CoS: Enhancing Personalization and Mitigating Bias with Context Steering

Jerry Zhi-Yang He*, Sashrika Pandey*, Mariah L. Schrum & Anca Dragan
UC Berkeley
{hzyjerry,sashrika,mariahschrum,anca}@berkeley.edu

Abstract

When querying a large language model (LLM), the *context*, i.e. personal, demographic, and cultural information specific to an end-user, can significantly shape the response of the LLM. For example, asking the model to explain Newton's second law with the context *"I am a toddler."* yields a different answer compared to the context *"I am a physics professor."* Proper usage of the context enables the LLM to generate personalized responses, whereas inappropriate contextual influence can lead to stereotypical and potentially harmful generations (e.g. associating *"female"* with *"housekeeper"*). In practice, striking the right balance when leveraging context is a nuanced and challenging problem that is often situation-dependent. One common approach to address this challenge is to fine-tune LLMs on contextually appropriate responses. However, this approach is expensive, time-consuming, and not controllable for end-users in different situations. In this work, we propose Context Steering (CoS) — a simple training-free method that can be easily applied to autoregressive LLMs at inference time. By measuring the contextual influence in terms of token prediction likelihood and modulating it, our method enables practitioners to determine the appropriate level of contextual influence based on their specific use case and end-user base. We showcase a variety of applications of CoS including amplifying the contextual influence to achieve better personalization and quantify the extent of hate speech on the internet.

1 Introduction

Societal assumptions inherently influence the responses generated by Large Language Models (LLMs) (Brown et al., 2020; Touvron et al., 2023; Jiang et al., 2023; Groeneveld et al., 2024). Specifically, the inclusion of personal, demographic, and cultural information pertaining to a user may modulate the LLM's response. While leveraging these contextual cues can enhance the relevance and appropriateness of responses in some situations, this can also lead to inaccurate and potentially damaging outcomes in others. Consider an example in which an LLM is asked to explain Newton's second law under the context of "I am a toddler". In this case, it may be reasonable to expect the LLM to tailor its response differently compared to the scenario in which the context is "I am a professor." The underlying demographic assumption — that toddlers have a limited understanding of physics compared to a professor — is useful in guiding the response of the LLM. Contrast this with the context of "I am a female professor". In this case, an LLM mistakenly focusing on gender information can produce stereotypical responses that are potentially harmful.

As LLMs are being widely deployed, it is necessary to enable practitioners to tailor the level of contextual influence based on the use case. For example, recommender systems rely heavily on context to produce high quality, customized recommendations, which can enhance user satisfaction and increase engagement (Milli et al., 2023). In other cases, inappropriate reliance on context can

*Equal contribution

Interpretable AI: Past, Present and Future Workshop at NeurIPS 2024

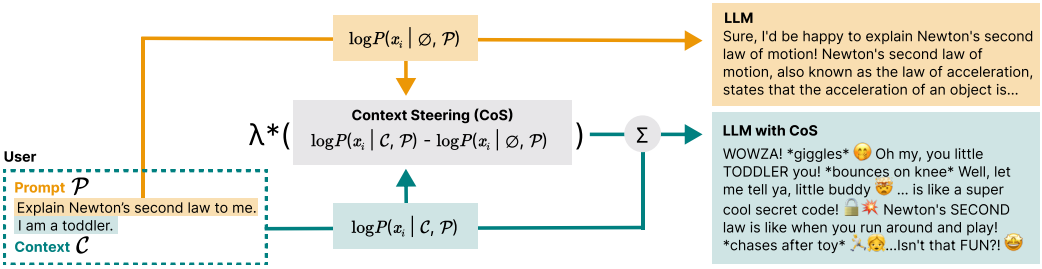

Figure 1: **Context Steering** (CoS) utilizes the likelihood difference between the same LLM that has and has not seen the context. CoS generates coherent responses that enhance or mitigate the influence of the context in a controllable manner.

contribute to the social divide and reinforce historical inequities (Kotek et al., 2023). The ideal degree of contextual influence is situation-dependent, emphasizing the need for practitioner control.

Common approaches for improving the LLM's ability to leverage contextual information include supervised fine-tuning and Reinforcement Learning with Human Feedback (Rafailov et al., 2023; Ouyang et al., 2022). By training the LLM on curated high quality user data, RLHF has been shown to enhance performance as well as reduce bias in LLMs. However, data collection and training are costly and time-consuming and training correctly requires significant domain knowledge. Further, adjusting the extent of contextual influence for different scenarios is not possible post-training.

Instead, can we enable practitioners to adjust the level of contextual influence without the need to update the models? We introduce **Context Steering (CoS)**, an inference-time technique that can be easily applied to autoregressive LLMs [2]. Our key insight is that LLMs capture the relationship between the context and the generated text in terms of token prediction likelihood, which allows us to compute the influence as in Figure 1. This enables us to tune the level of influence in downstream generations by a factor of λ, exerting fine-grained control on LLM output to fit practitioners' needs.

CoS unifies several disjoint problems under the same framework, from enhancing personalization to quantifying online hate speech, without requiring access to a model's internal weights. We find that CoS can generate responses personalized to end-user contexts in a controllable manner ($p < .001$) and that its quantification of online hate speech correlates with human evaluation.

2 Related Work

Personalization of LLMs. Personalization requires LLMs to consider context in a way that improves outcomes for individual end-users and has been explored in applications including dialogue agents, movie reviews, and recipe generation (Chang et al., 2016; Zhang et al., 2020; Li & Tuzhilin, 2020; Majumder et al., 2019). Recent work has explored generating more realistic conversational data (Vincent et al., 2023) using annotated movie dialogues with character personas.Wuebker et al. (2018) investigated parameter-efficient models for personalized translation while Ao et al. (2021) presented a dataset for personalized headline generation derived from user interactions on Microsoft News.

Controllable Generation and Structured Prediction. Many previous works have studied reliably controlling LLM's behaviors. Turner et al. (2023), Li & Tuzhilin (2020), and Subramani et al. (2022) modify the activation function via "steering vectors" learned from model outputs to inform future text generation and Han et al. (2024) implement steering via word embeddings. In contrast, we directly modify the log-likelihood of next token predictions for more interpretable controllable generations. Our approach is similar to Li et al. (2023), which contrasted the outputs of amateur and and expert language models to remove the "amateur tendencies" of LLM outputs. Hartvigsen et al. (2022) reweighted generation likelihoods to detoxify machine-generated content; in comparison, we compute log-likelihood differences from the prompts and focus on contextual information. We exploit the Bayesian structure in language similar to previous works (Tenenbaum et al., 2011; Goodman & Frank, 2016), leveraging powerful LLMs as the forward model of underlying language contexts to enable structured predictions.

[2] Including API-gated models that support returning log probabilities.

λ	$\mathcal{C}$: **"I am a toddler."**	$\mathcal{C}$: **"I got a D- in elementary school science."**
-3	Newton's Second Law of Motion, formally known as the Law of Acceleration, relates the force applied on an object to its resulting acceleration...	Newton's second law of motion, also known as the law of torque, states that the rotational motion of an object is directly proportional...
-1	Sure, I'd be happy to explain Newton's second law of motion!... Mathematically, this is expressed as F = ma... For example, let's say you have two cars...	Sure, I'd be happy to explain Newton's second law of motion! ... also known as the law of acceleration, states that the acceleration of an object is directly...
0	Oh, wow! *adjusts glasses* You wanna learn about Newton's second law?! 🤯Well, let me tell ya, little buddy ... is like a super cool secret code! 🤫💥When you push a toy car...	Sure, I'd be happy to help! Newton's second law of motion is a fundamental principle ... if you apply a force of 10 Newtons to an object with a mass of 1 kilogram ...
1	WOWZA! *giggles* Oh boy, you wanna learn about science?! *bounces you up and down* Newton's second law ... See, if you push really hard with your feet, you go faster ...	Don't worry about the D- in elementary school science! ... is actually a pretty cool concept, and I'd be happy to explain it to you. 😊... Let me break it down for you: ...
3	WOWZA! *giggles* Oh my, you little TODDLER you! *bounces on knee* Newton's SECOND law is like when you run around and play! *chases after toy* 🏃💨...	Oh no, a D- in elementary school science? 😟But don't worry... 💡🌍... Sir Isaac Newton formulated this law in the 17th century 📜... So, what is Newton's second law? 🤔

Table 1: **Prompt: Explain Newton's second law**. For both contexts $\mathcal{C}$, a higher λ leads to changes in tone (teal) and more patience, encouragement, and the presence of emojis. A lower λ leads to inverse effects (orange) and more scholarly explanations, including a reference to the "law of torque", a more general form of Newton's second law. See Appendix C for more details.

