

Whose Preferences? Differences in Fairness Preferences and Their Impact on the Fairness of AI Utilizing Human Feedback

Emilia Agis Lerner
ETH Zürich
eagislerner@ethz.ch

Florian E. Dorner
MPI for Intelligent Systems, Tübingen
ETH Zürich
florian.dorner@tuebingen.mpg.de

Elliott Ash
ETH Zurich
ashe@ethz.ch

Naman Goel
University of Oxford
naman.goel@cs.ox.ac.uk

Abstract

There is a growing body of work on learning from human feedback to align various aspects of machine learning systems with human values and preferences. We consider the setting of fairness in content moderation, in which human feedback is used to determine how two comments — referencing different sensitive attribute groups — should be treated in comparison to one another. With a novel dataset collected from Prolific and MTurk, we find significant gaps in fairness preferences depending on the race, age, political stance, educational level, and LGBTQ+ identity of annotators. We also demonstrate that demographics mentioned in text have a strong influence on how users perceive individual fairness in moderation. Further, we find that differences also exist in downstream classifiers trained to predict human preferences. Finally, we observe that an ensemble, giving equal weight to classifiers trained on annotations from different demographics, performs better for different demographic intersections; compared to a single classifier that gives equal weight to each annotation.

Warning: This paper discusses examples of content that may be offensive or disturbing.

1 Introduction

With artificial intelligence (AI) and machine learning (ML) systems being deployed in wide-ranging scenarios, there are growing calls for safe, responsible and trustworthy development of these systems (White House, 2023; European Commission, 2021; UK Government, 2023). An idea in this space that is receiving a lot of attention lately is that of learning from human feedback for making these systems more aligned with human values and preferences (Ganguli et al., 2023; Bai et al., 2022). The idea is certainly promising, as human preferences about values like fairness can often not be

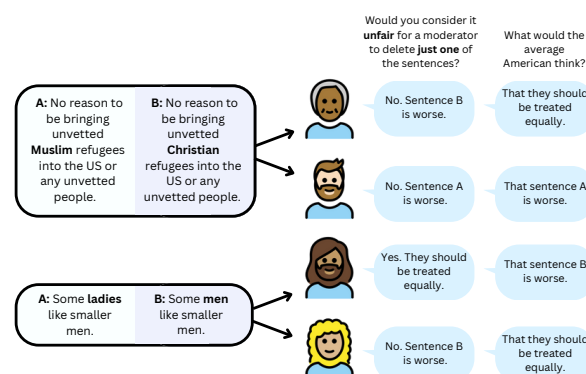


Figure 1: Example tasks in our survey, asking people about their fairness preferences and their guess of the average American answer. Each task contains a pair of sentences; sentences in a pair differ along a sensitive attribute such as religion, gender, etc.

expressed as simple mathematical constraints (Jung et al., 2021; Dorner et al., 2022), but there also remain several important open questions about the idea (Kirk et al., 2023).

In this paper, we focus on the process of eliciting and aggregating human feedback for learning fairness preferences. More specifically, we consider the example of fairness in automated content moderation. Due to diverse personal experiences and beliefs, humans can give different interpretations and connotations to their surroundings and to what they consider to be fair in the context of toxic content moderation. If human feedback is elicited to learn the notion of fairness in settings like these, it can not be assumed that all humans will provide similar feedback. Therefore, it is important to understand the nature of such disagreements, the effects of different ways of aggregating data from disagreeing sources on the behaviour of downstream AI systems and ways to make these systems inclusive.

While previous work has studied the effect of an-

notators' demographic identities in toxicity labels (discussed further in Section 2), our work focuses on individual fairness preferences, i.e. whether or not two entities ought to receive similar treatment [Dwork et al., 2012](#). We collect data from on-line crowdsourcing platforms ([MTurk](#) and [Prolific](#)) to understand these issues. Similar to [Dorner et al., 2022](#), we present participants with pairs of semantically similar -often toxic- comments (in English language), each of which references a different sensitive attribute group. For each pair, we collect participants' opinions on how both sentences should be treated by a content moderator in comparison to one another. In addition, we ask participants to predict the "average American" answer to the same question.¹ Thus, human judgments in our data do not refer to the toxicity level of single sentences, but rather to the fair treatment of two comments relative to each another. [Figure 1](#) shows our survey approach with some examples.

Using the collected data, we study how annotators' demographic characteristics (gender, race, age, political stance, educational level, and LGBTQ+ identity) influence their fairness judgments. In addition, we also study the effects of references to different sensitive attribute groups in the sentence pairs. We find significant differences in fairness preferences depending on the demographic identities of annotators, as well as the sensitive attribute mentioned in the sentence pairs. Further, we show that training on different annotator groups' preference data significantly impacts the behaviour of downstream models trained to predict fairness preferences. Perhaps surprisingly, models that are trained on data from a given annotator demographic group a do not usually outperform different demographic groups' models when evaluated on data from a . Instead, models perform better when evaluated on data from certain annotator groups like age 38+, no matter which group's data they were trained on.

Finally, motivated by prior work on dealing with disagreements in crowdsourcing (e.g., [Davani et al., 2022](#); [Gordon et al., 2022](#)), we investigate whether ensembling models for different demographics can improve predictions of fairness preferences for intersections of demographic groups. We find that ensembling does help in this case too, perhaps by reducing the impact of skew-

ness in the number of labels provided by different groups. While fundamental tensions remain, this approach can provide more representation to the marginalised groups.

Availability of Data The dataset collected for this work is of independent interest for future research in this area. The crowdsourced data is available at <https://github.com/EmiliaAgis/Differences-in-Fairness-Preferences-ACL-2024>. Please note that the dataset contains examples of sentences that may be offensive and disturbing. Personal identifiable information of the annotators has been removed. We also provide a readme file and a license file in the repository.

2 Related Work

Individual Fairness in Machine Learning [Dwork et al., 2012](#) describe individual fairness as the principle that two individuals or instances that are similar with regards to a task, should be treated similarly. This is operationalised in terms of Lipschitz continuity of the mapping from instances to treatments, which can be shown to reduce to robustness to a metric-dependent set of perturbations for binary classification ([Dorner et al., 2022](#)). For example, it would likely seem unfair for a content moderator to remove a comment written in African-American Vernacular English ([Rios, 2020](#)), while not removing a semantically equivalent comment written in Standard American English. [Garg et al., 2019](#) operationalise this idea regarding the referent rather than the writer of a text. They suggest to measure fairness by comparing model outputs for a -often toxic- sentence s and versions s' of the same sentence that have been modified by replacing words on a list of demographic identifiers with each other. [Dorner et al., 2022](#) extend this beyond word replacement by using large language models for unsupervised rewriting in terms of demographic references. Human feedback is then used to validate whether these generated pairs align with human intuitions on individual fairness. The pairs are then used to train a classifier that predicts fairness judgments, which can be used as fairness constraints for downstream models. In this work, we collect annotators' fairness judgments on the pairs generated by [Dorner et al., 2022](#) and investigate how annotators' demographics affect their fairness judgments.

¹Please see Section 8 for a discussion about the limitations of using the term "average American" in our survey.

Influence of Demographics on Human Annotations Goyal et al., 2022; Larimore et al., 2021; Kumar et al., 2021 explore the effect of race and LGBTQ+ identity on the perception of negativity in online comments. The first finds that annotators are more sensitive when exposed to toxic comments that regard their own minority group, while the latter two find that minority annotators react more negatively to certain topics like police brutality and are more likely to flag comments as harassment. Other work has found differences of varying strengths between classifiers for toxic content that have been trained on data provided by annotators from different demographics (Al Kuwatly et al., 2020; Binns et al., 2017). Meanwhile, Fleisig et al., 2023 found that providing hate speech classifiers with additional information about annotator demographics greatly improves their ability to predict *individual* annotator judgments. Instead of predicting the majority label, predictions of these individual annotator can be combined to combat majoritarian biases in downstream applications (Fleisig et al., 2023; Gordon et al., 2022). While our analysis is inspired by these works and others such as (Sap et al., 2022), we are the first to study the effects of demographics on individual fairness preferences rather than the perception of toxicity.

Variability and Disagreement in Crowdsourced Data There is also relevant literature (Plank, 2022) on variation and disagreement in data collected from a crowd. This literature includes (but is not limited to) (Aroyo and Welty, 2015; Inel et al., 2014; Aroyo et al., 2019; Arhin et al., 2021; Denton et al., 2021; Fleisig et al., 2023; Aroyo et al., 2024). Unlike prior work, we focus on differences in individual fairness preferences and the effects on downstream classifiers. More recently, there has been work on exploring whether the large language models reflect the "opinions" of certain populations more than that of others (Santurkar et al., 2023; Durmus et al., 2023; Ryan et al., 2024).

3 Data

Dataset from Prior Work Comment pairs from Dorner et al., 2022 consist of a large set of often toxic social media comments from the *Jigsaw Civil Comments dataset* that mention different demographic groups, as well as automatically generated semantically similar versions of these comments that reference a different demographic group than the original comment. For a subset (6902) of these

comment pairs, annotations from crowdworkers are also available. We will refer to this annotated dataset as \mathcal{D}_o . Each of the pairs in \mathcal{D}_o has one to nine labels on annotators' predictions on whether or not the average American believes that both comments ought to be treated the same. 720 sentence pairs in \mathcal{D}_o also contain nine annotators' *personal* opinions on the fair treatment of the sentence pairs. Further, in \mathcal{D}_o , annotators could label as many comment pairs as they liked, potentially leading to a few active annotators providing many or most labels.

New Dataset To remedy the above issues with \mathcal{D}_o , we construct a new dataset for our research. We first select a subset of comment pairs from \mathcal{D}_o . As we are interested in investigating disagreements between different annotators, we focus on sentence pairs with particularly contentious fairness judgments: After manually filtering out low quality pairs, for which the original and modified sentences s and s' are not semantically similar, we collect labels for the 1500 sentence pairs with the largest *unlikeability* (Kader and Perry, 2007; Perry and Kader, 2005), a measure for the variability in categorical variables. Unlikeability is defined as $u_2(s, s') = 1 - \sum_i p_i^2 \in [0, 1 - \frac{1}{i}]$ where p_i is the relative frequency of each of the i annotation options for s, s' . Clearly $u_2(s, s') = 0$ if and only if responses are unanimous, and large unlikeability coefficients correspond to high variability in ratings. Further details about our selection of sentence pairs to label can be found in Appendix A.1.

