversity and Bias](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 55734–55784. Curran Associates, Inc.

Tan Zhi-Xuan, Micah Carroll, Matija Franklin, and Hal Ashton. [Beyond preferences in AI alignment](#).

Banghua Zhu, Evan Frick, Tianhao Wu, Hanlin Zhu, Karthik Ganesan, Wei-Lin Chiang, Jian Zhang, and Jiantao Jiao. 2024. [Starling-7B: Improving Helpfulness and Harmlessness with RLAIIF](#). In *First Conference on Language Modeling*.

Banghua Zhu, Hiteshi Sharma, Felipe Vieira Frujeri, Shi Dong, Chenguang Zhu, Michael I. Jordan, and Jiantao Jiao. 2023. [Fine-Tuning Language Models with Advantage-Induced Policy Alignment](#).

A Details on constructing prompt manipulations

We derive the template and field values for the *persona* prompt condition from (Wang et al., 2024a) with additions made to the [race] field to cover all the United States census data categories and to the [occupation] field, to induce more diversity across our personas. We populate each field in the context by randomly selecting a value for each.

To create the *random* and *nonsense* prompt contexts, we first filtered English Wikipedia to remove articles about people (to avoid reduplicating persona prompting) as well as disambiguation pages. We obtained the random contexts by taking the first sentence of a randomly sampled article from the filtered dataset. We obtained the nonsense contexts by shuffling the words within each, and re-applying sentence casing to each resulting bag of words.

B Percentage of invalid responses

Figure 5 shows the percentage of total queried responses that constituted invalid responses for each model in each domain. We omit the results of models that have an average of 70% or more invalid responses across all prompt and temperature manipulations: Llama and Tulu on the color task, and Llama on the concept task. In general, we find that adding noise via high temperature and unrelated prompt contexts increases incoherence and thus invalidity, not meaningful diversity.

C Word-color association domain

C.1 Population vs internal ΔE

Figures 6 and 7 show the population ΔE versus internal ΔE for prompt- and temperature-based manipulations, respectively. In all settings, the points are near the line of unity, suggesting low population heterogeneity.

C.2 Comparing model and human color associations

To validate models’ ability to perform word-color associations, we report results comparing the similarity between models’ and humans’ word-color associations. Figure 8 shows the Jensen-Shannon divergence between models’ and humans’ color response distributions, averaged over all words.

C.3 Visualizing color associations for a variety of word types

Figure 9 shows models’ and humans’ color associations for a variety of word types. These are well-aligned in some cases, like “tomato”, which is strongly associated with red. Interestingly, for more abstract words like “optimism” or “fame” (not pictured), humans have heterogeneous color distributions, whereas models’ responses are less diverse but still interpretable – for example, “fame” is associated with gold and “jealousy” with green. For the word “skin”, we see humans’ color associations feature the lighter skin tones that one might expect from a Western population. In contrast, some of the models (e.g. Openchat, Starling, and persona-prompted Gemma-Instruct) actually do a better job of diversifying these associations. This might suggest that models’ word-color associations largely rely on distributional semantics (e.g., common co-occurrences such as “green with jealousy”), whereas humans might have other associations with abstract concepts that aren’t commonly expressed in text corpora.

D Conceptual similarity judgements domain

D.0.1 Between-subjects reliability

As an additional measure of population heterogeneity in the conceptual similarity domain, we calculate the between-subjects reliability for the human baseline and model data. Figure 10 shows the probability that two random subjects will agree on the similarity judgements for a given concept.




Figure 5: Number of invalid responses for the word-color associations task (top) and conceptual similarity judgements task (bottom) for each tested combination of prompting condition and temperature manipulation. We omit the results of Llama and Tulu on the color task, and for Llama on the concept task given the low average valid responses across prompt and temperature manipulations. In general, we find that adding noise via high temperature and unrelated prompt contexts increases incoherence, not meaningful diversity.