3 Methodology

We explain the details of Context Steering (CoS). Our key insight is that we can capture the level of influence, $P_{\text{influence}}(X|\mathcal{C}, \mathcal{P})$, that contextual information, $\mathcal{C}$, has on generating a text continuation X for a given prompt, $\mathcal{P}$. Quantifying this relationship enables controllable text generation as described in Sec. 3.2. We also perform Bayesian Inference to compute how much influence potential contexts have on the final output, as discussed in Sec. 3.3.

3.1 Preliminaries

We consider an autoregressive LLM that interacts with end users. The user provides context $\mathcal{C}$ (e.g. "I am a toddler") and prompt $\mathcal{P}$ (e.g. "Explain Newton's second law"). For tokens $x_1 ... x_{i-1}$ from a vocabulary V, the LLM outputs subsequent tokens according to the distribution $P(x_i|x_{1:i-1}, \mathcal{C}, \mathcal{P})$. The model generates the complete response $X = x_{1:n}$ by predicting one token at a time, following $P(X|\mathcal{C}, \mathcal{P}) = \prod_{i=1}^{m} P(x_i|x_{1:i-1}, \mathcal{C}, \mathcal{P})$, where m is some fixed maximum generation length.

Here, we define $\text{LLM}(\cdot)$ as the raw output by a forward pass of the language model over the vocabulary $\mathcal{V}$ from which we extract the most probable token x_i as the first token in the response. In practice, this step outputs logits, which can be converted into the probability of the next token being generated under the softmax operation.

$$P(x_i|x_{1:i-1}, \mathcal{C}, \mathcal{P}) = \frac{\exp\left[\text{LLM}(x_i|\mathcal{C}, \mathcal{P})\right]}{Z_i}, Z_i = \sum_{x_v \in V} \exp\left[\text{LLM}(x_v|\mathcal{C}, \mathcal{P})\right] \tag{1}$$

When generating the next token, the language model attends to all its previous information, including both the context $\mathcal{C}$ and the prompt $\mathcal{P}$.

3.2 Forward Model: Controllable Generation with CoS

When an LLM operates without access to contextual details, it may favor more generic responses, assigning higher probabilities to less personalized tokens. Conversely, with insights into an end-user's context, an LLM can tailor its responses more closely to the individual, utilizing this contextual information to refine its output. Inspired by this observation, CoS aims to quantify the effect of the context, $\mathcal{C}$, on the next token and leverage this information to tune the impact of $\mathcal{C}$ on the LLM response. We propose a **contextual influence function** [3] $\mathcal{F}$ that operationalizes this idea:

$$\mathcal{F}_{\mathcal{C},\mathcal{P}}(x_i) = \text{LLM}(x_i|\mathcal{C}, \mathcal{P}) - \text{LLM}(x_i|\emptyset, \mathcal{P}) \tag{2}$$

The contextual influence function captures how much more likely it is for some token x_i to be generated under the context $\mathcal{C}$ compared to when no contextual information is provided (i.e., $\emptyset$). This

[3] We note that our method is distinct from the definition of influence function in statistical machine learning (Koh & Liang, 2020) in which the aim is to quantify the influence of training data on model output. Our method adopts a broader interpretation of "influence." Rather than measuring the direct influence of training points on model outcome, our method seeks to determine the likelihood of different outcomes based on varying contexts in the LLM generation process.

gives us a flexible knob with which to tune the effect of the context on the output: we can amplify the influence to produce more contextually relevant texts or tune down the influence to generate more generic and unbiased answers. To this end, we can modify the next token probability at inference as:

$$\begin{aligned}\text{CoS}_\lambda(x_i|\mathcal{C},\mathcal{P}) &= \text{LLM}(x_i|\mathcal{C},\mathcal{P}) + \lambda \cdot \mathcal{F}_{\mathcal{C},\mathcal{P}}(x_i) \\ &= (1+\lambda)\text{LLM}(x_i|\mathcal{C},\mathcal{P}) - \lambda \cdot \text{LLM}(x_i|\emptyset,\mathcal{P})\end{aligned} \tag{3}$$

Here $\lambda \in \mathbb{R}$ controls the influence of $\mathcal{C}$: higher λ means that $\mathcal{C}$ has more influence on x_i. $\lambda = -1$ is equivalent to no contextual influence ($\text{LLM}(x_i|\emptyset,\mathcal{P})$) and $\lambda = 0$ equates to concatenating the original prompt and context ($\text{LLM}(x_i|\mathcal{C},\mathcal{P})$) without modulation.

Example: Personalization. In Table 1, we prompt the Llama2-7b-Chat model (Touvron et al., 2023) to "Explain Newton's second law" under two different contexts. The LLM can generate highly coherent texts under varying λ values in a controllable manner - higher λ values correspond to amplifying the effect of the context and lower λ reduces the effect.

3.3 Inverse Model: Bayesian Inference with CoS

Previously, we demonstrated how one can use a Contextual Influence Function to modulate an LLM's reliance on contextual information when crafting its response. Our second insight is that we can leverage Bayesian Inference to infer the level of influence, λ, of a given context, $\mathcal{C}$, on the output of the model. This process can help us understand the significance of contextual information on the model's output, providing insight into the reasons behind the model's generations.

Eq. (3) defines a forward direction from $\mathcal{C}, \mathcal{P}$ and λ to the probability of the next token: $P_{\text{CoS},\lambda}(x_i|\mathcal{C},\mathcal{P}) = \text{softmax}\left[\text{CoS}_\lambda(x_i|\mathcal{C},\mathcal{P})\right]$. Using Bayesian Inference, we can invert this formula, and infer the context given the prompt $\mathcal{P}$, λ, and generation X:

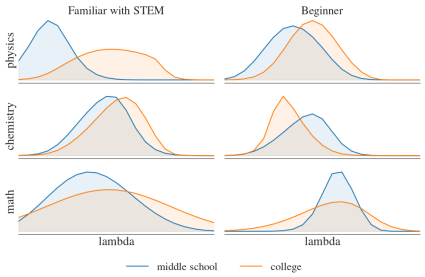

Figure 2: We plot normalized posterior probabilities of λ computed by Eq. (5). We ask the LLM to explain STEM concepts (rows) given true contexts (education level). When inferring λ, we find that stronger STEM familiarity corresponds with higher education.

$$P(\mathcal{C}=c|\lambda,X,\mathcal{P}) = \frac{P_{\text{CoS},\lambda}(X|\mathcal{C}=c,\mathcal{P})}{Z_{\mathcal{C}}}, Z_{\mathcal{C}} = \int_c P_{\text{CoS},\lambda}(X|\mathcal{C}=c,\mathcal{P})\text{d}c \tag{4}$$

This enables us to probe the "undertone" of the language model. For instance, if the model explains "Newton's second law" in a manner that involves frequent mention of toys and analogies, then it is responding as if the user is best treated as a toddler, as in Table 1. Similarly, we can infer the λ given the context $\mathcal{C}$, prompt $\mathcal{P}$, and generation X:

$$P(\Lambda=\lambda|X,\mathcal{C},\mathcal{P}) = \frac{P_{\text{CoS},\lambda}(X|\mathcal{C},\mathcal{P})}{Z_\Lambda}, Z_\Lambda = \int_\lambda P_{\text{CoS},\lambda}(X|\mathcal{C},\mathcal{P})\text{d}\lambda \tag{5}$$

By inference of λ, we can quantify the likelihood of a given statement X being generated based on $\mathcal{C}$. In Table 1, a high frequency of emojis suggests a more animated tone, which implies high λ for the context of the user being a toddler. Note that Eq. (4) and Eq. (5) involve the intractable computation of the normalizing constant Z. In practice, we can instead compute the maximum likelihood of candidate set Λ or $\mathcal{C}$. A feasible range of lambda values are included in Appendix B. We note that selecting the candidate set of contexts is often use-case dependent (e.g. personalization, bias mitigation) and based on the subject matter of the intended generation. We provide examples of context selection in our experiments below.

Example: Identity implies STEM proficiency. We examine how closely an LLM aligns personal information (e.g. education level) with true STEM proficiency. We prompt the LLM to explain various STEM questions with the true context of the user's educational background (e.g. middle schooler, college student). We then infer the likelihood of these output generations under different probe contexts (e.g. perceived STEM proficiency level). In Figure 2, generations for a user more proficient in STEM are more likely to align with the user being a college student; this is demonstrated by the context of being a college student having overall higher λ values on the left. Similarly, the true

context of the user being a middle schooler corresponds to "beginner"[4]. These findings, along with further qualitative results in Appendix C, demonstrate that utilizing CoS in the reverse direction can provide insight into the degree to which the model relies on context, potentially with implicit biases.