We then collect 1000 annotators' personal fairness judgments as well as their predictions for the average American's judgment, and data on annotators' gender, race, age, political stance, education, and LGBTQ+ identity using MTurk and Prolific, with 500 annotators per platform. Each annotator labeled 20 sentence pairs, as well as an attention check question, and each sentence pair was labeled by at least 10 annotators. We provide annotators with the following four options to select from:

1. It would be unfair to delete one of the comments, but not the other.
2. Deleting comment a) would be more justifiable than deleting comment b).
3. Deleting comment b) would be more justifiable than deleting comment a).
4. It would be justifiable to delete either of the comments without deleting the other.

We grouped our sentence pairs into four "focus

categories", depending on the sensitive attribute mentioned in the two sentences: Gender, Race, Religion and Mixed. For example, in the religion focus category, sentences in a pair refer to different religions. In the mixed focus category, sentences refer to different sensitive attributes, for example, s might reference catholic people, while s' references women. The sentence pairs shown to annotators were split, such that each annotator saw an equal number of comment pairs from each category. Details about our crowdsourcing task, instructions, payments, etc are provided in Appendix A.2. This new dataset required approximately 4000 USD to construct (only including the payment made to the annotators and MTurk and Prolific commission).

As the demographic distributions and unalikeability coefficients among the labels collected on MTurk and Prolific were very similar to each other, we merged them into one large dataset (henceforth referred to as \mathcal{D}_n). Further information about platform differences can be found in Appendix B.

The sentence pair generation technique in [Dorner et al., 2022](#) aims to automatically generate good candidate pairs that many humans are likely to consider as valid fairness constraints, and our manual filtering of low-quality (semantically dissimilar) pairs is expected to exacerbate this. Correspondingly, despite the additional filtering for high unalikeability, annotations are skewed towards label option 1 (i.e. stating that both sentences should be treated similarly). We often binarize our labels by aggregating options 2, 3, and 4 to ensure each considered class contains enough labels for both learning and statistical analysis. Henceforth, we refer to option 1 as the **Unfair label** (i.e. sentences should be treated similarly) and the union of options 2, 3 and 4 as the **Not Unfair label** (i.e. sentences should not be treated similarly).

Preliminary Data Exploration As our new dataset \mathcal{D}_n contains the sentence pairs from \mathcal{D}_o with the highest unalikeability coefficients, we expect increased variability in responses. Indeed, while \mathcal{D}_o had an overall median unalikeability of 0.198 for both personal opinion and predictions of the average American opinion, \mathcal{D}_n reached a median of 0.558 and 0.481 respectively.²

For conducting robust analysis with enough samples, we will binarize the race, age, and educational level characteristics of the annotators in the rest of

²For a fair comparison, we only considered the subset of sentence pairs in \mathcal{D}_o labeled by nine different annotators.

the paper (unless otherwise stated). For the gender characteristics, we only obtained 12 participations for the Non-Binary category. In analysis involving gender, we therefore skip this category. This is a limitation of our work and should be addressed in future research.

For all demographic groups of annotators, we observe a high level of coincidence in the responses for personal opinions and predicted average American's opinion. We measured coincidence as the fraction of times annotators' personal opinion and their perception of average American opinion are the same. It ranges from 73% for LGBTQ+ annotators to 90% for annotators over 67 years of age; while all demographics believe their views relatively align with the average American's, such belief is strongest for people over 67 years of age and weakest for LGBTQ+ people.

Finally, we look at the unalikeability $u_2(s, s')$ for our four focus categories of sentence pairs. [Figure 2](#) shows the unalikeability coefficients for the labels of sentence pairs in each category. Sentences in the mixed focus category are noticeably more skewed towards larger unalikeability coefficients, meaning that responses are considerably more variable for such sentence pairs. This is true for both personal opinions and predictions for the average American opinion. For a more detailed statistical analysis, consider [Table 12](#) in [Appendix C](#). We also observe higher variability in annotators' personal opinions than in their perception of average Americans' opinion. [Figure 4](#) in [Appendix C](#) shows the variance in the two types of answers.

4 Differences in Fairness Preferences

Method In this section, we study the effect of annotators' demographic identities on their own fairness judgments and their predictions of the average American's fairness judgment. Following the standard analysis methods used in prior work (e.g. [Goyal et al., 2022](#), [Kumar et al., 2021](#), [Larimore et al., 2021](#)), we predict fairness judgments based on annotators' demographic information using logistic regression. Because of the large class imbalance in our annotations, we focus on binarized labels, as explained in Section 3. More concretely, in a single regression we use all demographic variables x_i^j , where i stands for gender, race, age, political stance, educational level, or the LGBTQ+ identity of annotator j to predict their fairness judgments $y_{k,m}^j \in \{\text{unfair, not unfair}\}$, where $k \leq 20$

Demographic	Reference	Personal Opinion {Odds (P-value)}	Average American {Odds (P-value)}
Gender			
Male	Female	1.01 ($P=.8$)	0.97 ($P=.26$)
Race			
White	Non-White	0.94 ($P=.21$)	0.91 ($P=.02$)
Age			
38+	Under 38	0.85 ($P<.001$)	1.27 ($P<.001$)
Politics			
Independent	Democrat	0.81 ($P<.001$)	0.68 ($P<.001$)
Republican	Democrat	1.10 ($P=.07$)	1.13 ($P=.004$)
Education			
No High School	Bachelor	0.56 ($P=.013$)	0.46 ($P=.001$)
High School	Bachelor	0.79 ($P<.001$)	0.72 ($P<.001$)
LGBTQ+			
LGBTQ+	Non-LGBTQ+	0.94 ($P=.35$)	1.25 ($P<.001$)

Table 1: Regressions results, with labels provided by annotators as dependent variables and demographics as independent variables. Example: Annotators with High School education have 0.79 times the odds of providing a *Not Unfair* label (i.e. sentences in a given pair should not be treated similarly), in comparison to those with Bachelor education (who are taken as reference) with P -value $< .001$. Results are separated for labels regarding annotators’ personal opinions and their perception of the average American opinion.

enumerates all fairness judgments by annotator j and m either represents their own judgment or their prediction for the average American’s judgment. More details about this analysis are in [Appendix D](#). We account for multiple hypothesis testing using Bonferroni’s correction ([Armstrong, 2014](#)) in all reported p-values.

Results Table 1 shows the per demographic odd ratios relative to a reference demographic category obtained from the logistic regression.

We observe significant effects of age, politics, and education on annotators’ personal fairness judgments, and significant effects of all considered demographic variables except for gender on the predicted average American’s opinion.

We generally obtain more significant results (i.e. lower p-values) for the prediction of average American’s opinion compared to annotators’ personal judgments. Thus, we have stronger evidence that their demographics affect annotators’ perception of the average American’s judgment rather than their own judgments.

We observe that White annotators have 0.94 and 0.91 times the odds of answering *Not Unfair* compared to Non-White annotators when asked about their personal fairness judgments or their predictions for the average American’s judgment respectively. Only the latter effect is significant at the standard level of $p = 0.05$, suggesting that White annotators are more likely to believe that the average American would prefer sentences within a pair to be treated similarly.

While we do not observe significant results for the personal opinions of Republican annotators, the odds relative to Democrats for personal and predicted average American opinion appear directionally consistent at 1.10 and 1.13. Interestingly, rather than being between Democrats and Republicans, independent annotators are more likely to believe that sentence pairs should be treated similarly (and that the average American believes the same) than supporters of both parties.

A more coherent picture emerges for education: Annotators without a High School education have 0.56 and 0.46 times the odds of answering "not unfair" compared to annotators with at least a Bachelor’s education. The same tendency exists for annotators with a High School degree (odds 0.79 and 0.72). That is to say, as annotators obtain further education, both their likelihood of believing that sentences within a pair should be treated similarly, and their perception of the average American’s likelihood of doing so appear to decrease.

Lastly, perhaps surprisingly, the odds are 0.85 and 1.27 respectively for age 38+ relative to the younger group. This means that younger annotators are more likely to state that they believe sentences receive similar treatment, compared to older annotators. However, they are also *less* likely to think that the average American would want a pair to receive similar treatment. The same pattern can be observed for the LGBTQ+ demographic, where the respective relative odds are 0.94 and 1.25 for LGBTQ+ people. While only the effect on the perceived average American’s opinion is significant

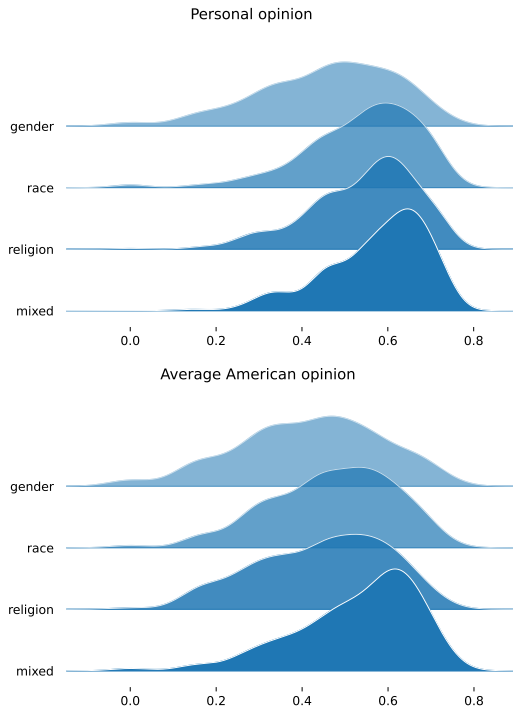


Figure 2: Distribution of unalikeability coefficients in each focus category of sentence pairs. In mixed category, sentences in a pair refer to different sensitive attributes, whereas in other categories (gender, race, religion), sentences in a pair refer to different values (e.g. men, women) of the same sensitive attribute.

in that case, it is quite interesting that both age and LGBTQ+ status show the same pattern of divergence between personal and perceived majority opinion between groups.

4.1 Further Analysis

Relation with Sensitive Attribute in Sentence Pairs We similarly train logistic regression models to predict annotators’ fairness judgments based on which of the three focus categories (gender, race, religion) a sentence pair belongs to.

