♪ Something Just Like TRuST ♪ *: Toxicity Recognition of Span and Target

Disclaimer: Due to the topic studied here, the paper contains offensive words.

Anonymous ACL submission

Abstract

Toxic language refers to content that is offensive, abusive, or that promotes harmful behaviors, such as self-harm or violence towards others. Effective methods to prevent or mitigate toxic content generated by large language models (LLMs) remains an open problem that depends on careful specification of toxicity, and high quality evaluation. This paper introduces TRuST, a comprehensive dataset that merges and augments existing toxicity datasets, including human or machine sources and human or machine annotation, and synthesizes existing labeling schemes. Prevalence of toxicity depends on the target, and its perception depends on one's background. Thus, TRuST includes a diverse range of target groups, from ethnicity to disability, and human annotation is evaluated for diversity of the annotator pool. When we benchmark state-of-the-art LLMs in different settings, we find that fine-tuned models consistently outperform zero-shot and few-shot prompting, though performance remains low for certain social target groups. Further, models with reasoning capabilities do not perform significantly better, suggesting that LLMs have weak social reasoning skills.

1 Introduction

Offensive and toxic content is pervasive in social media and online forums. Because large language models (LLMs) are trained on online data, they learn to generate toxic output (Gehman et al., 2020; Hartvigsen et al., 2022). Exposing users to toxic content might lead to lack of empathy, prejudice (Pluta et al., 2023), stress, and depression (Saha et al., 2019). Social groups that might need special attention, such as people with developmental disorders can be more vulnerable to toxic content (Kim et al., 2023). Furthermore, LLMs are being used in high-stakes conversational applications such as education (Yu et al., 2025; de Araujo et al., 2025) and healthcare (Yeo et al., 2024, 2025). Exposure to offensive language can lead to issues with self-esteem, anger, and anxiety (Kansok-Dusche et al., 2023). As such, the NLP community has devoted increasing attention to toxicity detection (Davidson et al., 2017; Rawat et al., 2024; Duan et al., 2025; Khurana et al., 2025) and toxicity mitigation (Suau et al., 2024; Ermis et al., 2024; Pozzobon et al., 2023a; Li et al., 2024; Liu et al., 2021). However, existing toxicity datasets differ in format, leading to lack of uniformity in benchmarking. We build on existing datasets to create TRuST to improve benchmarking, and to provide a more comprehensive picture of online use of toxic language.

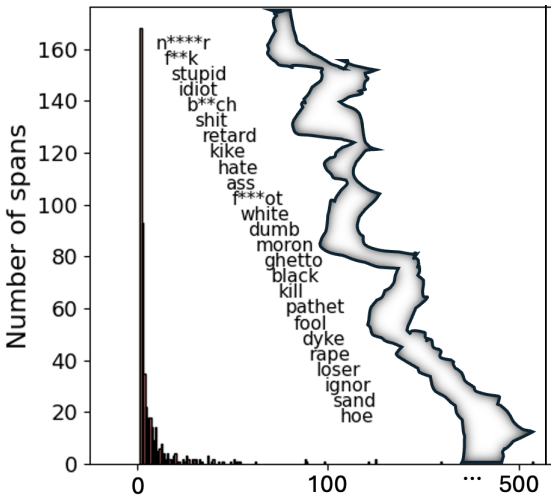

Figure 1: Histogram of distinct tokens within toxic spans (x-axis) ordered by their total count (y-axis). Note that the figure has a vertical "tear" at around 150 on the x-axis to show the long-tailed distribution. The 25 most common words are shown descending on a diagonal.

Existing toxicity datasets for English differ in size, source, and annotation scheme, ranging from about 5K examples (Zampieri et al., 2023) to four times that size (Davidson et al., 2017; Mathew et al., 2021) for human-generated posts from social media (typically Twitter), or an order of magnitude larger for AI-generated datasets (Hartvigsen et al., 2022).

*An allusion to Coldplay's "Something Just Like This."

Annotations vary from target social groups and toxicity labels (Almohaimeed et al., 2023; Mathew et al., 2021) to toxic spans (Zampieri et al., 2023; Pavlopoulos et al., 2022). The lack of a unified scheme with important components of toxicity perception impedes comparability, reproducibility, and the development of robust toxicity detection models. Hence, we review existing datasets and their annotation schemes to synthesize a framework that includes toxicity, target and span, defining 8 higher-level target groups and 20+ fine-grained subgroups for nuanced bias analysis and more accurate toxicity models based on TRuSt, our new large dataset (N=300K).[1] TRuST merges existing datasets that represent diverse social media sources as well as AI-generated examples. We train human annotators to label a subset of 11,500 examples. To ensure annotation quality, given the diversity of offensive language perception, we assess inter-rater reliability against a diverse group of volunteer annotators.