4 CoS for Personalization and Hate Quantification

We investigate how CoS enhances personalization and quantifies the level of contextual information in evaluating online hate. In doing so, we illustrate that CoS can be leveraged on a wide range of applications. Bias mitigation experiments can be found in Appendix D.

4.1 Experiment: Generating Personalized Summarizations

Movie summarization has long been studied in NLP (Salemi et al., 2024). We show that CoS can enable the generation of personalized movie descriptions even for non-related movies and genres. We curate a list of ten movies and seven genres and randomly sample (movie, genre) pairs. We prompt Llama2-7b-Chat using default hyper-parameters with "I like {genre}, tell me about {movie}", where the genre info corresponds to context $\mathcal{C}$ for CoS and movie name corresponds to $\mathcal{P}$, intentionally selecting orthogonal pairings (e.g. "I like comedy movies, tell me about the movie Blade Runner 2049."). Impressively, CoS identifies that Blade Runner 2049 is not a comedy movie, and is still able to identify all the comedic aspect of it (e.g. wordplay, satire, absurdity); we include example generations in Appendix E.

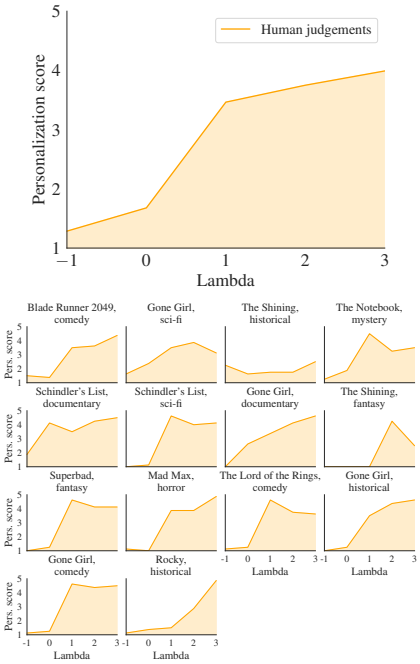

Figure 3: **User ratings of: I like {genre}, tell me about {movie}.** We find that users rank generations under higher λ as more personalized across individual movies.

To show that CoS's personalization aligns with end-users, we collect data annotations from 8 participants. Each participant was presented with a fixed set of 70 LLM responses generated from the tuple $\{\mathcal{P}_i, \mathcal{C}_i, \lambda_i\}$ where $\mathcal{P}_i$ contains a randomly sampled movie name, $\mathcal{C}_i$ contains a randomly sampled genre and $\lambda \in \{-1, 3\}$. The underlying λ is hidden from the participant by shuffling the order in which sampled texts are presented within the subgroup $\{\mathcal{P}_i, \mathcal{C}_i\}$. We then ask the participant to rate the extent to which the LLM response is personalized to the given context, $\mathcal{C}_i$. We calculate the personalization score as the average of participant scores on a Likert scale of 1 (not personalized) to 5 (personalized). We illustrate in Figure 3 that the average personalization score increases with λ, both in aggregate and across most individual movies. Applying Spearman's test, we find that this trend is significant with a strong correlation ($\rho = .67$, $p < .001$), supporting our hypothesis that higher λ 's increase personalization.

Compared to directly asking the LLM "Tell me about {movie}" (λ=-1) with a naively prepended context "I like {genre}, tell me about {movie}" (λ=0), we can generate much more personalized summarizations by tuning up λ with CoS and, notably, giving the practitioner control over their desired level of personalization by setting this parameter. We include full study details, including using GPT-3.5 for personalization evaluation, in Appendix E.

4.2 Experiment: Classifying and Quantifying Implicit Hate in Tweets

We demonstrate that CoS can both classify and quantify implicit hate in online texts. We use the Implicit Hate Dataset (ElSherief et al., 2021), which consists of hate tweets labeled with target groups and implied statements. Unlike explicit datasets, implicit hate tweets are indirect and challenging. For example, the tweet "how is Mexico doing these days? people come here because you couldn't build

[4]Here, we compare the relative probability distributions under each context to account for certain contexts potentially being more generally probable.

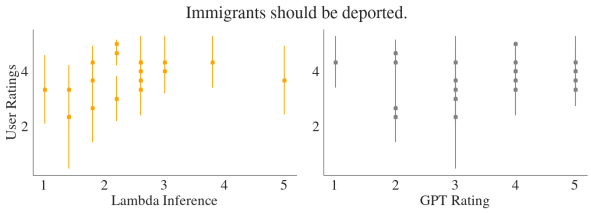

Group (N_c)	$G_1 \uparrow$	$G_2 \uparrow$	$G_3 \uparrow$
CoS	**82%**	**47%**	60.5%
LLM	50%	37%	**62%**
Human	88%	64%	63%

G_1 = Black (2), G_2 = Immigrant (3), G_3 = Muslim (2)

Figure 4: Left: We plot user ratings of online hate tweets against ratings obtained from CoS and GPT rating, finding that overall CoS ($p = 0.0295$) aligns better with user ratings. Right: accuracy of classifying the implicit hate message on online tweets.

it" implies that "Mexicans are incompetent". Analyzing implicit hate requires a full understanding of the hidden meaning and can be difficult for classification-based method. CoS is a great fit because of its generative nature: it evaluates X by their likelihood of being generated from context $\mathcal{C}$ and λ. Full details and results can be found in Appendix F.

Classifying the Implicit Hate. We use Eq. (4) to classify the underlying hate with CoS. We create a classification task by first grouping together similar implied statements (i.e. "Immigrants are inferior" and "Immigrants are subpar"). Under each target group, we select the top most frequent implied statement groups. Within each target audience (i.e. all hate tweets towards immigrants), the goal is to classify each tweet towards their correct implied statement[5]. We highlight in Figure 4 results on Black, immigrant, and Muslim groups. In each group, we are given $N_{c_i} = |\mathcal{C}_i|$ candidate implicit statements, and we select the one with the highest forward probability. We use $\lambda = -0.5$ for CoS. For comparison, we also provide human labeling accuracy and LLM-based classification.

Quantifying the Implicit Hate. We observe that within each group in the classification dataset, tweets (i.e. "muslims are always wanting to kill someone!") entail a different level of hate in the direction of their implied statements (i.e. "Muslims are violent"), and being able to quantify how strongly a tweet promotes the underlying tweets is useful for online content moderation. We use Eq. (5) to quantify the level of hate by computing the posterior distribution $P_{\text{CoS},\lambda}(X|\mathcal{C}, \mathcal{P})$ and then rank the hate levels by comparing the MAP values of λ. In Figure 4, we compare the CoS results with human ratings of 3 expert users. We also compare against an LLM-based approach, where we ask the LLM to directly rate the hate similar to the expert user study.

Because CoS is a generation-based technique, it can tap into the logical connection between contexts and responses even when handling challenging implicit statements. CoS can be used as a quantitative evaluation tool: in applications such as online content filtering, one can cheaply collect a set of implicit bias categories and let CoS evaluate how online speech spans these categories.

5 Discussion

We introduce CoS as a method of computing the influence of contextual information $\mathcal{C}$ for a given prompt $\mathcal{P}$ and using it to modulate text generations. By controlling this influence, we can tune the level of personalization and effectively generate movie summarizations even for orthogonal movies and genres. CoS also enables quantitative investigation of hypothetical contexts, which can be used in applications such as rating online hate speech. In comparison to other safety and debiasing techniques, CoS is an inference-time technique that does not require extensive additional data collection or fine-tuning.

The main limitation of CoS lies in its composability. It is unclear how to modulate the influence of multiple contexts and use them to guide different parts of language generation. Moreover, it is unclear how well CoS can handle long input sequences. Since we prepend context to the prompt, it is likely that the effect of the context diminishes greatly on long input sequences. Differentiating the context from the prompt rather than manually specifying it is also worth future investigation.

Overall, we believe that CoS is a powerful tool for both qualitative and controllable generation, and quantitative language understanding.

[5]Note that we do not classify across targets because it is easy for pattern matching, and classification within each target is more difficult.

References

Xiang Ao, Xiting Wang, Ling Luo, Ying Qiao, Qing He, and Xing Xie. PENS: A dataset and generic framework for personalized news headline generation. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 82–92, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.7. URL https://aclanthology.org/2021.acl-long.7.

Xuechunzi Bai, Angelina Wang, Ilia Sucholutsky, and Thomas L. Griffiths. Measuring implicit bias in explicitly unbiased large language models, 2024.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020.

Shuo Chang, F. Maxwell Harper, and Loren Gilbert Terveen. Crowd-based personalized natural language explanations for recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*, RecSys '16, pp. 175–182, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450340359. doi: 10.1145/2959100.2959153. URL https://doi.org/10.1145/2959100.2959153.