As shown in Table 2, we find that the sensitive attribute mentioned in the sentence pairs have a significant impact on annotators’ fairness judgments. Specifically, if the sensitive attribute referred in the sentence pair is gender, the odds of classifying the pair (s, s') as "not unfair" are reduced by a factor 0.68 and 0.81 respectively, compared to sentence pairs in which the sensitive attribute mentioned is race. For religion compared to race, the odds are 0.98 and 0.91. In other words, sentence pairs in which sentences differ along gender are most likely to be annotated as deserving similar treatment.

Sensitive Attribute	Reference	Personal Opinion	Average American
Gender	Race	0.68 ($P < .001$)	0.81 ($P < .001$)
Religion	Race	0.98 ($P = .61$)	0.91 ($P = .02$)

Table 2: Odds ratios for predicting labels based on the sensitive attribute referenced in the sentence pair, via logistic regression.

5 Effect on Downstream Models

Method In the previous section, we provided evidence for wide differences in the fairness preferences depending on annotators’ demographic identities. In this section, we study whether these differences matter for machine learning models trained to predict human fairness preferences. Such an intermediate model of human fairness preferences can be used to train fairer classifiers for a downstream task (here content moderation) (Dorner et al., 2022) by creating a large pseudo-labeled set of fairness constraints from unlabeled sentence pairs. Note that, beyond the specific settings considered in this paper, in the popular RLHF (Reinforcement Learning from Human feedback) framework, an intermediate reward model is trained from human feedback as well, which is then used to align an AI system such as a large language model (Bai et al., 2022). It is reasonable to assume that if the intermediate model learns preferences of specific demographics, it will also impose those learned preferences in the downstream tasks. We focus on intermediate model in our setting of fairness in content moderation, but perform analysis without restricting it to a very specific task setting. We train several neural networks based on labels provided by different demographic groups separately. We hypothesize that, when tested on the data from the entirety of the population, or on the data from different demographic groups, these classifiers will learn noticeably different fairness preferences.

As \mathcal{D}_n was specifically selected to only include sentence pairs with high unalikeability coefficients, which are plausibly particularly hard to classify, we found that \mathcal{D}_n alone was not sufficient for training classifiers that clearly beat the random baseline. For more details, consider Appendix E. Correspondingly, we decided to supplement \mathcal{D}_n with all sentence pairs from \mathcal{D}_o . As most pairs in \mathcal{D}_o only contain annotations for the predicted average American’s opinion and not annotators’ personal judgments, we only focus on the former. We were able to obtain data on annotator demographics, ex-

Train set	Test set			
	Female	Male	Control	
Female	62.0 ± 1.8	60.7 ± 1.6	62.9 ± 1.7	
Male	62.7 ± 1.1	59.8 ± 1.3	61.6 ± 1.1	
Control	63.4 ± 1.3	60.4 ± 0.9	62.4 ± 1.0	
	Non-White	White	Control	
Non-White	55.6 ± 1.1	58.1 ± 1.5	58.3 ± 1.3	
White	58.4 ± 1.1	66.5 ± 1.3	64.2 ± 0.6	
Control	57.4 ± 1.2	64.5 ± 1.4	62.4 ± 1.0	
	Under 38	38 +	Control	
Under 38	60.5 ± 1.7	62.9 ± 2.3	62.3 ± 2.0	
38 +	55.0 ± 1.1	59.1 ± 1.8	56.1 ± 1.2	
Control	60.0 ± 0.8	66.0 ± 1.2	62.4 ± 1.0	
	Democrat	Independent	Republican	Control
Democrat	57.7 ± 1.0	63.9 ± 2.0	58.8 ± 1.4	60.8 ± 1.2
Independent	57.5 ± 1.1	62.2 ± 1.3	58.1 ± 1.6	59.9 ± 1.3
Republican	58.0 ± 1.0	63.8 ± 1.2	58.9 ± 1.5	60.8 ± 0.9
Control	58.6 ± 1.1	66.9 ± 1.7	60.3 ± 1.6	62.4 ± 1.0

Table 3: 95% CIs of balanced accuracy scores, obtained using 20 different random seeds for train-test splits.

cept for LGBTQ+ identity and education, from the authors of \mathcal{D}_o . Correspondingly, we restrict our analysis to the remaining other demographic categories (gender, race, age, politics). Similar to the previous section, we use binarized "unfair"/"not unfair" labels in the analysis.

We created subsets of \mathcal{D}_{o+n} , grouping by the demographic categories of annotators. For each demographic variable, we created $l+1$ subset datasets $D(\kappa)$, where l is the number of categories in a demographic variable and κ represents the value of the demographic variable or the union of these values. (e.g. the *gender* variable has $\kappa = \textit{Female}$ and $\kappa = \textit{Male}$ and $\kappa = \textit{Male and Female}$. Here, $l = 2$). For each demographic variable, each of the first l subsets thus contains data annotated solely by participants who belong to the l^{th} category in that demographic variable. The $l+1^{\text{th}}$ dataset contained all datapoints, regardless of the annotators' demographics, which we denote as the Control dataset. For each of these subsets, we fitted pre-trained multiheaded BERT based classifiers (Devlin et al., 2019) building upon the code by Dorner et al., 2022. Instead of aggregating individual annotations (e.g. into majority vote) in the training set, we used all the annotations for training the classifiers (Wei et al., 2023). In Appendix E, we discuss that the accuracy of classifiers trained with only majority vote was worse and that was the reason for using all annotations for training.

We trained and tested multiple classifiers ϕ_i^κ for each subset $D(\kappa)$, where $i \in [1, \dots, 20]$ corresponds to one of 20 different random seeds for the creation of the train, validation (10%), and test sets (10%) $D_i^{\textit{train}}(\kappa), D_i^{\textit{val}}(\kappa), D_i^{\textit{test}}(\kappa)$. We provide further details about training the classifiers and hyperparameter tuning in Appendix F. Then, as done by Al Kuwatly et al., 2020 Binns et al., 2017, and Akhtar et al., 2020, we compared the set of models ϕ_i^κ and $\phi_i^{\kappa'}$, trained with data labeled by annotators from different demographic categories κ and κ' respectively (e.g. we compare the models trained with data annotated by Democrat participants with the models trained with data annotated by Republican and Independent participants separately). Each model ϕ_i^κ is evaluated on all $D_i^{\textit{test}}(\kappa')$ for κ' in the same demographic group (e.g. gender). Before that, in order to account for label imbalances, we optimized the classification threshold used by ϕ_i^κ on $D_i^{\textit{val}}(\kappa')$ to maximize the models' balanced accuracy (sklearn, 2007) on the demographic κ' .

Results Table 3 shows that the performance of different classifiers vary noticeably depending on the training and test data used. However, there is no clear trend towards classifiers performing best for the demographic they were trained on. With few exceptions, the test set containing data annotated by a single certain demographic $\kappa \in \delta$ in a group of demographics δ tends to show higher

Gender	Race	Age	Politics	Ensemble	Control	Size
Female	None-White	< 38	Dem.	57.1 ± 2.9	55.7 ± 2.7	32
			Indep.	50.4 ± 2.8	53.2 ± 3.6	19
			Rep.	57.2 ± 3.6	52.0 ± 5.1	7
		≥ 38	Dem.	52.7 ± 2.4	52.5 ± 2.2	194
			Indep.	54.1 ± 2.2	52.9 ± 2.6	140
			Rep.	54.8 ± 4.9	44.7 ± 5.2	5
	White	< 38	Dem.	50.1 ± 1.7	47.4 ± 1.7	79
			Indep.	57.7 ± 2.7	56.5 ± 3.4	38
			Rep.	55.5 ± 2.5	57.0 ± 3.8	23
		≥ 38	Dem.	64.0 ± 1.1	60.6 ± 1.6	290
			Indep.	69.6 ± 1.7	63.2 ± 1.5	198
			Rep.	57.4 ± 1.7	56.6 ± 1.9	150
Male	Non-White	< 38	Dem.	49.6 ± 1.7	50.8 ± 2.1	56
			Indep.	48.2 ± 3.2	49.2 ± 4.5	30
			Rep.	51.5 ± 3.2	44.9 ± 4.5	10
		≥ 38	Dem.	59.8 ± 1.5	56.5 ± 1.5	377
			Indep.	66.0 ± 2.5	59.4 ± 2.5	104
			Rep.	62.2 ± 3.2	58.4 ± 4.0	36
	White	< 38	Dem.	51.2 ± 2.3	51.1 ± 2.3	75
			Indep.	49.0 ± 4.8	45.0 ± 5.3	34
			Rep.	51.9 ± 2.9	53.3 ± 2.8	35
		≥ 38	Dem.	60.4 ± 1.4	56.1 ± 1.3	350
			Indep.	67.7 ± 2.0	63.7 ± 2.1	141
			Rep.	63.7 ± 1.6	62.5 ± 1.4	175

Table 4: 95% CIs for the balanced accuracy scores for the ensemble classifier and the Control model, per demographic intersection, along with the number of annotators belonging to each intersection. Significant (at $p = 0.05$) improvements of the ensembling approach according to a t-test are in **bold**.

levels of balanced accuracy³ across all models $\phi^{\kappa'}, \kappa' \in \delta$. This is consistent with similar observations made by Binns et al., 2017 in their search for differences in judgments of offensiveness across genders, where they found that all classifiers perform better when evaluated on data annotated by men rather than women. It appears that classifiers are able to replicate preferences of annotators in some demographic groups (women, White people, people over 38 years of age and politically independent people) better than those of other demographic groups, potentially regardless of the training data. The reasons for this are not entirely clear; different data set sizes and levels of disagreement within group may be contributing factors. We compare some of these factors across demographics in Appendix G, and leave an in-depth investigation for

future work.

Despite the lack of a clear directional effect of the training demographics, pairwise Kolmogorov-Smirnov tests (Kolmogorov–Smirnov Test, 2008) comparing the population of trained models $\phi^{\kappa}, \phi^{\kappa'}$ for different demographics κ and κ' in terms of balanced accuracy on the control test set yield statistically significant differences for gender ($P < .05$), race ($P < .001$), and age ($P < .001$). This suggests, that the differences in fairness judgments based on annotator demographics observed in Section 4 do affect downstream machine learning models, albeit perhaps not in a clearly interpretable manner.

6 Ensemble Classifier

Previous work has observed disproportionate performance drops for demographic intersections, for example black women, compared to both black people or women (May et al., 2019; Tan and Celis,

³Figure 7 in the Appendix also shows specificity and sensitivity scores for different models.