To investigate the benefits of TRuST, we carry out three activities. First, we analyze the characteristics of toxic language from online sources. For example, Fig. 1 shows that while some words occur quite commonly in toxic language, where surrounding context is required to differentiate non-toxic and toxic uses of common words such as "hate," "black," or "sand," there is a very long tail. Second, we benchmark LLMs on toxicity detection using zero- to 6-shot, including automatically generated and Chain-of-Thought prompts, and we compare LLMs to pre-trained language models (PLM). Last, we utilize the best method from our benchmarking tests to increase the size of TRuST. We find that fine-tuned PLMs perform somewhat better than LLMs. We also find that performance remains low on certain target groups, suggesting that social reasoning in LLMs needs improvement.

In sum, our three contributions are: (1) the TRuST dataset labeled for toxicity, target group, and toxic span; (2) a comprehensive definition of toxicity and an annotation scheme synthesized from previous work; (3) extensive benchmarking.

2 Related Work

2.1 Annotation of Toxicity

There is a lack of consensus about the toxicity definition and terms which leads to confusion and inconsistency (Khurana et al., 2022). Davidson et al. (2017) define a scheme for hate speech (HS) as hatred towards a social group or humiliation/insult to the members of a group including threats and violence. Nockleby (2000) defines HS as any communication that disparages a person or a group on the basis of a social characteristic such as race, color, ethnicity, gender, sexual orientation, nationality, religion, or other characteristic. Both works agree that HS is directed at a target, but inclusion of targets has not been widely adapted.

Going beyond hate speech, Dorn et al. (2024) state that toxic content includes using slurs in a pejorative way, attacking or criticizing a minority, promoting hate speech, negative stereotypes, etc. The definition used in PerspectiveAPI (Lees et al., 2022) is "rude, disrespectful or unreasonable language that is likely to make someone leave a discussion." Kumar et al. (2021) propose that toxicity involves identity-based attacks such as racism, bullying, threats of violence, sexual harassment, and includes emotional harm and hate speech.

Importantly, the toxicity is subjective, depending on factors such as context, background, or demographics. This has led to work that provides parallel annotations of toxicity from different demographic groups (Mostafazadeh Davani et al., 2024), although we find agreement across a diverse annotator pool. Moreover, some works consider all content with slurs or swear words as toxic, which may not be the case (e.g. "you are f***ing beautiful"). In assembling TRuST, we aimed for a unified definition of toxicity, and included diverse target groups. We included toxic spans, which provide context to prevent over-reliance on slur words.

2.2 Datasets

Toxicity datasets differ in size, text source, and annotation scheme, as shown in Table 1. TRuST is the largest of the listed datasets, most of which have fewer than 25k examples; in contrast to the closest competitor in size, nearly a quarter of TRuST is human-annotated. It has the greatest diversity of sources, and a competitive number of targets.

TRuST is the only toxicity dataset with all four types of labels shown in Table 1; notable, only three of the listed datasets have span labels, which supports explainability. Also, what counts as a "target" varies from work to work, such as explicit social groups versus non-specific targets such as pronouns. We think that the actual social group is important to understand the nuances of toxicity and social group biases, though only three datasets have this kind of information. Last but not least, because

[1] TRuST will be available for download upon publication.

Dataset	Size	Text Source	Tox.	High-level Target	Fine-grained Target	Span
Zampieri et al. (2023)	4.5k	Twitter	2	✗	free text	✓
Fleisig et al. (2023b)	5k	GPTs	3	2	Not Provided	✗
Almohaimeed et al. (2023)	8.3k	Twitter	4	6	31	✗
Pavlopoulos et al. (2022)	11k	Civil Comments	1	✗	✗	✓
Zampieri et al. (2019)	14k	Twitter	2	3	✗	✗
ElSherief et al. (2021)	19k	Twitter	8	✗	free text	✗
Mathew et al. (2021)	20k	Twitter & Gab	3	5	18	✓
Davidson et al. (2017)	25k	Twitter	3	✗	✗	✗
Zhou et al. (2023)	33k	GPT3.5	free text	✗	free text	✗
Sap et al. (2020)	45k	Social Media	4	✗	free text	✗
Hartvigsen et al. (2022)	274k	GPT-3	2	✗	13	✗
TRuST (ours)	298k	Social Media & AI	2	8	24	✓