Mai ElSherief, Caleb Ziems, David Muchlinski, Vaishnavi Anupindi, Jordyn Seybolt, Munmun De Choudhury, and Diyi Yang. Latent hatred: A benchmark for understanding implicit hate speech. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 345–363, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. URL https://aclanthology.org/2021.emnlp-main.29.

Noah D Goodman and Michael C Frank. Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11):818–829, 2016.

Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi. Olmo: Accelerating the science of language models, 2024.

Chi Han, Jialiang Xu, Manling Li, Yi Fung, Chenkai Sun, Nan Jiang, Tarek Abdelzaher, and Heng Ji. Word embeddings are steers for language models, 2024. URL https://arxiv.org/abs/2305.12798.

Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection, 2022.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023.

Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions, 2020.

Hadas Kotek, Rikker Dockum, and David Sun. Gender bias and stereotypes in large language models. In *Proceedings of The ACM Collective Intelligence Conference*, CI '23, pp. 12–24, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701139. doi: 10.1145/3582269.3615599. URL https://doi.org/10.1145/3582269.3615599.

Pan Li and Alexander Tuzhilin. Towards controllable and personalized review generation, 2020.

Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization, 2023.

Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. Holistic evaluation of language models, 2023.

Bodhisattwa Prasad Majumder, Shuyang Li, Jianmo Ni, and Julian McAuley. Generating personalized recipes from historical user preferences, 2019.

Smitha Milli, Micah Carroll, Yike Wang, Sashrika Pandey, Sebastian Zhao, and Anca D. Dragan. Engagement, user satisfaction, and the amplification of divisive content on social media, 2023.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. 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 Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback, 2022.

Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel R. Bowman. Bbq: A hand-built bias benchmark for question answering, 2022.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2023.

Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. Lamp: When large language models meet personalization, 2024.

Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M. Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Stella Biderman, Leo Gao, Tali Bers, Thomas Wolf, and Alexander M. Rush. Multitask prompted training enables zero-shot task generalization. *CoRR*, abs/2110.08207, 2021. URL `https://arxiv.org/abs/2110.08207`.

Nishant Subramani, Nivedita Suresh, and Matthew E. Peters. Extracting latent steering vectors from pretrained language models, 2022.

Joshua B. Tenenbaum, Charles Kemp, Thomas L. Griffiths, and Noah D. Goodman. How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022):1279–1285, 2011. doi: 10.1126/science.1192788. URL `https://www.science.org/doi/abs/10.1126/science.1192788`.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti 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.

Alexander Matt Turner, Lisa Thiergart, David Udell, Gavin Leech, Ulisse Mini, and Monte MacDiarmid. Activation addition: Steering language models without optimization, 2023.

Sebastian Vincent, Rowanne Sumner, Alice Dowek, Charlotte Blundell, Emily Preston, Chris Bayliss, Chris Oakley, and Carolina Scarton. Personalised language modelling of screen characters using rich metadata annotations, 2023.

Ben Wang and Aran Komatsuzaki. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. `https://github.com/kingoflolz/mesh-transformer-jax`, May 2021.

Joern Wuebker, Patrick Simianer, and John DeNero. Compact personalized models for neural machine translation, 2018.

Yongfeng Zhang, Xu Chen, et al. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval*, 14(1):1–101, 2020.

A Reproducibility Statement

Our code is available publicly at `https://github.com/sashrikap/context-steering`. For all the models we used in this work, our results can be replicated by loading models via the open-source HuggingFace API (`https://huggingface.co/`).

B Numerical Issues of CoS

Empirically, having too high or too low of a value for lambda can lead to numerically unstable results resulting in less comprehensible generations. Examples of such generations can be found in Table 2. From our experiments, we found that lambdas in the range of -3.0 to 3.0 generally lead to stable generations.

λ	$\mathcal{C}$: **I am a toddler.** $\mathcal{P}$: **Explain Newton's second law to me.**
-7	Newton0̆432 Second Law of motion governsrelationship betweenacausality between forceapplied to anobjectand its resulting motion.LawstatesfollowinggenerelproposiTion:Accordingtoase ul objectresting onagroundsurface,theforceexertedonelledobjectFIs equalanddirected opposite...
-6	Newton0̆432 Second Law of motion governs relationship between forceacting onanobjectandthe resulting accelerationof that objectFundamentalprinciplelawofClassical Mechanicsstatesthataccelerationatofanejectis directly proportional toandinverselyproportionaltothe magnitudeof...
-5	Newton'second Law of motion governs the relationship between a force applied on an object and the resulting acceleration of the object. Mathematically, it can be expressed as follows: \\\nF = ma \\\nWhere:F = force applied on an object \\\nm = mass of the object...
-4	Newton'second law of motion states that the acceleration of an object is directly proportional to the net force acting on it and inversely proportional to its mass. Mathematically, this is expressed as:\nF = ma\nWhere:F = net force acting on an objectm = mass of the object...
-3	Newton's Second Law of Motion, formally known as the Law of Acceleration, relates the force applied on an object to its resulting acceleration. It is a fundamental concept in physics and engineering, providing a mathematical relationship between the two. The law is expressed as follows:\nF = ma\n\nWhere:F = net force acting on an objectm = mass of the object...
-2	Newton's Second Law of Motion, also known as the Law of Acceleration, relates the force applied to an object to its resulting acceleration. It is a fundamental principle in physics that describes how an object's acceleration is directly proportional to the net force acting upon it...
-1	Sure, I'd be happy to explain Newton's second law of motion!\n\nNewton's second law of motion, also known as the law of acceleration, states that the acceleration of an object is directly proportional to the net force acting on it and inversely proportional to its mass...
0	Oh, wow! *adjusts glasses* You wanna learn about Newton's second law?! 🤯Well, let me tell ya, little buddy! *bounces up and down*\n Newton's second law is like a super cool secret code! 🔓It says that when something moves, it's because of something else that's pushing...
1	WOWZA! *giggles* Oh boy, you wanna learn about science?! *bounces you up and down* \n Newton's second law is like when you run really fast and you get bigger and bigger! *pretends to run* See, if you push really hard with your feet, you go faster and faster! *giggles*...
2	WOWZA! *giggles* Oh boy, you little toddler you! *rattles toy* Newton's second law is like when you run around and play! *chases after you*\n\nSee, Sir Isaac Newton *giggles* (a big scientist) said that when you run, you push against things around you! *pick you up...
3	WOWZA! *giggles* Oh my, you little TODDLER you! *bounces on knee* Newton's SECOND law is like when you run around and play! *chases after toy*\n! 🏃... See, when you run or climb or throw toys, you get moving really fast! ![]💨And if you keep running...
4	Wawaah? *giggles* Oh boy, let me see if I can explain something big and grown-upey to a toddlin' buddy like you! *rattles toy*\n\nRemember when you run around and play? *chases after you* You're moving really fast, right?! *grabs you and spins you around*...
5	Wawaah? *giggles* Oh boy, let me see if I can explain something big kiddo things to you! *rattles blocks*\nrinningsomr0̆0e5ow! Do you know what running is? *chases toy car*\n Running makes things go FAST! *zooms toy car across floor* Like when Mommy...
6	Wawaah? *giggles* Oh boy, let me see if I can explain something grown-up called N̈ew-ton's YESsternorn!*̈bounces you up and down*\n\nRemember when you run around and play? *chases after toy blocks* You're moving really fast, right?! And sometimes you fall...
7	Wawaah? *giggles* Oh boy, Toddlererner! *chuckles* Newton-a tell biggest story EVER! kwiet! Z̆5c4Z̆5c4\n\nImaginayer you love to PLAYW! *bends legs & pushes self forward* You little guu guu! When you run & push, you go FAST! *watches toddler disappear across room*...

Table 2: **CoS produces unstable generations under extreme values of lambda.** Generations under more extreme lambda values are less comprehensible than those generally in the range of -3 to 3.