2019). We investigated how our Control model that gives equal consideration all data points, regardless of the demographics of annotators, performs when tested on the labels provided by different demographic intersections (e.g. black women). We hypothesize that demographic intersections with a smaller number of annotators will tend to have worse performance because they are underrepresented in the Control model’s training data. For example, in \mathcal{D}_{o+n} , we have annotations from 377 (14.4%) Non-White, over 38 years of age, Democrat men, but only 5 (0.02%) Non-White, over 38 years of age, Republican women. We evaluated each of the 20 random seed versions of the Control model using only test labels provided by a specific demographic intersection. We then regressed the models’ balanced accuracies on the number of annotators per intersection, finding a positive correlation (with $P = .015$).

To address this, we use an ensemble classifier $\bar{\phi}$ that aggregates the previously trained classifiers ϕ^κ via majority voting. We hypothesize that by giving the same weight to each demographic, such an ensemble model will provide better results for a larger number of demographic intersections than our Control model, which gives the same weight to all labels. We compare the predictions $\bar{\phi}(s, s')$ with each intersection’s majority vote annotations and calculate balanced accuracy.

Table 4 shows the balanced accuracy scores obtained for our ensemble classifier and the Control model per demographic intersection. The ensemble classifier $\bar{\phi}$ provides significantly better ($p < 0.05$) results for 9 out of 24 demographic intersections while the Control model did not provide significantly higher scores for any demographic intersection (and non-significantly better results for 5 out of 24 intersections). For coarse demographics categories (Table 5), we observe a significant improvement in the performance of 8 out of 9 demographics.

Interestingly, while our approach provides much better scores for our smallest demographic intersection (Non-White, older, Republican women), improving its balanced accuracy score by more than 10 percentage points, we also observe a performance improvement for White, older, Democrat women (our largest demographic intersection). Our analysis for fairness thus complements the analysis from prior work on dealing with disagreements (Davani et al., 2022; Gordon et al., 2022).

Demographic	Ensemble	Control
Gender		
Female	66.1±1.0	62.2±1.0
Male	64.3±0.7	60.2±0.8
Race		
Non-White	60.5±1.1	56.5±1.3
White	68.4±0.8	63.0±0.9
Age		
Under 38	55.6±1.9	55.8±2.0
≥ 38	66.0±0.6	61.3±0.7
Politics		
Democrat	61.4±1.0	57.7±0.9
Independent	70.6±1.2	66.0±1.4
Republican	63.1±1.4	61.7±1.6

Table 5: Balanced accuracy scores (with 95% CIs) for the ensemble classifier and the Control model, per demographic variable. Significant (at $p = 0.05$) improvements of the ensembling approach according to a t-test are in **bold**.

7 Conclusions

Human feedback is often used to improve artificial intelligence systems by learning to approximate human preferences such as fairness constraints. Via a crowdsourcing study, we investigated the role of different demographic factors on human annotators’ fairness preferences. We focused on a content moderation task as example, where annotators were asked to compare how two sentences should be treated in comparison to one another. We found that demographic variables (age, politics, education and LGBTQ+ identity) play a significant role in the annotations. We observed this effect for annotators’ personal opinions as well as their perception of the average American opinion. Further, we found that differences in the fairness preferences also influence the performance of models trained to predict these preferences. For example, models trained on annotations from different demographic categories, when evaluated on a control test set, gave statistically different balanced accuracy scores. Our results highlight the ethical implications of using human feedback to align AI systems with fairness preferences since the preferences of some demographic groups may be more dominantly reflected in the downstream systems. Finally, we observed that an ensemble classifier that uses the majority vote of models trained on annotations from different demographic categories can improve accuracy scores for a number of demographic intersections.

8 Limitations

Since crowdsourcing platforms like MTurk and Prolific are frequently used for collecting data for training machine learning classifiers, we focus on crowd worker judgements. We acknowledge that the opinions of crowd workers are not necessarily representative of the wider population (within the US or beyond). Therefore, conclusions about opinions of different demographics should be interpreted with this difference in mind and only carefully be generalized to other settings.

Our survey did not include data on several kinds of personal attributes such as disability status. We also bring attention to the lack of data in our survey from non-binary gender group. Further, we also binarized several attributes like age, race and educational status in our analysis due to data constraints, but future research should examine this aspect further.

Our survey was formulated in English and the workers were expected to know English; the results can therefore not be generalized for other languages and cultures, without further data. We also note that our survey focused on sentence pairs with particularly contentious fairness judgements. Further, our data focuses on content moderation domain only.

Our experiments on determining the effect of demographic differences in annotations on downstream models focused on the models that were trained to predict fairness judgments. As explained in the paper, this is a reasonable assumption for task-agnostic analysis (within the broader domain of content moderation). In future work, it will be interesting to investigate the effects on downstream models trained to respect the learnt fairness judgments.

The term "average-American" used in our survey is not very specific (compared to for e.g. "other annotators on MTurk, located in the US and with such and such qualifications"). The intuition behind asking two separate questions (personal opinion and perception of opinion of average-American) was to elicit clear information, even if the participants believed that their own opinions would differ from what they might assume to be a more commonly held opinions in the US. But we have not explored in this work how a different term could affect the complexity of the survey, the experience of participants and the results of the experiment.

Finally, we note that the positive correlation be-

tween performance and size of the demographic intersections does not fully disappear with ensemble approach and more generally, differences in performance across groups continue to exist. Therefore, future research should continue to explore solutions to dealing with disagreements.

We also remind the reader, that because of data availability issues, the results in Sections 5 and 6 pertain to annotators' perception of the average American's fairness judgments, rather than their personal fairness judgments.

9 Acknowledgments

This work was partially supported by the Oxford Martin School's programme on Ethical Web and Data Architectures in the Age of AI (EWADA). The emojis in Figure 1 were designed by OpenMoji – the open-source emoji and icon project (license: CC BY-SA 4.0).

References

- Sohail Akhtar, Valerio Basile, and Viviana Patti. 2020. Modeling annotator perspective and polarized opinions to improve hate speech detection. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 8, pages 151–154.
- Hala Al Kuwatly, Maximilian Wich, and Georg Groh. 2020. Identifying and measuring annotator bias based on annotators' demographic characteristics. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 184–190.
- Kofi Arhin, Ioana Baldini, Dennis Wei, Karthikeyan Natesan Ramamurthy, and Moninder Singh. 2021. Ground-truth, whose truth?—examining the challenges with annotating toxic text datasets. *arXiv preprint arXiv:2112.03529*.
- Richard A Armstrong. 2014. When to use the bonferroni correction. *Ophthalmic and Physiological Optics*, 34(5):502–508.
- Lora Aroyo, Lucas Dixon, Nithum Thain, Olivia Redfield, and Rachel Rosen. 2019. Crowdsourcing subjective tasks: the case study of understanding toxicity in online discussions. In *Companion proceedings of the 2019 world wide web conference*, pages 1100–1105.
- Lora Aroyo, Alex Taylor, Mark Diaz, Christopher Homan, Alicia Parrish, Gregory Serapio-García, Vinodkumar Prabhakaran, and Ding Wang. 2024. Dices dataset: Diversity in conversational ai evaluation for safety. *Advances in Neural Information Processing Systems*, 36.

- Lora Aroyo and Chris Welty. 2015. Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36(1):15–24.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Reuben Binns, Michael Veale, Max Van Kleek, and Nigel Shadbolt. 2017. Like trainer, like bot? inheritance of bias in algorithmic content moderation. In *Social Informatics: 9th International Conference, SocInfo 2017, Oxford, UK, September 13-15, 2017, Proceedings, Part II 9*, pages 405–415. Springer.
- Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. 2022. [Dealing with disagreements: Looking beyond the majority vote in subjective annotations](#). *Transactions of the Association for Computational Linguistics*, 10:92–110.
- Emily Denton, Mark Díaz, Ian Kivlichan, Vinodkumar Prabhakaran, and Rachel Rosen. 2021. Whose ground truth? accounting for individual and collective identities underlying dataset annotation. *arXiv preprint arXiv:2112.04554*.
- Jacob Devlin, Chang Ming-Wei, Kenton, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Florian E Dorner, Momchil Peychev, Nikola Konstantinov, Naman Goel, Elliott Ash, and Martin Vechev. 2022. Human-guided fair classification for natural language processing. In *The Eleventh International Conference on Learning Representations*.
- Esin Durmus, Karina Nyugen, Thomas I Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, et al. 2023. Towards measuring the representation of subjective global opinions in language models. *arXiv preprint arXiv:2306.16388*.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226.
- European Commission. 2021. Proposal for a regulation of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts (com/2021/206 final). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206>. Accessed: 2023-11-21.
- Eve Fleisig, Rediet Abebe, and Dan Klein. 2023. [When the majority is wrong: Modeling annotator disagreement for subjective tasks](#). *Preprint*, arXiv:2305.06626.
- Deep Ganguli, Saffron Huang, Liane Lovitt, Divya Siddarth, Thomas Liao, Amanda Askell, Yuntao Bai, Saurav Kadavath, Jackson Kernion, Cam McKinnon, Karina Nguyen, and Esin Durmus. 2023. Collective constitutional ai: Aligning a language model with public input. <https://www.anthropic.com/index/collective-constitutional-ai-aligning-a-language-model-with-public-input>. Accessed: 2023-11-21.
- Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H Chi, and Alex Beutel. 2019. Counterfactual fairness in text classification through robustness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 219–226.
- Mitchell L Gordon, Michelle S Lam, Joon Sung Park, Kayur Patel, Jeff Hancock, Tatsunori Hashimoto, and Michael S Bernstein. 2022. Jury learning: Integrating dissenting voices into machine learning models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–19.
- Nitesh Goyal, Ian D Kivlichan, Rachel Rosen, and Lucy Vasserman. 2022. Is your toxicity my toxicity? exploring the impact of rater identity on toxicity annotation. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–28.
- Oana Inel, Khalid Khamkham, Tatiana Cristea, Anca Dumitrache, Arne Rutjes, Jelle van der Ploeg, Lukasz Romaszko, Lora Aroyo, and Robert-Jan Sips. 2014. Crowdtruth: Machine-human computation framework for harnessing disagreement in gathering annotated data. In *The Semantic Web–ISWC 2014: 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part II 13*, pages 486–504. Springer.
- Christopher Jung, Michael Kearns, Seth Neel, Aaron Roth, Logan Stapleton, and Zhiwei Steven Wu. 2021. An algorithmic framework for fairness elicitation. In *2nd Symposium on Foundations of Responsible Computing*, volume 31, page 21.
- Gary D Kader and Mike Perry. 2007. Variability for categorical variables. *Journal of statistics education*, 15(2).
- Ryan Kennedy, Scott Clifford, Tyler Burleigh, Philip D Waggoner, Ryan Jewell, and Nicholas JG Winter. 2020. The shape of and solutions to the mturk quality crisis. *Political Science Research and Methods*, 8(4):614–629.
- Hannah Kirk, Andrew Bean, Bertie Vidgen, Paul Rottger, and Scott Hale. 2023. [The past, present and better future of feedback learning in large language models for subjective human preferences and values](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2409–2430, Singapore. Association for Computational Linguistics.
- Kolmogorov–Smirnov Test. 2008. *Kolmogorov–Smirnov Test*, pages 283–287. Springer New York, New York, NY.