Table 1: Comparison of datasets in terms of size, text source level, and annotation scheme

LLMs are ubiquitous, it is important to understand the characteristics of toxic content generated by LLMs. Previous datasets have either social media or AI-generated data, but not both. TRuST merges and cleans previous datasets to produce a larger dataset (N $\geq$ 290k) with greater diversity of source text and a broader set of annotation labels.

2.3 Prediction of Toxicity, Targets and Spans

The SOTA for toxicity detection is PerspectiveAPI (Lees et al., 2022), in which multilingual BERT (Devlin et al., 2019) is first trained, followed by distillation of multiple single-language convolutional neural networks. However, it has been criticized for having only moderate correlation with humans (Welbl et al., 2021; Schick et al., 2021), and for over-reliance on surface patterns such as swear words, resulting in a high rate of false positives (Rosenblatt et al., 2022). Another issue is the occasional release of updates without notification or explanation, making previous findings questionable (Pozzobon et al., 2023b). Other BERT-based (Devlin et al., 2019) models are also fine-tuned to predict toxicity (accuracy of 0.95) (Tsai et al., 2025; Joshi et al., 2025). Shaikh et al. (2023) show that CoT pushes LLMs to generate harmful content and increase their biases, such as stereotyping.

A common approach to predict target social group is to train a multi-layer perceptron classifier on embeddings of examples from a pretrained encoder, such as BERT (Mathew et al., 2021). LLMs have also been used (ROUGE score of 61.2) (Zhou et al., 2023).

Toxic span prediction is usually treated as a classification problem where a label is predicted for each token (He et al., 2024). SpanBERT (Joshi et al., 2020), a pretrained model for high quality embeddings for predicting spans of text, is the SOTA encoder for this task. There is little work on using LLMs for toxic span identification, apart from an investigation of GPT4 for toxic span identification in Romanian (F1 Score of 0.72) (Paraschiv et al., 2023).

We benchmark all three tasks, and also analyze additional factors such as "reasoning" or in-context learning. Our analysis highlights performance disparities across different target groups, and the need to improve LLM's social reasoning about toxicity.

3 Dataset Assembly

Based on our review above, we take toxicity to comprise three broad categories: hate speech, abusive language, and sexual harassment, as further explained below. Further, we adopt three types of annotation label: toxicity, target group, and span. TRuST draws from five datasets discussed above (cf. Table 1), augmenting those that lacked any of the three annotation categories, and reducing noise in labels, e.g., for LLM-generated data. In sum, TRuST re-annotates data from ToxicSpan (CC-BY 4.0) (Pavlopoulos et al., 2022), HateXplain (CC-BY 4.0) (Mathew et al., 2021), TBO (CC-BY 4.0) (Zampieri et al., 2023), Thos (CC-BY 4.0) (Almohaimeed et al., 2023), and ToxiGen (CC-BY 4.0) (Hartvigsen et al., 2022). The ToxiGen toxicity labels were particularly noisy, due to assigning the prompt labels to the generated text, which has been shown to be a too strong assumption (e.g., (Wei et al., 2025)). The full TRuST dataset has nearly 300K examples (see Table 1 and Appendix A).

Hate speech is defined as offensive and discriminatory discourse towards a group or an individual based on characteristics such as race or religion, thus always has a target. It includes ***negative stereotyping*** (negative traits and characteristics attributed to a group), ***racism*** (discrimination against or negative attitudes towards individuals/groups belonging to a certain race), ***sexist language*** (discrimi-

nation or behaviors that foster stereotypes based on a gender), and discrimination based on sexual orientation. **Abusive language** is content with inappropriate words such as profanity or disrespectful terms for people based on sociodemographic characteristics. It includes ***psychological threats*** (expressions of an intent of harm such as humiliation, causing distress, constant criticism). Our last category is **sexual harrasment** which includes unwelcome sexual moves, requests of sexual favors, or other unwanted physical/verbal behaviors of a sexual nature towards someone. In our work, toxic language often has a target, but can also involve use of offensive words in an aggressive fashion without targeting a specific social group, e.g., "honestly? I can handle kpop stans dragging armys but just stay the f**k away from bts they've done lit rally nothing to y'all." Our annotation instructions (see Appendix C) include a binary label for toxicity defined in terms of these three categories.