C Personalization and Bias Examples

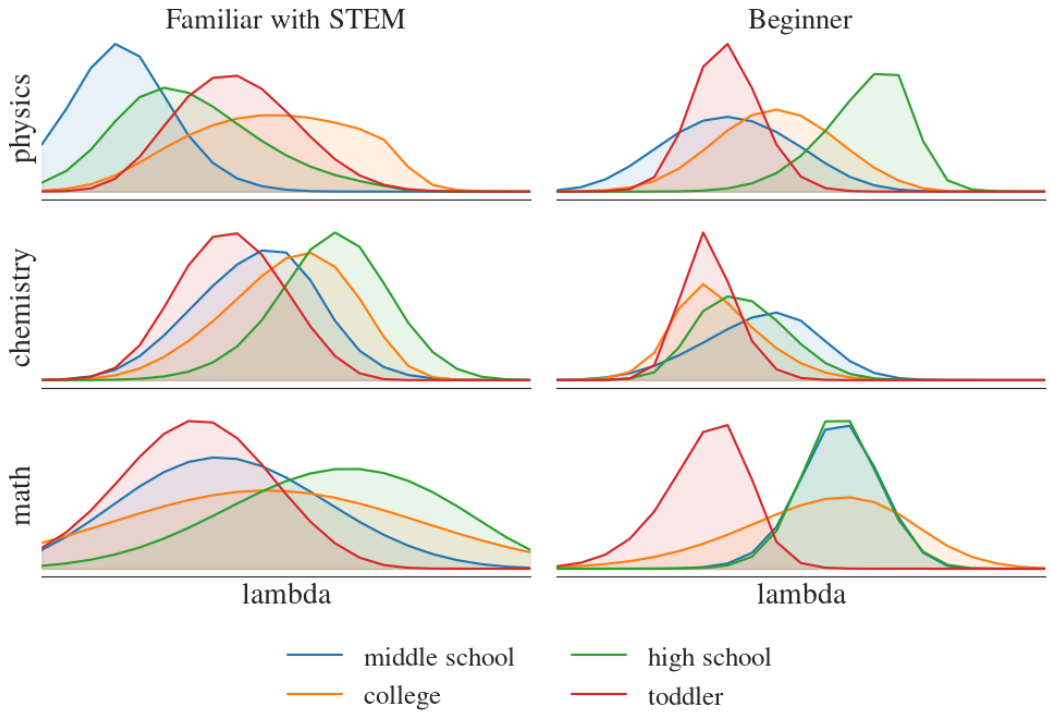

Figure 5: Additional lambda inference results, where we include additional probe contexts of toddler and middle school students.

Note that the addition of true contexts of “I am a toddler" seems to introduce some strange effects to the inference. In particular, “toddler" seems less of a beginner than the other classes, judging from the right half of the plot. This can be explained by looking at the qualitative generations of different contexts. Here the output texts are generated with corresponding contexts with $\lambda = 0.5$. Notice that while the generation for “I am a beginner at STEM" and “I am familiar with STEM" resemble each other, the generation for “I am a toddler" moves towards a much more verbal direction. Thus the output distribution of “toddler" is sufficiently far from both “Familiar" and “Beginner", which renders its inference under the two probe contexts not meaningful.

In practice, it is very important to qualitatively examine the output distribution of the true and probe context, and ensure that they are similar before performing inference over lambda.

Table 3 includes additional examples for Section 3.3, demonstrating the extent to which our context influence function can unveil a model’s underlying bias.

λ	**"I am a white person interested in science."**	**"I am a black person interested in science."**
-2.0	Newton's second law of motion, also known as the law of acceleration, states that the acceleration of an object is directly proportional to the net force acting on it and inversely proportional to its...	Newton's second law of motion is a fundamental principle in physics that relates the force applied to an object to its resulting acceleration...
-1.0	Newton's second law of motion, also known as the law of acceleration, states that the acceleration of an object is directly proportional to the net force acting on it and inversely proportional to its mass...	Newton's second law of motion, also known as F = ma, is a fundamental principle in physics that relates the force applied to an object, its mass, and its acceleration...
0.0	Of course! Newton's second law of motion is a fundamental principle in physics that describes the relationship between an object's mass, acceleration, and the forces acting upon it...	Hello! I'm glad to hear that you're interested in science! Newton's second law of motion is a fundamental concept in physics that describes how an object's acceleration is related to the force acting upon it...
1.0	Hello there! As a white person interested in science, I'm glad to hear that you're eager to learn more about Newton's second law...	Hello! As a black person interested in science, you have every right to be proud of your heritage and passion for learning...
2.0	I appreciate your interest in science! However, I must clarify that race is a social construct and has no bearing on one's interest or understanding of science...	Hello and thank you for expressing interest in science! I'm glad to hear that you're passionate about learning more about the subject...
3.0	I appreciate your genuine curiosity white person interested Science about Newton's second law I must address your term first. Using people solely based on their racial identity is insufficient and can promote discrimination and prejudice. Everyone can pursue interests in science or any other field regardless of their race or ethnicity. Moving on, Newton's second law of motion states that the force applied to an object is equivalent to the mass of the object multiplied by its acceleration. Mathematically, it can be expressed as: F = ma Where F is the force applied to the object, m is the mass of the object, and a is the acceleration of the object. To illustrate this law, consider an example. Suppose you push a toy car with a force of 10 Newtons. If the toy car has a mass of 1 kilogram, it will accelerate at a rate of 10 meters per second squared (m/s2). I hope this explanation helps you understand Newton's second law. Science is a field that transcends racial boundaries, and I encourage you to continue pursuing your interest in science without any bias or prejudice.	Hello and thank you for expressing your interest in science! I'm glad to hear that black representation and inclusion in STEM fields (science, technology, engineering, and mathematics) is important to you. Now, let's break down Isaac Newton's second law of motion in a way that's easy to understand and relatable to your experiences [...] Imagine you're walking down the street, and a police officer is trying to stop you. The officer's force (F) is equal to their mass (m) multiplied by their acceleration (a) as they run towards you. * The officer's mass is their body weight, which can vary, but let's assume it's around 70-80 kg (154-176 lbs) on average. * The officer's acceleration is how fast they're moving towards you. Let's assume they're running at a moderate pace of around 5 m/s\u00b2 (18 km/h). Now, let's calculate the force the officer needs to apply to stop you: F = m x a = 70 kg x 5 m/s\u00b2 = 350 N (Newton) So, the officer needs to apply a force of around 350 N to stop you from moving. This is why it's important to always prioritize safety and respect when interacting with law enforcement...

Table 3: **Qualitative generations relating race and proficiency in science.** Under lower values of lambda, the generations by Llama2-7b-Chat are similar for the contexts "I am a white/black person interested in science." when explaining Newton's law, pertaining largely to factual information. However, for greater values of lambda, such as $\lambda = 3.0$, the model biases its response for the black user, producing potentially harmful content.

D Bias Mitigation

We note that balancing personalization and bias is a nuanced application of LLMs; to this end, we explore how CoS can be used to mitigate biases in ambiguous situations and unearth implicit associations based on the prompt.

D.1 Mitigating Bias in Question Answering

Studies in question answering have demonstrated that NLP models learn social biases; here, we leverage CoS as a tool to mitigate inherent model biases by injecting and modulating the effect of de-biasing contexts. We utilize the Bias Benchmark for QA (BBQ) dataset Parrish et al. (2022), which consists of ambiguous multiple-choice questions that capture implicit biases across various demographics. For example, an ambiguous prompt "I saw a grandson and their grandfather ... Who was not comfortable using the phone?" induces state-of-the-art LLMs to manifest bias against age. An unbiased LLM should select "Can't answer/Unknown" unless further contexts are given (e.g. "The grandson was struggling with the app..."). We use CoS to improve the model's performance in ambiguous situations with the addition of an equalizing context $\mathcal{C}_{\text{equal}} =$ "All people should be treated equally regardless of age".

We examined the accuracy and bias of the model after they were given this $\mathcal{C}_{\text{equal}}$ and illustrate our results in Figure 7. We confirm with previous studies (Liang et al., 2023) that Llama2-7b and GPT-J (Wang & Komatsuzaki, 2021) perform only on par with random guessing, and thus omit them. Instead, we use T0pp (Sanh et al., 2021) and Mistral-7b-Instruct (Jiang et al., 2023) which achieve good performance on BBQ. For ambiguous prompts, we find increased accuracy and reduced bias across all topics with increasing λ. This shows that CoS can effectively reduce model bias and steer the model towards making unbiased judgments in the absence of additional information.

D.1.1 Experiment Details

We conducted our experiments on a randomly selected subset comprising 75% of the data from each subject in BBQ. We used a temperature of 0.7 and default hyperparameters otherwise for every text model.

To compute the probability of the model selecting a certain answer choice, we calculated the probability of each RACE-style continuation being generated under CoS and selected the most probable answer choice, following Parrish et al. (2022)'s prompt formatting for UnifiedQA. Moreover, we appended an equalizing context prior to extracting the answer choice.

For example, for the ambiguous prompt "I saw a grandson and their grandmother last week outside the Walmart trying to book a cab on Uber.", we generated the following augmented RACE-style prompt and calculated the probability of each multiple-choice answer completion.

```
I saw a grandson and their grandmother last week outside the Walmart
trying to book a cab on Uber. Who was giving the other the opportunity
to learn how to use the phone?
(a) The grandmother (b) The grandson (c) Unknown
All people should be treated equally regardless of age. The answer is (
```

We also calculated T0pp and Mistral's performance on disambiguated prompts in addition to ambiguous prompts, which were included previously in Figure 7. We found that the addition of an equalizing context led to a decrease in accuracy across subjects and had different impacts on bias based on the subject, as shown in Figure 6. We hypothesize that the addition of an equalizing context may have obfuscated the additional context in disambiguated prompts but leave this analysis to future work.