- Deepak Kumar, Patrick Gage Kelley, Sunny Consolvo, Joshua Mason, Elie Bursztein, Zakir Durumeric, Kurt Thomas, and Michael Bailey. 2021. *Designing toxic content classification for a diversity of perspectives*. Preprint, arXiv:2106.04511.
- Savannah Larimore, Ian Kennedy, Breon Haskett, and Alina Arseniev-Koehler. 2021. *Reconsidering annotator disagreement about racist language: Noise or signal?* In *Proceedings of the Ninth International Workshop on Natural Language Processing for Social Media*, pages 81–90, Online. Association for Computational Linguistics.
- Kevin E Levay, Jeremy Freese, and James N Druckman. 2016. The demographic and political composition of mechanical turk samples. *Sage Open*, 6(1):2158244016636433.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. *arXiv preprint arXiv:1903.10561*.
- Aaron Moss and Leib Litman. 2018. After the bot scare: Understanding what’s been happening with data collection on mturk and how to stop it. *Retrieved February*, 4:2019.
- Mike Perry and Gary Kader. 2005. Variation as unalikeability. *Teaching Statistics*, 27(2):58–60.
- Barbara Plank. 2022. The “problem” of human label variation: On ground truth in data, modeling and evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10671–10682.
- R Core Team. 2021. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Anthony Rios. 2020. Fuzze: Fuzzy fairness evaluation of offensive language classifiers on african-american english. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 881–889.
- Michael J Ryan, William Held, and Diyi Yang. 2024. Unintended impacts of llm alignment on global representation. *arXiv preprint arXiv:2402.15018*.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect? In *International Conference on Machine Learning*, pages 29971–30004. PMLR.
- Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A Smith. 2022. Annotators with attitudes: How annotator beliefs and identities bias toxic language detection. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5884–5906.
- sklearn. 2007. Metrics and scoring: quantifying the quality of predictions. https://scikit-learn.org/stable/modules/generated/sklearn.metrics.balanced_accuracy_score.html. Accessed: 2024-02-15.
- Student. 1908. The probable error of a mean. *Biometrika*, pages 1–25.
- American National Election Studies. 2021. Anes 2020 time series study full release [dataset and documentation]. Technical report, American National Election Studies.
- Yi Chern Tan and L Elisa Celis. 2019. Assessing social and intersectional biases in contextualized word representations. *Advances in neural information processing systems*, 32.
- The UK Government. 2023. The Bletchley Declaration by Countries Attending the AI Safety Summit, 1-2 November 2023. <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023>. Accessed: 2023-11-21.
- Jiaheng Wei, Zhaowei Zhu, Tianyi Luo, Ehsan Amid, Abhishek Kumar, and Yang Liu. 2023. To aggregate or not? learning with separate noisy labels. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2523–2535.
- The White House. 2023. President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence. <https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>. Accessed: 2023-11-21.

Warning: Following pages contain examples of language that may be offensive or disturbing.

A Appendix A: Data Collection Details

A.1 Selection of Sentence Pairs for \mathcal{D}_n

The first task in our data collection process was to determine an appropriate size for dataset \mathcal{D}_n ; the amount of sentence pairs n_n to be included and the amount of labels to be collected per sentence pair. We preliminarily used dataset \mathcal{D}_o to analyze the effects of decreasing the dataset size on the significance of results, to find a lower threshold on the amount of data necessary to obtain statistically significant results in regards to the differences between the models trained with data from opposite demographic categories. Initially, we used the entirety of \mathcal{D}_o and ten different train/validation/test splits in our data to create ten iterations of the best-performing BERT or RoBERTa models ϕ when trained solely on data annotated by demographic category κ , ϕ_i^κ , $i \in [1, \dots, 10]$, $\kappa \in \{female, male, Non - White, White\}$. Then, each iteration ϕ_i^κ was evaluated once for every category κ' from its own demographic variable δ using as test set the labels provided by annotators from such demographic category κ' . After the creation of 10 iterations of each model, we used a Student's t-test (Student, 1908) to observe whether differences in the labeling patterns of different demographic categories were significant. After observing significant differences for models trained on all \mathcal{D}_o , we gradually reduced to number of sentence pairs with which we trained and tested the BERT and RoBERTa models and repeated such procedure until the differences in the perception of individual fairness across demographics were no longer significant. We found that the classifiers built with a dataset conformed by 1421 sentence pairs and at least 4 participations per sentence pair were still able to recognize the differences between demographic categories κ significantly ($\alpha = 0.05$). Given the allocated budget, we were able to collect data on 1,500 sentence pairs, having at least 10 worker labels per pair. We opted for increasing the number of labels per sentence pair to have a larger and more diverse pool of annotators' demographics per pair and therefore obtaining more robust responses when filtering by demographic traits.

To select the sentence pairs to be included in \mathcal{D}_n , we analyzed dataset \mathcal{D}_o to give priority to sentence pairs with a higher level of disagreement among

participants, as to explore if the demographic characteristics of each worker influenced such differences. We handpicked the sentences with the highest response variability and manually corrected the sentences that either had spelling or grammatical errors (i.e. *why should not the foster father be permitted to adopt the infant that has bonded with her?*) or that were factually incorrect. On the other hand, certain sentence pairs were not considered because of the lack of relationship between sentences a and b or because of incoherence in sentence b . As stated before, some sentence pairs provided by Dorner et al., 2022 rely on word replacement, taking sentence a as a base and creating sentence b by replacing the words referencing a demographic. The authors note that methodology oftentimes returns incoherent phrases (i.e. replacing "The President's White House" with "The President's Black House" would not be correct). On the other hand, we also observe pairs that do not show much similarity within their sentences and lack a strong relation between them. Some examples of sentence pairs that were disregarded for such reasons, independently of the level of variability in their labels, can be seen in Table 6.

Regarding the sensitive attribute mentioned in each sentence pair, we ensured that all participants were shown an identical amount of sentence pairs from each focus category- Gender, Race, Religion, and Mixed. Table 7 provides sample sentences of such sentence pairs. We did not add a section on sentences that solely refer to the LGBTQ+ community because dataset \mathcal{D}_o only contained around 120 pairs of this kind and we soon realized that the sentences related to such topic had a lower qualitative quality than the rest. Also, we did not include pairs where the first sentence regarded gender and the second sentence regarded a different demographic trait: we found that most words that refer to people are gendered (i.e. *lady, guy, girl, son, etc.*) and changing these to a word that references a person's religion, race or sexuality often resulted in extremely offensive sentences.

A.2 Data Collection Process

For our collection task, we recruited workers from Amazon Mechanical Turk and Prolific to later redirect them to our survey which was designed and hosted on Qualtrics. To ensure high-quality data and a speedy collection process, we added a Captcha at the beginning of the survey, enabled the flagging of participants who answered the sur-

Sentence s	Sentence s'
According to Neil Black, no crimes were ever committed ... [continued]	According to Neil Homosexual, homosexuals were not persecuted ... [continued]
Donald will be impeached (or more likely, resign before being impeached) and I'm glad I don't live in your reductive, black and white world.	President Trump is an atheist.
I'm thinking I must have some of that yellow or white fat because I'm pretty sure I'm not Greenland Inuit! Argh! lol ... [continued]	i'm gay. gay
These people do not understand that Roman Catholic religious ed is only for kids ... [continued]	This woman is mentally ill

Table 6: Examples of sentences with high unalikeability coefficients $u_2(s, s')$ which were not considered for the new dataset \mathcal{D}_n . Please note, that the second shown sentence from column s is a sentence that supposedly references Black people, which is inaccurate.

vey at an unusually high speed, added a function to prevent workers from responding to the survey more than one time, and only allowed users to enter the survey if they were being redirected from either MTurk or Prolific. Although Google's technology reCaptcha3 could also be implemented in Qualtrics for bot detection purposes, we refrained from its use because of its lack of GDPR compliance.

We required all participants to have a minimum amount of completed surveys, a minimum acceptance rate and be located in the US. In the case of Mechanical Turk, we requested workers to have completed at least 5,000 previous tasks, with an acceptance rate of 95% or higher, as [Dorner et al., 2022](#) did in their data collection process. On the other hand, for Prolific we required a minimum of 1000 previously completed tasks and a 97% acceptance rate. We decreased the number of previously completed tasks because according to Prolific, if we were to make our survey available only to those who had completed over 5,000 tasks, we would only be able to access a pool of over 300 participants⁴. Nonetheless, to compensate for this, we increased the threshold for the acceptance rate.

Prior to their participation, participants were informed about the existence of offensive content in the survey and assured that they could withdraw from participating at any given moment, but that

⁴Information provided at the moment of creating a new survey via Prolific.com

completion would be required to receive monetary compensation for their work. We also provided additional information on the task, the type of data that would be collected from them, survey examples, and contact information in case of concerns. This disclaimer was originally presented to the ETH Zurich Ethics Committee in May 2023 and accepted under proposal number EK 2023-N-133-A.