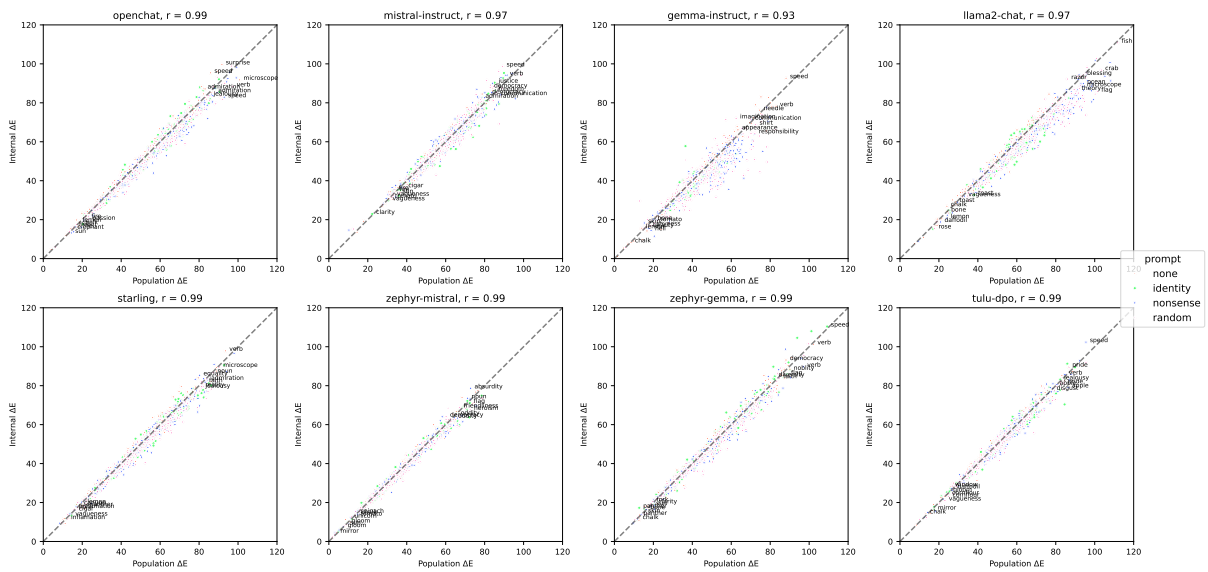


Figure 6: Population vs. Internal ΔE for prompting manipulations.

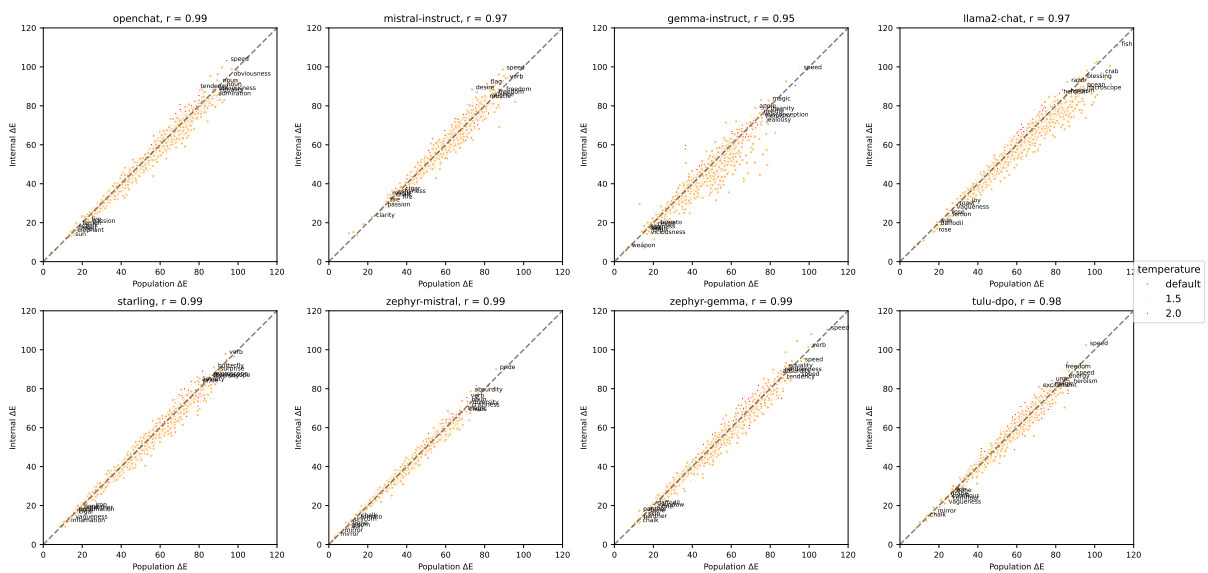


Figure 7: Population vs. Internal ΔE for temperature manipulations.




Figure 8: Jensen-Shannon divergence between models' color response distributions and humans' for each word.




Figure 9: Models' and humans' color associations for a variety of word types. For words where there is high human variation (e.g. abstract words like "optimism") the models' responses appear to collapse to just handful of colors, but these colors are easily interpretable (e.g. *jealousy* ≈ green)




Figure 10: Between-subjects reliability for model data and human baseline in the conceptual similarity domain.