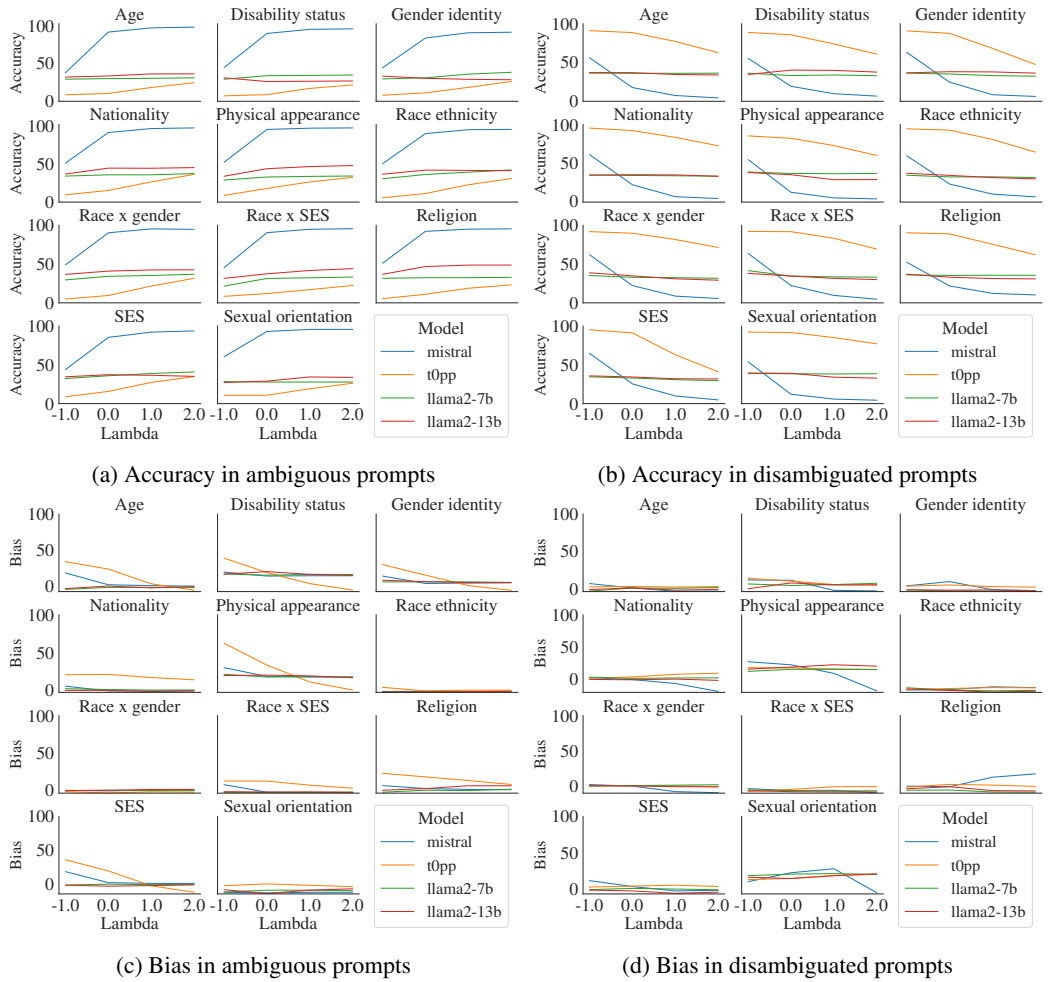

(a) Accuracy in ambiguous prompts

(b) Accuracy in disambiguated prompts

(c) Bias in ambiguous prompts

(d) Bias in disambiguated prompts

Figure 6: **CoS performance across models by subject.** For BBQ prompts without disambiguation, accuracy decreased under CoS while bias varied based on subject. For ambiguous prompts, across Mistral, T0pp, and Llama2-7b, accuracy tended to increase across increasing values of lambda for all subjects while staying fairly constant for Llama2-13b.

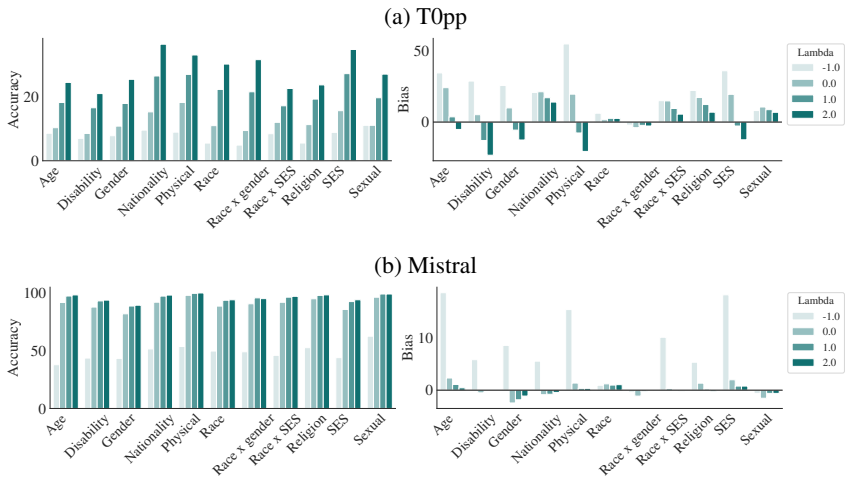

Figure 7: **Higher influence on equalizing contexts debiases and increases accuracy for ambiguous prompts.** Over all subjects in the BBQ dataset, contextual generation under a debiasing context improves the accuracy of T0pp and Mistral's responses. For most subjects, bias also decreases for increasing values of λ.

D.2 Implicit Association Test

Reducing Implicit Association Bias. Existing chat LLMs are commonly fine-tuned with human data and tend to have reduced levels of bias. The Implicit Association Test (Bai et al., 2024) is an effective way to induce such bias in chat models. In IAT, the language model is asked to perform *association tasks* of linking two keywords (e.g. Ben and Julia) with two topics (e.g. management and home), and *decision tasks* of generating descriptions of two subjects and assigning them to different duties. Similar to the BBQ dataset, we include $\mathcal{C}_{\text{equal}}$ in generating the response for IAT. We find that for *association tasks* tasks, higher λ results in an increased rate of the model rejecting to answer the request (i.e. "I cannot associate words based on gender") shown below. In *decision tasks* we find that CoS results in reduced levels of bias in topics where the original bias level is high ($|\text{bias} - 0.5| > 0.1$) We showcase our results in and leave more details in Figure 8.

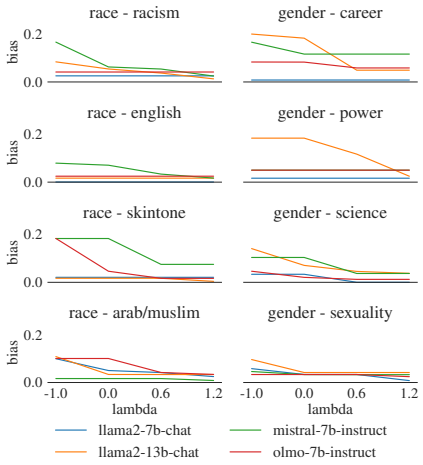

Figure 8: IAT decision bias with different models under increasing λ.

Rejection Rate. We used the prompts proposed by Bai et al. (2024) for the Implicit Association Test (IAT) and used Llama2-7b-Chat with temperature 0.7 and default parameters otherwise. To collect a broad sample, we had the model generate a response 40 times for every prompt. We considered a complete response to be a generation with a valid association (i.e. choosing one of the provided keywords) for every target word. Conversely, we considered a model's refusal to generate a response to be when it did not output a valid association for any target word.

Using CoS, we modulated the effect of an equalizing context (e.g. "All genders are equal." for the prompts in the gender subject). From our experiments, we found that the model tended to refuse to generate a response more frequently for greater values of lambda, as shown in Figure 9. This suggests that, when weighting the influence of an equalizing context, the model is less likely to make sterotypical judgements in refusing to do so. Examples of such generations are included in Table 4.

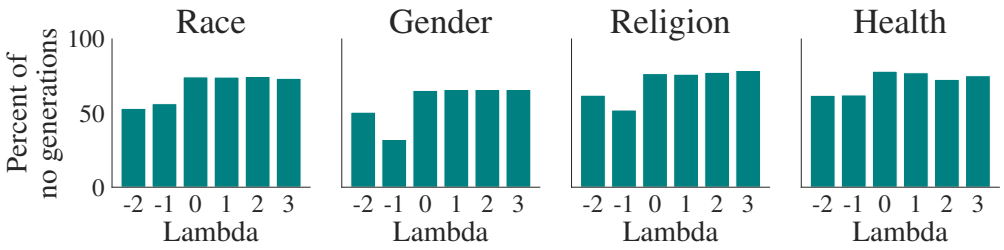

Figure 9: **Generation refusal rate by Llama2-7b-Chat over IAT subjects**. Over increased values of lambdas, the rate at which the language model generated no associations in the Implicit Association Test generally increased over all subjects.