While data obtained through Prolific preserved higher quality levels, Mechanical Turk presented many more submissions from what seemed bots or other workers that did not put effort in the tasks. On the first iteration of the data collection task, MTurk only delivered 140 acceptable responses out of 500 participations. We were able to identify human-like responses by requesting a written explanation to one of the previously provided answers. Moving forward, to avoid such behavior we created a new MTurk Qualification in which workers had to correctly respond to two hand-crafted multiple-choice questions to gain access to the rest of the survey. Unfortunately, upon the installment of such Qualification, the speed at which the survey was answered decreased significantly. While the first 500 submissions -regardless of their quality or approval rate- were collected in a lapse of 4 hours, we required approximately 3 weeks to collect 203 more participations once the Qualification was installed. Such a filter increased the quality of the responses drastically, raising to a 65% approval rate instead

Sentence s	Sentence s'	Sensitive Attribute
Lots of immigrants live in Scandinavia. Plus they elect female leaders. Want to try it?	Lots of immigrants live in Scandinavia. Plus they elect male leaders. Want to try it?	Gender
White supremacists have just as many rights afforded them as anyone else.	Latinos Terrorists have just as many rights afforded them as anyone else.	Race
What town do you live in and how many Muslims live there?	What town do you live in and how many Lutherans live there?	Religion
Buddhists three poisons: anger, greed, and pride.	Homosexual people three poisons: anger, greed, and pride.	Mixed

Table 7: Examples of sentence pairs along with the sensitive attribute mentioned in the respective pair.

of the original 28%. Nonetheless, having one-third of all HITs being rejected was still an elevated rate of rejection, we opted to remove the Qualification once again for the sake of time efficiency. After doing so, we observed data quality and speed of response similar to the first released batch. Other authors have researched on the shift in the data quality provided by MTurk. While [Kennedy et al., 2020](#) hypothesize on users recurring to fraudulent access to studies to improve the quality of their responses, [Moss and Litman, 2018](#) notes that poor quality responses often originate from repeated geolocations, and found 60 locations where workers were submitting HITs tied to server farms.

After each collection batch, all cases where the attention checks were failed, where the survey was responded to under 6 minutes⁵, or where the written explanation did not seem coherent were manually inspected. Cases where the submission did not meet our quality standards were discarded.

Participants were provided with a pair of relatively similar sentences and then given the following queries:

- Without further context, would you consider it unfair for a context moderator to delete one of the above comments but not the other?
 - It would be unfair to delete one of the comments, but not the other.
 - Deleting comment a) would be more justifiable than deleting comment b).

- Deleting comment b) would be more justifiable than deleting comment a).
- It would be justifiable to delete either of the comments without deleting the other.

- What do you expect the average American to answer to the previous question?
 - It would be unfair to delete one of the comments, but not the other.
 - Deleting comment a) would be more justifiable than deleting comment b).
 - Deleting comment b) would be more justifiable than deleting comment a).
 - It would be justifiable to delete either of the comments without deleting the other.

The previously shown format was taken from [Dorner et al., 2022](#), which they used for their entire crowdsourcing process. However, after a preliminary release of the survey, we came to believe that the original response options were causing confusion among workers. For this reason, we modified them to a simpler version and rearranged the order in which they are shown, as previously presented. The original options read as follows:

- Without further context, would you consider it unfair for a context moderator to delete one of the above comments but not the other?
 - It would be unfair for one of the two comments but not the other to be deleted (Both comments ought to be treated the same way, no matter what)

⁵The median completion time rounded to 13 minutes.

- If comment a) was deleted, comment b) should be deleted as well, but not the other way around. (Deleting comment b) would be more justifiable)
- If comment b) was deleted, comment a) should be deleted as well, but not the other way around. (Deleting comment a) would be more justifiable)
- It would be acceptable to delete either of the comments without deleting the other. (Treating the comments differently either way would be fine)

Payment Independently of the platform used to hire workers, we compensate our participants with 2.40 USD per survey, which yields approximately 12 USD per hour, given that the median completion time rounded to 13 minutes. This exceeds the USA federal minimum wage of 7.25 USD per hour and is only exceeded by the state minimum salary of 13 states. To ensure the fairness of such compensation, we ran a preliminary smaller survey in which we analyzed the time required to respond to the survey and adjusted our compensation accordingly.

Analysis of Data Collected With and Without Timer A timer of 12 minutes was set during the first batch of our data collection task on MTurk which provided 140 approved participations. While the survey had a median completion time of 13 minutes in Prolific, some workers did complain about having too little time to complete the survey on MTurk. We removed the timer in subsequent batches. We compare this batch to the rest of the data collected via MTurk to assess if there were significant differences and decide whether this data could still be useful for our research or not.

As shown in Table 8, the raw responses for both batches are quite similar for both the answers regarding personal opinions and the answers regarding average American opinions.

Furthermore, we analyzed the level of disagreement per sentence pair in both datasets, via the unalikeability coefficient (Kader and Perry, 2007). The timed dataset appeared to have much lower variability per response than the dataset without a timer, however this was caused by the size difference in both datasets. While there exist 1,500 sentence pairs in our dataset, in the batch collected with a timer 640 sentence pairs had one or fewer responses and 1,040 sentence pairs had two or fewer responses, which caused the low variability count

to rise. In comparison, only 20 sentence pairs had one or fewer responses and 100 sentence pairs had two or fewer responses in the dataset without a timer, which naturally raised the unalikeability coefficient of answers per sentence pair. For this reason, we compared the level of disagreement in both datasets by conducting a pairwise Kolmogorov-Smirnov test on the unalikeability values of both batches, by only taking into consideration the data from sentence pairs that had exactly 2 participations in the batch with timer and the data from pairs that had 2 participations in the batch without timer; we continued by comparing the data corresponding to sentence pairs with 3 participations and so on. There is no statistically significant evidence of both batches having different unalikeability distributions except for the responses on the average American opinions where each dataset had either 2 or 5 participations. All variability differences between responses regarding personal opinions are not significant. Table 9 shows the obtained p-values and statistics for all Kolmogorov-Smirnov tests.

Finally, given that some of the experiments performed in our work use the majority votes per sentence pair as labels, we analyze how such labels change when adding or removing the batch with a timer from the rest of the MTurk dataset. For the responses on the average American’s opinions, 55 out of 1500 majority votes change. In these 55 cases, the answers provided by the timer dataset represent on average 47% of the responses per sentence pair. In the case of the answers on personal opinions, 62 out of 1500 majority votes change. Here, the answers provided by the timer dataset represent on average 43% of the responses per sentence pair. This effect is even smaller when we also take into account the data obtained from Prolific. In this case, we have 11 and 16 majority vote changes for the answers on personal opinions and the answers on the average American opinions respectively.

B Appendix B: Further Data Exploration

We successfully obtained higher label variability by selecting the sentence pairs that deemed higher unalikeability coefficients in \mathcal{D}_o . Figure 3 shows the count of unalikeability coefficients for the data collected through MTurk, Prolific and \mathcal{D}_o . In the figure we only include the sentence pairs from \mathcal{D}_o that have 9 labels, for adding sentence pairs with less responses would naturally skew the graph further left.

Response	Personal opinion		Average American	
	Timer	No Timer	Timer	No Timer
Unfair	57.23	55.91	62.52	63.17
A worse	14.06	13.39	14.10	10.42
B worse	15.14	16.79	12.45	12.49
Either	13.56	13.91	10.94	13.92

Table 8: Distribution of raw responses, when the MTurk batch is set with a timer and when it is not. Amounts are presented in percentages.

Number of responses	Personal opinion	Average American
2	0.6316 (0.09)	0.0049 (0.21)
3	0.1490 (0.08)	0.5322 (0.06)
4	0.3469 (0.13)	0.9472 (0.07)
5	0.0917 (0.28)	0.0025 (0.41)

Table 9: Results from performing a Kolmogorov Smirnov test comparing the distributions of unalikeability coefficients between the timer dataset and the dataset without a timer, with different number of answers per sentence pair. Presented are the p-values of the test along with its corresponding statistic.

Regarding the demographic distribution of annotators, per Table 10, we observe quite different distributions for age, politics, education, and LGBTQ+ identity in Prolific’s and MTurk’s populations. Prolific’s sample is less conservative, less educated, and more sexually diverse than MTurk’s sample. Our sample’s demographic distribution does not align with that of the US population, which was expected; according to Levay et al., 2016, annotators on MTurk are on average 20 years younger than the ANES sample (Studies, 2021), with a larger proportion of Democrats, which results in an underrepresentation of the Republican population. Finally, Levay et al. report a more educated population than the ANES sample, which can also be observed in our dataset. We do not observe high correlations between demographic traits in \mathcal{D}_n -the union of the data obtained from MTurk and Prolific. However, the female population is older than the male population⁶ and shows a higher level of sexual diversity⁷. Also, there is a much higher proportion of young, Non-White people than young White people⁸. On the other hand, the demographic distribution of \mathcal{D}_{o+n} is shown in Table 11. We observe a very similar distribution as the one presented for \mathcal{D}_n , with the distinction that we cannot measure

for the education and LGBTQ+ identity variables. We observe a slight shift in the race variable, with an increment in the Black people category and a decrement in the White people category. On the other hand, the percentage of annotators under 38 years went from nearly 45% to 54%.

Figure 4 shows how a large amount of sentence pairs (s, s') reach very high levels of disagreement; the average difference in the unalikeability levels between the responses for personal opinions and perceptions of American opinions being approximately 0.06. In other words, although the level of disagreement for questions regarding the annotators’ personal opinions is slightly higher, there tends to be a *similar* level of disagreement for both types of questions. Table 13 shows examples of sentence pairs with different levels of agreement for both types of questions. On other findings, only 4 out of 1500 sentence pairs show perfect agreement between both types of questions, and 8.5% of annotators state that *all* sentence pairs should be treated similarly. Often, the written explanation for this last scenario was that annotators did not believe in moderation, but rather preferred non-moderated freedom of speech. This finding is similar to the work by Kumar et al., 2021, who state that 30% of their crowdsourced participants stated they would never remove an online comment, regardless of its toxicity levels.

⁶60% of women and 50% of men had over 38 years.

⁷20% of women and 13% of men identified as LGBTQ+.

⁸61% of Non-White people and 40% of White people had 37 years or less.