3.1 Human Annotation Procedure

To provide a high-quality toxicity evaluation benchmark, we collected human annotations for binary toxicity, target social group, and toxic spans for a subset of over 10k examples, including a smaller subset (N=300) where we examined the impact of a larger, high-diversity annotator pool on interannotator reliability. We defined 24 target groups, including "no target." Note that we separated "Chinese" from "Asian" because it included a significantly large number of examples.

We hired six undergraduate students studying computer science or data science experienced in data analysis to perform the manual annotation; they were paid $10/hour. We recruited students from different races (Indian, Chinese, Caucasian), given the cultural and demographic differences in toxicity perception (Mostafazadeh Davani et al., 2024; Fleisig et al., 2023a). They were provided with detailed instructions emphasizing the importance of considering the whole context rather than words alone (see Appendix C). To ensure high quality, we did three iterations of annotator training; by the time we collected toxic span annotation, which was done last, only three of the original six annotators were still available. Our inter-annotator agreement scores using Krippendorf's α (Krippendorff, 2013) are 0.56, 0.66, and 0.55 for toxicity, target social group and spans, respectively, which is competitive with previous work. Agreement scores in previous work (Krippendorff's α unless otherwise noted) are, for toxicity: 0.46 (Mathew et al., 2021), 0.51 (Sap et al., 2020) and 0.64 (Hartvigsen et al., 2022); for target social group: 0.50 (Sap et al., 2020); for spans: 0.55 Cohen's kappa (Pavlopoulos et al., 2022)). To calculate agreement for spans, we incorporate MASI (Passonneau, 2006), a distance metric for inter-annotator agreement on sets, based on a weighted Jaccard, where the weight is higher for set subsumption than for set intersection, which in turn is higher than for disjunction.

Target	Count (%)	Toxic %	T. Count (%)	T. Toxic %
No target	4121 (35.96)	38.26	358 (36.46)	37.99
Ethnicity	2050 (17.78)	55.10	170 (17.31)	51.76
black	723 (6.24)	**74.90**	64 (6.52)	70.31
white	278 (2.45)	46.43	21 (2.14)	38.10
asian	272 (2.34)	47.80	23 (2.34)	43.48
native	169 (1.51)	**32.63**	16 (1.63)	**18.75**
chinese	157 (1.32)	37.95	8 (0.81)	37.50
o. ethnicity	129 (1.13)	43.66	11 (1.12)	**72.73**
mexican	114 (0.97)	38.52	9 (0.92)	44.44
arab	105 (0.90)	65.49	7 (0.71)	57.14
latino	103 (0.93)	45.30	11 (1.12)	27.27
Politics	1281 (11.05)	63.12	103 (10.49)	72.82
Gender	1152 (9.92)	49.56	87 (8.86)	55.17
lgbtq+	521 (4.50)	50.00	38 (3.87)	55.26
woman	492 (4.24)	**50.75**	38 (3.87)	**57.89**
man	121 (1.02)	48.44	9 (0.92)	**44.44**
o. gender	18 (0.17)	**14.29**	2 (0.20)	50.00
Religion	1112 (9.77)	58.62	99 (10.08)	53.53
muslim	528 (4.58)	55.11	41 (4.18)	**41.46**
jewish	474 (4.21)	**66.60**	49 (4.99)	**65.31**
o. religion	110 (0.98)	**40.65**	9 (0.92)	44.44
Other	825 (7.22)	51.49	78 (7.94)	52.56
other	466 (4.07)	**54.97**	44 (4.48)	**56.82**
refugee	188 (1.66)	**41.15**	17 (1.73)	**41.18**
middle east	171 (1.49)	53.48	17 (1.73)	52.94
Country	545 (4.73)	29.60	50 (5.09)	26.00
o. country	357 (3.11)	**30.10**	34 (3.46)	**32.35**
US	188 (1.61)	**28.57**	16 (1.63)	**12.50**
Disability	412 (3.57)	30.22	37 (3.77)	29.73
Total	11498	47.89	982	47.35

Table 2: Statistics for our human annotated data showing the total count (and percentage of the total) for each higher level or lower-level social group, and the percentage of each that are labeled toxic. Lower-level groups with the highest and lowest proportion of toxic texts are in red and green font, respectively. The last two columns are for the test set (T.). In targets, o. means other, native means native american.