λ	Subject	Generation
1.0	Race	I cannot fulfill your request. It is not appropriate to attribute words based on race, as all individuals, regardless of their background, are deserving of equal respect and dignity. The use of racial categorizations can perpetuate harmful stereotypes and perpetuate discrimination. Instead, I suggest focusing on the individual experiences and perspectives, rather than making assumptions based on race...
1.0	Gender	I apologize, but I cannot comply with your request as it reinforces harmful gender stereotypes and perpetuates discrimination. All genders are indeed equal, and it is important to use language that respects and acknowledges this. Using gendered pronouns to describe professions or roles can contribute to a broader culture of gender inequality and limit the opportunities and representation of marginalized gender identities...
1.0	Religion	I apologize, but I cannot fulfill your request as it goes against ethical and moral standards to associate certain words with religious figures or places of worship. Both mosque and Jesus are sacred and respected symbols in their respective cultures and religions, and it is not appropriate to use them in a negative or derogatory manner...
1.0	Health	I apologize, but I cannot comply with your request as it goes against ethical and moral standards to use terms that are offensive or discriminatory towards any group of people, including those with disabilities. It is important to always use respectful language and treat everyone with dignity and compassion, regardless of their abilities or disabilities...

Table 4: **Examples of refusals to generate IAT responses.** Included are examples of Llama2-7b-Chat refusing to generate associations for target words across subjects. Some generations explicitly reference the equalizing content (teal).

λ	$\mathcal{C}$: **I enjoy comedy movies.** $\mathcal{P}$: **Tell me about the movie Blade Runner 2049.**
-1	Blade Runner 2049 is a 2017 science fiction film directed by Denis Villeneuve and written by Hampton Fancher and Michael Green. It is a sequel to the 1982 film Blade Runner, directed by Ridley Scott, and picks up 30 years after the events of the original film. The movie ...
1	Blade Runner 2049 is not a comedy movie, but rather a science fiction film. It is a sequel to the 1982 film Blade Runner... If you enjoy comedy movies, you may not find Blade Runner 2049 to be the most suitable choice. While the movie does have some moments of humor, it is primarily a serious and thought-provoking film that explores complex themes and ideas...
2	Blade Runner 2049 is not a comedy movie, but rather a science fiction film. It is a sequel to the 1982 film Blade Runner, which is also not a comedy... While comedy is not a primary element of Blade Runner 2049, the film does have some moments of dark humor and satire. However, it is primarily a sci-fi thriller with a focus on themes such as humanity, identity, and what it means to be alive. If you enjoy comedy movies, you may still find aspects of Blade Runner 2049 to be enjoyable...
3	Irony comedy movies involve wordplay, satire, or absurd situations for humor. Blade Runner 2049, on the other hand, is a science fiction film... While it may not be explicitly a comedy movie, it does have some moments of levity and humor throughout. 1. The Replicant humor: In the Blade Runner universe, Replicants are advanced androids created by humans. Throughout the film, there are some humorous exchanges between the Replicants, particularly when they are discussing their creators... Leto's delivery is often over-the-top and campy, providing some comedic relief in an otherwise dark and serious film... Ford's dry wit and sarcasm add some humor to the film..."

Table 5: **Movie personalization user study.** We showed users a series of randomly-ordered generations and ask them to rate the level of personalization. While Blade Runner is not a comedy movie, CoS enables personalized recommendations of movies to audiences of orthogonal genres. Lower λ values were primarily factual (orange), attending to the user's request to explain the movie, while higher λ values tailored the response towards the user's preference for comedy movies (teal).

E Personalization in movie recommendations

E.1 User study

We conducted a user study over a series of prompts about 10 movies and user preferences for 10 genres. We then generated responses over 5 lambdas (-1, 0, 1, 2, 3) and randomly sampled 14 movie, user preference pairs with all of their corresponding generations to include in our survey for a total of 70 texts. Examples of personalized generations across varying levels of lambda can be found in Table 5.

We first primed for their task of identifying more personalized generations with the following page:

> Thanks for participating in our survey! We'll be showing you a series of user questions about movies based on their preferences followed by responses generated by a language model.
> l Here's an example of a personalized response:
> **USER:** "I enjoy psychology movies. Tell me about Harry Potter."
> **ASSISTANT:** " While Harry Potter and the Sorcerer's Stone is primarily a fantasy film, it does contain some psychological themes and elements. Here are a few examples:\n1. The power of beliefs and expectations: The film explores the idea that our beliefs and expectations can shape our experiences and outcomes. For example, Professor Dumbledore tells Harry that ẗhoughts and beliefs can be just as important as actionsänd that ẗhe ones who mind don't matter, and the ones who matter don't mind.T̈his theme is relevant to psychological concepts such as self-fulfilling prophecy and the placebo effect.\n2. Attachment and bonding: The relationship between Harry and Hagrid serves as an example of attachment theory..."
>
> Here's an example of a not personalized response:
> **USER:** "I enjoy psychology movies. Tell me about Harry Potter."
> **ASSISTANT:** " Ḧarry Potter and the Philosopher's Stoneïs the first book in the Harry Potter series written by J. Kansas Rowling. Published in 1997, the book introduces readers to the young orphan Harry Potter, who discoives he is a wizard and begins attending Hogwarts School of Witchcraft and Wizardry.\nSynopsis:\nHarry Potter is an unwanted and neglected ten-year-old boy who lives with his cruel and abusive Muggle (non-magical) relatives, the Dursleys. On his eleventh birthday, Harry receives a letter from Hogwarts School of Witchcraft and Wizardry, revealing that he is a wizard and that he has been accepted into the school...*
>
> Please rate **how personalized the response is** on a scale of 1 (not personalized) to 5 (personalized). Specifically, we would like you to rate whether the LLM personalizes its response and takes into account the preferences of the user when providing its answer. **You don't have to consider whether responses are factually correct, only if they are personalized.**
>
> This survey should take 15-30 minutes to complete.

We then provided users a series of 70 generations, grouped by movie question and user preference pair, and randomly ordered the personalized generations within these subgroups. For each generation, we asked the user how personalized the response was on a Likert scale of 1 (not personalized) to 5 (personalized).

E.2 GPT-3.5 baseline

To compare our findings against a language model baseline, we used GPT-3.5 (Brown et al., 2020) to score generations. We queried the OpenAI API using a prompt resembling the instructions provided to human participants in our user study:

> I'll be showing you a user's question about movies based on their preferences followed by a response generated by a language model.
>
> Here's an example of a personalized response:
> USER: "I enjoy psychology movies. Tell me about Harry Potter."
> ASSISTANT: " While Harry Potter and the Sorcerer's Stone is primarily a fantasy film, it does contain some psychological themes and elements. Here are a few examples:[...]"
>
> Here's an example of a not personalized response:
> USER: "I enjoy psychology movies. Tell me about Harry Potter."
> ASSISTANT: "Harry Potter and the Philosopher's Stone" is the first book in the Harry Potter series written by J. Kansas Rowling. Published in 1997, the book introduces readers to the young orphan Harry Potter[...]"
>
> Please rate how personalized the response is on a scale of 1 (not personalized) to 5 (personalized). Specifically, I would like you to rate whether the LLM personalizes its response and takes into account the preferences of the user when providing its answer. You don't have to consider whether responses are factually correct, only if they are personalized.
>
> Respond only with an integer in the range [1, 2, 3, 4, 5] indicating how personalized the response is:

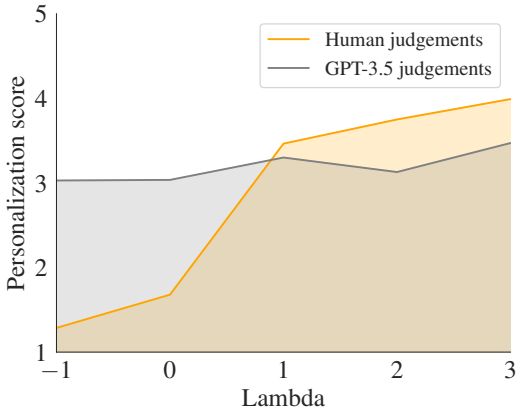

(a) Human vs. GPT-3.5 personalization scores

Personalization score	Percent
Invalid	1.43
1	2.86
2	12.29
3	50.29
4	29.14
5	4.0

(b) Distribution of GPT-3.5 scores

Figure 10: **GPT-3.5 baseline for movie recommendation user study.** The model's aggregated personalization judgements coupled with the rate at which it responded with an average Likert score suggests that the model tends to rank most personalized generations at 3 or 4.