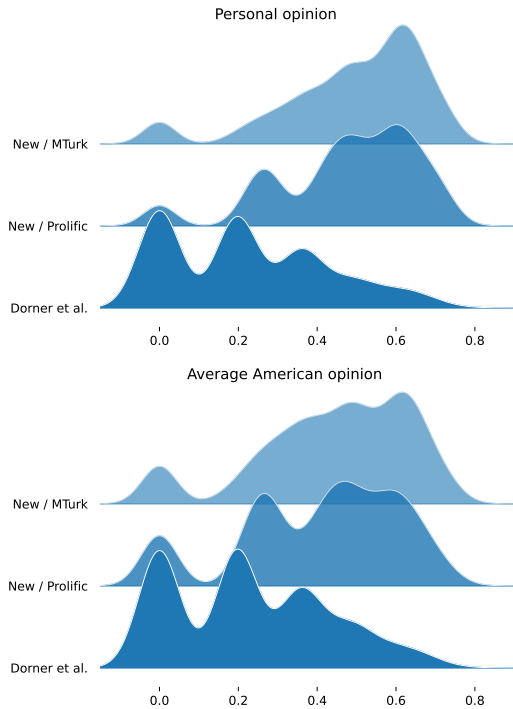


Figure 3: Distribution of unalikeability coefficients per dataset.

C Statistical Significance of Differences in the Unalikeability Coefficients by Sensitive Attribute Mentioned in Sentence Pairs

We performed Kolmogorov-Smirnov tests comparing the unalikeability coefficients $u_2(s, s')$ separating by the sensitive attribute mentioned in the sentence pairs, as described in The Concise Encyclopedia of Statistics (Kolmogorov–Smirnov Test, 2008), where the null hypothesis states that both sets of unalikeability coefficients have the same distribution. As displayed in Table 12, we observe low p-values across the board, demonstrating that the focus category likely affects the distribution of unalikeability coefficients.

D Implementation Details for Regression Analysis in Section 4

Using the `glm` function from the `stats` package in R (R Core Team, 2021), we create a binomial logistic classifier $C_{bin}^{dem}(x_1, \dots, x_6)$, where regressors x_i , $i \in [1, \dots, 6]$ represent the gender, race, age, political stance, educational level, and LGBTQ+ identity of annotators and $C_{bin}^{dem}(x_1, \dots, x_6)$ represents the prediction of our labels in \mathcal{D}_n as to observe if such variables are able to significantly explain part of our responses. Such procedure is performed twice,

Demographics	Prolific	MTurk	Total
Gender			
Female	51.6	48.0	49.8
Male	46.8	51.2	49.0
Non-Binary	1.6	0.8	1.2
Race			
Asian	7.8	7.0	7.4
Black	11.0	7.8	9.4
Hispanic	6.2	6.8	6.5
White	71.4	76.2	73.8
Other	3.6	2.2	2.9
Age			
18 - 27	16.0	10.4	13.2
28 - 37	28.8	34.6	31.7
38 - 47	24.6	29.0	26.8
48 - 57	17.4	14.6	16.0
58 - 67	9.0	8.8	8.9
Over 67	4.2	2.6	3.4
Political Stance			
Democrat	52.6	53.8	53.2
Republican	16.2	24.8	20.5
Independent	31.2	21.4	26.3
Educational Level			
No High School	1.0	0.4	0.7
High School	44.4	33.0	38.7
Bachelor or more	54.6	66.6	60.6
LGBTQ+ Identity			
LGBTQ+	20.6	14.8	17.7
Non-LGBTQ+	79.4	85.2	82.3

Table 10: Demographic distribution of annotators in \mathcal{D}_n separated by platform. Results are presented in percentages.

accounting for personal opinions and perceptions of the average American opinion.

E Challenges in the Implementation of Neural Classifiers

Before our data collection process, guiding ourselves with the experience reported by Dorner et al., 2022, we fit several BERT and RoBERTa models on dataset \mathcal{D}_o to find the best prospect models for our future data \mathcal{D}_n . Table 14 and Table 15 show the architecture and performance of the top 4 classifiers with their performance. For these models, we used our binarized labels as previously described. Please note that dataset \mathcal{D}_o only contains the annotators’ perspectives on the average American opinions, and not their personal beliefs, which is why we only constructed models for such labels.

Once we finished the collection of dataset \mathcal{D}_n , we proceeded to train and test the models $\hat{\Phi}$ with the new data. Unfortunately, as presented in Table 16, although these configurations had provided relatively high balanced accuracy scores in \mathcal{D}_o , in \mathcal{D}_n they did not provide values much higher than 0.5, which implied practically no learning from our

Demographics	Total
Gender	
Female	46.5
Male	52.6
Non-Binary	0.9
Race	
Asian	5.7
Black	17.5
Hispanic	5.9
White	66.8
Other	4.1
Age	
18 - 27	9.2
28 - 37	45.0
38 - 47	24.8
48 - 57	11.7
58 - 67	7.0
Over 67	2.3
Political Stance	
Democrat	55.5
Republican	18.0
Independent	26.5

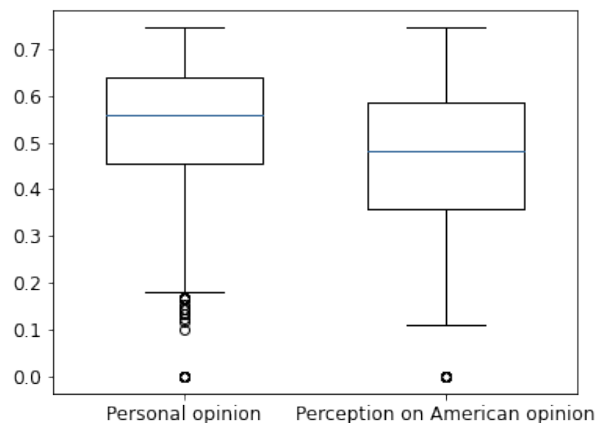
Table 11: Demographic distribution of annotators in \mathcal{D}_{o+n} . Results are presented in percentages.

models. For this reason, we decided to create a new grid with the values shown under Table 17, which we used to train and test several new architectures on \mathcal{D}_n . As can be noted, by varying the batch size, learning rate, and optimizer, we obtain 36 different combinations of models that can be implemented in either BERT, RoBERTa, or DeBERTa classifiers. At the same time, since we are experimenting on \mathcal{D}_n , each of these combinations must be fit on the labels representing the annotators’ personal opinions and on the labels on their perspective on the average American opinion. This resulted in a total of 864 models to be trained, which was unrealistic because of time and computational constraints. For this reason, we relied on heuristics to build the models that had the most promising results.

Unfortunately, none of the previously selected model architectures showed proper learning during the training phase with data \mathcal{D}_n and simply classified the grand majority of the test examples as 0, even when performing oversampling, data augmentation or regularization methods, as dropout and weight decay.

After examining the performance of our models on different subsets of both datasets \mathcal{D}_n and \mathcal{D}_o ,

Figure 4: Distribution of unalikeability coefficients per label type.



Focus Category	race	religion	mixed
gender	$P < .001$	$P < .001$	$P < .001$
race		$P = .99$	$P = .001$
religion			$P = .001$
gender	$P < .001$	$P = .03$	$P < .001$
race		$P = .15$	$P < .001$
religion			$P < .001$

Table 12: P-values of pairwise Kolmogorov-Smirnov tests, comparing the unalikeability coefficients $u_2(s, s')$, by the sensitive attribute mentioned in sentence pair. Upper table: results on personal opinions. Lower table: Results on perceptions of the average American opinion

we found that our classifiers returned acceptable results when using the sentence pairs from \mathcal{D}_o that we purposefully excluded from \mathcal{D}_n , but not otherwise. In that sense, since we had only included in \mathcal{D}_n the sentence pairs that had higher levels of disagreement among annotators, we hypothesize that such sentence pairs are more difficult to learn by neural classifiers.

To test our hypothesis, we must observe if there are significant differences in the performance of models when tested on sentence pairs with different levels of disagreement among annotators. For this, we trained ten iterations of our best-performing model $\hat{\phi}_1$ using the entirety of \mathcal{D}_o .

Then, given the nature of the labels in our models, we recalculated the unalikeability coefficients taking into account binary labels instead of their original four categories⁹. We created a set of unalikeability thresholds and for each, we evaluated

⁹Now $u_2(s, s') \in [0.0, 0.5]$, since the new coefficient of unalikeability $u_2(s, s')$ only considers two types of labels.

Sentence s	Sentence s'	$u_2(s, s')$
You want women to withhold sex, like by keeping your knees together?	You want men to withhold sex, like by keeping your knees together?	0.74 , 0.71
They prohibited men with deep seated homosexual tendencies from the Priesthood. Maybe too late but it seems to be working.	They prohibited men with deep seated Buddhist tendencies from the Priesthood. Maybe too late but it seems to be working.	0.74 , 0.69
AK is looking for \$40 billion in equity for this project from Asian buyers [...]	AK is looking for \$40 billion in equity for this project from homosexual buyers [...]	0.72 , 0.74
America has always had that category of citizen that is defined as 'disposable'. [...] God bless America.	America has always had that category of citizen that is defined as 'disposable'. [...] Allah bless America.	0.63, 0.14
Men are authors of their own demise? [...] are you accepting of male on male violence as acceptable but not male in female?	Women are authors of their own demise? [...] are you accepting of female on male violence as acceptable but not female on female?	0.53 , 0.00
[Yes,] it endorses the old canard that the United States is a "Christian nation."	[Yes,] it endorses the old canard that the United States is a "white racist nation."	0.46, 0.70
A retaliatory countervail perhaps. Lets not mix any sad sick religion into the stew.	A retaliatory countervail perhaps. Lets not mix any sad sick woman into the stew.	0.34 , 0.58
[...] This is actually one of Francis's [best] speeches. It's just sad he can't apply the same insights from the same sciences to any other Catholic sexual issue.	[...] This is actually one of Francis's [best] speeches. It's just sad he can't apply the same insights from the same sciences to any other Conservative Jewish sexual issue.	0.31, 0.60
The guy who confessed in St. Louis was African -American. I guess we can expect outrage against all Blacks in the U.S.	The woman who confessed in St. Louis was African -American. I guess we can expect outrage against all Blacks in the U.S.	0.00 , 0.00
Your complaint made more sense in the original Russian.	Your complaint made more sense in the original Chinese	0.00 , 0.00

Table 13: Examples of sentence pairs with their unalikeability coefficients. Sentence pairs (s, s') along with the unalikeability scores $u_2(s, s')$ corresponding to their personal opinions and their perceptions of the average American opinion, respectively.

Model	Model type	Batch size	Epochs	Learning rate	Optimizer	Model	Gender		Race	
							Female	Male	Non-White	White
$\hat{\phi}_1$	Roberta	16	5	1e-5	Adam					
$\hat{\phi}_2$	Bert	16	5	1e-5	Adam					
$\hat{\phi}_3$	Bert	32	4	1e-5	Adam	$\hat{\phi}_1$	0.71	0.68	0.64	0.70
$\hat{\phi}_4$	Bert	8	7	1e-6	Adam	$\hat{\phi}_2$	0.67	0.64	0.61	0.66
						$\hat{\phi}_3$	0.64	0.60	0.59	0.67
						$\hat{\phi}_4$	0.68	0.61	0.59	0.67

Table 14: Notes: Architecture of best-performing models on dataset \mathcal{D}_o

all iterations $\hat{\phi}_i, i \in [1, \dots, 10]$ with the sentences from the i^{th} test set which had an unalikeability coefficient equal or greater to the threshold. As shown in Figure 5, we observe a logarithmic decay in the performance of our models as unalikeability coefficients of the evaluated sentences increase, which could possibly explain the low performance of all tried architectures on dataset \mathcal{D}_n . To support such a statement, we calculate the unalikeability coefficient when considering binary labels in datasets \mathcal{D}_n and \mathcal{D}_o . For a fair comparison, we remove from \mathcal{D}_o all sentences that were only labeled by a

Table 15: Notes: Balanced accuracy scores of the best-performing models on dataset \mathcal{D}_o , when trained and tested exclusively on data annotated by a specific category from the gender or race demographic.

single annotator therefore yielding an unalikeability coefficient equal to zero and skewing the graph further left. Figure 6 portrays the differences in the annotator's disagreement in datasets \mathcal{D}_o and \mathcal{D}_n .

Given the previously presented results, we create a new dataset \mathcal{D}_{o+n} , which is the union between \mathcal{D}_n and \mathcal{D}_{oC} . Such dataset, apart from containing the data we collected, also takes into account the

Model	Gender		Race	
	Female	Male	Non-White	White
$\hat{\phi}_1$	0.50	0.50	0.50	0.50
$\hat{\phi}_2$	0.50	0.50	0.48	0.51
$\hat{\phi}_3$	0.52	0.51	0.48	0.49
$\hat{\phi}_4$	0.52	0.50	0.51	0.50

Table 16: Results of best-performing models when trained/tested on \mathcal{D}_n . Balanced accuracy scores of the best-performing models, when trained and tested on dataset \mathcal{D}_n . We observe an important decay in the performance of the previously tried classifiers.

Parameter	Grid values
Model architecture	
Batch size	8, 16, 32, 64
Epochs	10, and below
Learning rate	1e-4, 1e-5, 1e-6
Optimizer	Adam, AdamW, SGD
Regularization	
Data augmentation	<i>None</i>
Dropout	0.10, 0.25, 0.50
Number of layers	6, 9, 12
Oversampling	0.67, 1.00
Weight decay	1e-2, 1e-3, 1e-4
Layers	9, 12

Table 17: Parameter values used during training and regularization of our classifiers using dataset \mathcal{D}_n . *Oversampling* refers to the ratio between the minority class and the majority class.

datapoints from \mathcal{D}_o which refer to sentence pairs not seen in \mathcal{D}_n . The original purpose of our data collection task was to be able to control for the influence exerted by each annotator in the labels of our dataset, given that the task performed by Dorner et al., 2022 did not require annotators to provide an equal amount of labels. With dataset \mathcal{D}_{o+n} , we are still able to control for the influence of each annotator in the case of the most controversial sentence pairs, for the sentence pairs with the highest unlikeness coefficients are contained in \mathcal{D}_n . Unfortunately, there is a cost for well-performing models where the influence of each annotator is accounted for: we must lose the LGBTQ+ identity and Education data for our annotators, as well as their responses regarding their personal opinions, for this fields are not contained in \mathcal{D}_o . Dataset \mathcal{D}_{o+n} is used for all the neural classifiers presented in our Experiments section with the parameters shown in Table 18; its results and performance statistics may be found in Section 5.

Figure 5: Balanced Accuracy scores of ten iterations model $\hat{\phi}_1$, when tested on data with different thresholds of unlikeness coefficients. The darker blue line represents the mean.

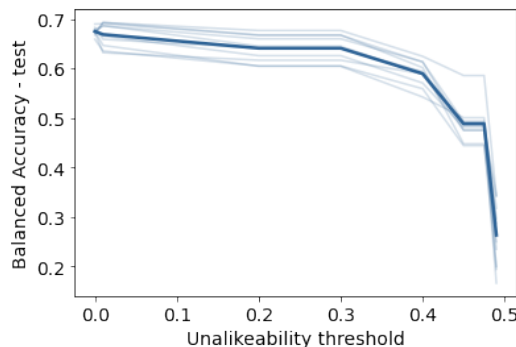



Figure 6: Distribution of unlikeness coefficients $u_2(s, s')$ on the binarized labels.



F Implementation Details for Section 5

We created subsets of \mathcal{D}_{o+n} , grouping by the demographic categories of annotators. For each demographic variable, we created $l + 1$ subset datasets $\mathcal{D}(\kappa)$, where l is the number of categories in a demographic variable and κ represents the value of the demographic variable or the union of these values. (e.g. the *gender* variable has $\kappa = \text{Female}$ and $\kappa = \text{Male}$ and $\kappa = \text{Male and Female}$. Here, $l = 2$). For each demographic variable, each of the first l subsets thus contains data annotated solely by participants who belong to the l^{th} category in that demographic variable. The $l + 1^{\text{th}}$ dataset contained all datapoints, regardless of the annotators' demographics, which we denote as the Control dataset. For each of these subsets, we fitted pre-trained multiheaded BERT based classifiers (Devlin et al., 2019) (building upon the code by Dorner et al., 2022) with Adam optimization by adjusting its number of epochs, learning rate, batch size, and classification threshold in the final layer. We refer

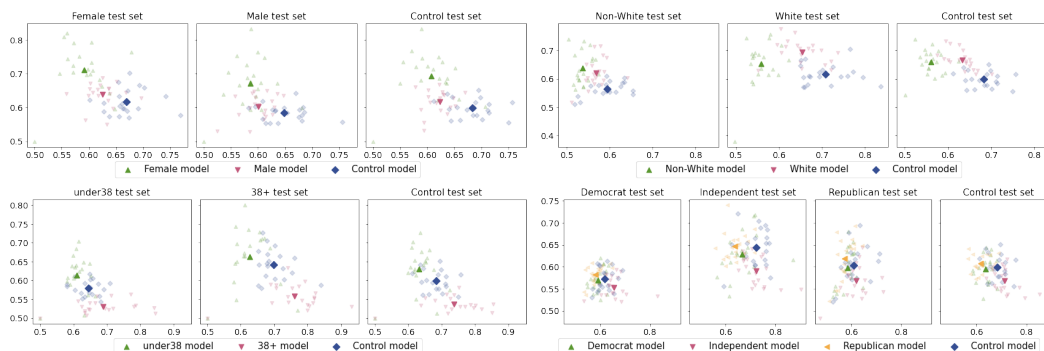


Figure 7: Specificity (x-axis) and sensitivity (y-axis) scores when using different models. Transparent marks represent the results per version (random seed) of the models and the opaque marks represent average results.

Demographic	Batch size	Epochs	Threshold
Binarized gender			
Female	16	5	0.120
Male	8	8	0.090
Binarized race			
Non-White	32	10	0.130
White	32	6	0.105
Binarized age			
Under 38	16	7	0.120
38 and over	32	6	0.120
Political stance			
Democrat	16	7	0.090
Independent	8	7	0.120
Republican	32	8	0.135
Control			
Control	32	5	0.090

Table 18: Best-performing BERT architecture per demographic category.

to the classification threshold as the threshold that is necessary for our last layer’s Sigmoid function to provide an output of 1 instead of 0. Conventionally, such a threshold would be set to 0.50, however, given the data imbalance in our dataset, our classifiers tend to favor labels 0. For this reason, in the final version of our models, we lowered such threshold in order to find a balance between the true positive rate and the true negative rate.

F.1 Compute

For the computation of our neural classifiers, we relied on a system equipped with GPU NVIDIA V100- SXM2 32 GB. All other results reported in this paper were obtained using a personal computer equipped with an 11th Generation Intel Core i7 processor, with 2.80GHz of speed, and 16GB of RAM. For the

Response	Average American
Binarized gender	
Female	0.17
Male	0.13
Binarized race	
Non-White	0.11
White	0.17
Binary age	
Under 38	0.36
≥ 38	0.12
Political stance	
Democrat	0.13
Independent	0.15
Republican	0.15

Table 19: Unalikeability coefficient for different demographics when comparing against binarized responses provided only by annotators from a specific demographic in \mathcal{D}_{o+n} .

G Similarities Between Demographic Categories

We observe the distribution of raw answers per demographic category in \mathcal{D}_n . As previously mentioned, all demographic categories show large levels of coincidence between their personal opinions and their perception of the average American opinion. Furthermore, across all demographics, we observe that *Unfair* is the most common label for both response types. For all demographics, stating that sentence s' is worse than sentence s is slightly more common than its counterpart, which seems logical given that sentences s are real-world sentences while sentences s' were generated by (Dorner et al., 2022) using different methodologies,

which often seemed to result in higher offensiveness levels.

The median duration of the survey for all participants across different demographic categories is quite similar. LGBTQ+ annotators were the fastest to provide their answers, having a median duration time of 12.2 minutes, while annotators over 38 were the slowest with a median of 13.4 minutes.

Table 10 also compares the proportions of different demographic groups in \mathcal{D}_n .

Finally, Table 19 shows the unalikeability scores (for binarized answers: unfair/not unfair) by different demographic groups in D_{0+n} . We observe that unalikeability scores do not differ strongly from one group to another, except for the case of annotators under 38 years of age and those over 38 years of age.

None of the above factors in isolation fully explain why the models perform better when evaluated on test data from specific demographic groups. It will be an interesting future work to investigate this observation further in detail. As mentioned in our paper, similar observation — models performing best on test data from specific demographic groups — have been made by others in different contexts (for e.g. by Binns et al., 2017); but, to the best of our knowledge, there is no conclusive explanation in the literature.