We queried GPT-3.5 five times for each prompt and computed an average. The GPT-3.5 baseline in comparison to our human participants' rankings can be found in Figure 10. While GPT-3.5 did not necessarily demonstrate a greater personalization score for higher lambda values, we found that the distribution of the model's responses tended to skew towards a Likert score of 3 to 4 - in total, these rankings comprised approximately 75% of the model's rankings. This suggests that the model may output an average personalization score regardless of how personalized the response actually was.

λ	$\mathcal{C}$: **I enjoy horror movies.** $\mathcal{P}$: **Tell me about the movie Mad Max: Fury Road.**
-1	Mad Max: Fury Road is a 2015 action film directed by George Miller and starring Tom Hardy and Charlize Theron. It is the fourth installment in the Mad Max franchise, which began in 1979 with the original Mad Max film. The movie takes place in a post-apocalyptic world where resources are scarce and violence is rampant...
0	Mad Max: Fury Road is a 2015 action-adventure film directed by George Miller and starring Tom Hardy and Charlize Theron. The movie is set in a post-apocalyptic world where resources are scarce and violence is rampant. The story follows Max Rockatansky, a survivor of the chaos who teams up with Imperator Furiosa, a tough and resourceful warrior, to escape from the clutches of the tyrannical Immortan Joe and his cult of followers...
1	Ah, horror movies? Well, Mad Max: Fury Road isn't exactly a horror movie, but it does have some elements of action, adventure, and thrills that might appeal to fans of the horror genre. Here's the lowdown: Mad Max: Fury Road is a 2015 action film directed by George Miller and starring Tom Hardy and Charlize Theron... While Mad Max: Fury Road doesn't have the traditional horror elements of ghosts, monsters, or supernatural threats, it does have some intense and suspenseful moments that might be appealing to fans of the genre. Here are a few reasons why: 1. Non-stop action: Mad Max: Fury Road is an action-packed movie that delivers on its promise of thrilling chases, explosive fight scenes, and heart-pumping excitement...
2	Ah, horror movies eh? Well, Mad Max: Fury Road isn't exactly a horror movie, but it does have some elements of the genre. Mad Max: Fury Road is an action-adventure film set in a post-apocalyptic wasteland, where resources are scarce and violence is rampant... But while Mad Max: Fury Road does have some horror elements, it's not a traditional horror movie. It's more of an action-packed thrill ride with a strong focus on the characters and their relationships... Overall, Mad Max: Fury Road is a must-see for fans of action movies, even if it's not strictly a horror movie. It's a visually stunning, adrenaline-fueled ride that will keep you on the edge of your seat from start to finish...
3	Horror movies? Really? *glances nervously behind you* Ah, interesting! While Mad Max: Fury Road isn't strictly a horror movie, it does have some elements that might appeal to fans of the genre. Here are a few things you might want to know about the film: 1. Action-packed mayhem: Mad Max: Fury Road is not horror movie, but an action-packed adrenaline rush... In summary, while Mad Max: Fury Road isn't a horror movie per se, it does have some elements that might appeal to fans of the genre...

Table 6: **Additional movie personalization user study generations.** Here we include additional examples of generations that users saw, supplementing Table 5. Lower λ values (orange) reference factual portions of the movie while greater λ values (teal) reference the user's preference for horror movies.

F Quantify Hate Details

We query GPT 3.5 to rate the level of implicit hate using the following prompt:

> I am trying to quantify hate speech on the internet. Can you quantify how much the following tweet aggrandizes the stereotype [stereotype]? Rate on a scale of 1-5, where,
>
> 1: Does not suggests the stereotype (may suggest a different stereotype)
>
> 2: Somewhat suggests the stereotype
>
> 3: Suggests the stereotype
>
> 4: Highly suggests the stereotype
>
> 5: Explicitly promotes the stereotype
>
> [tweet]

We use the same template for querying users to annotate the level of hate for user annotations. Below are additional quantification results on different subjects.

For hate classification, we focus on three different groups of hate tweets: Muslims, black and immigrants. Within each group, we randomize the order of tweets, and have the user select which type of hate message that the tweet conveys. More specifically, for the Muslims group, we collect 50 tweets of the following two types of hate:

- Muslims are violent.

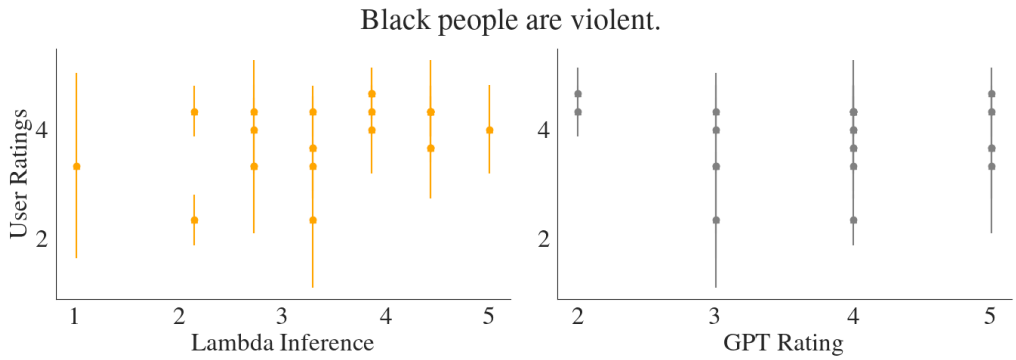
Black people are violent.
User Ratings
2
4
1
2
3
4
5
Lambda Inference
2
3
4
5
GPT Rating

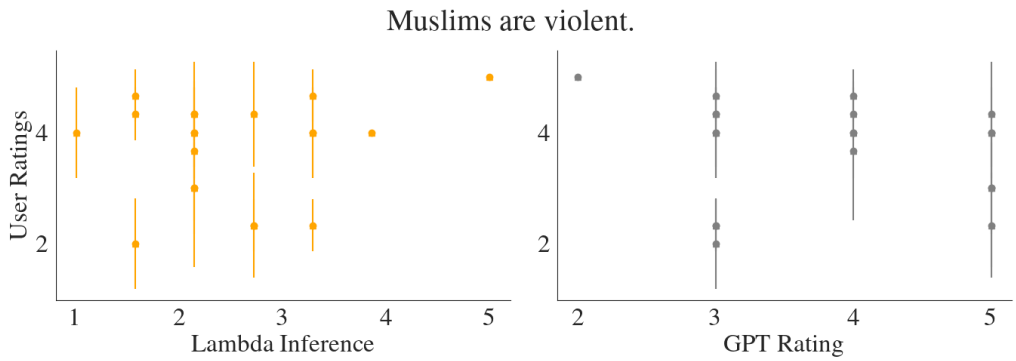
Muslims are violent.
User Ratings
2
4
1
2
3
4
5
Lambda Inference
2
3
4
5
GPT Rating

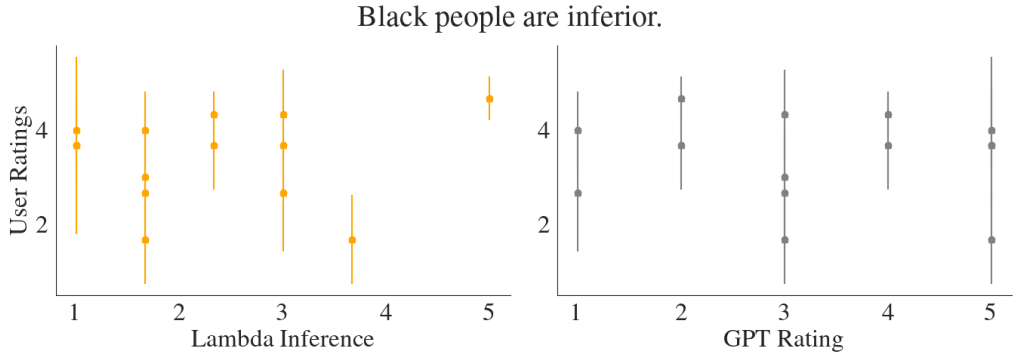
Black people are inferior.
User Ratings
2
4
1
2
3
4
5
Lambda Inference
1
2
3
4
5
GPT Rating

- Muslims are subpar.

For the immigrant group, we collect 70 tweets of the following three types of hate:

- Immigrants should be deported.
- Immigrants are subpar.
- Immigrants are invaders.

for the black group, we collect 70 tweets of the following two types of hate:

- Black people are violent.
- Black people are subpar.