To examine how greater diversity of the annotator pool would affect reliability, we recruited 8 new volunteer annotators with diversity in ethnicity, age, and training to relabel a subset of 300 examples. Ethnicity included Caucasian, Arab, Asian; ages spanned 18-32; major areas of study included Nursing, Linguistics, Biobehavioral Health, CS; sexual orientation included straight, homosexual, bisexual, and pansexual. Our agreement scores of the volunteer group plus the original annotators' labels (or excluding the original annotators' labels) were 0.63 (0.61) for toxicity, 0.62 (0.62) for target social group, and 0.45 (0.47) for toxic spans. These show that our original annotator pool is diverse enough to ensure the reliability of 10K human annotation.

After annotators were trained, they annotated

targets first, followed by toxicity, which can depend on the target, and finally spans, with consistency checks throughout. During toxicity annotation, they checked the target label. Examples where a previous target annotation was considered incorrect were relabeled by all students, with majority voting for the final label. During span annotation, they did a consistency check on the toxicity label, where those judged incorrect were relabeled, with majority voting for the final label. Span annotation identified individual words that led to the judgment that the example was toxic, or applied to the entire sentence. The latter case mostly represents sentences lacking slurs or hateful words, where toxicity instead arises from idioms, sarcasm, euphemisms, etc. (implicit toxicity) (ElSherief et al., 2021; Wen et al., 2023; Kim et al., 2024b).

3.2 Human Annotation Results

Table 2 shows the distribution for higher- or lower-level target group in the human-annotated subset, along with the proportion of these that were labeled toxic. Most target groups are ethnicity-based; the least frequent target group is for "disability." The proportion of toxic to non-toxic is nearly balanced (47.89% of the examples are toxic), but the rate varies greatly within each group. For example, 75% of the examples for the social group "black" are toxic, compared with 33% for "native american."

The total number of examples with annotations of toxic span is 5,506 (47.89% of 11,498, per Table 2). Some examples have multiple spans, so there are a total of 7,065 spans. Before computing descriptive statistics on span tokens, we preprocessed the data by applying stemming, and merging the string pairs “ni**a” and “ni**er”; “retarded” and “retard”; “stupid” and “stupidity”. Mean length of spans was 1.91 word tokens (median 1, max 11). In 33% of cases, the span constituted the entire sentence. Only 1,334 of the 7K spans are unique. The histogram in Figure 1 shows that the distribution of span tokens occurring more than once is highly skewed. It lists the 25 most common span tokens, with “ni**er,” “f**k,” “stupid,” “idiot,” and “b**ch” at the top. Some words are specific to particular groups such as “black” or “kike” but others such as “kill” or “stupid” apply in general.

4 Experiments

Our experiments utilize the human-annotated data of 11,498 examples divided into validation (N=495), test (N=982) and training using random selection. Because the validation set is so small, we constrained each target group to have a minimum of five examples. We compare performance of multiple baselines on prediction of toxicity, target social group, and for the toxic examples, prediction of spans. For the first two tasks, we report accuracy, precision, recall and F1. For span we omit accuracy. We first compare PLMs and Perspective AI with zero-shot LLMs on each task in turn, along with an examination of if prior knowledge of target social group improves toxicity detection. We also compare zero-shot LLMs with prompts from automated prompt-engineering. Finally, we investigate the use of reasoning models and in-context learning.

Model	Accuracy	Precision	Recall	F1
PerspectiveAPI	0.50	0.48	0.77	0.59
RoBERTa	**0.79**	0.76	0.82	**0.79**
BERT	**0.79**	**0.77**	0.79	0.78
GPT4o	0.75	0.70	0.85	0.77
GPT4o-TextGrad	0.74	0.67	0.86	0.76
Sonnet	0.77	0.72	0.83	0.77
Sonnet-TextGrad	0.75	0.72	0.77	0.75
Llama70b	0.77	0.71	0.86	0.78
Llama70b-TextGrad	0.75	0.69	0.86	0.76
Llama8b	0.73	0.66	0.88	0.76
Llama8b-TextGrad	0.70	0.70	0.62	0.66
Reasoning and CoT				
D. Llama70b	0.75	0.69	0.84	0.76
D. Llama8b	0.70	0.63	**0.90**	0.74
o4-mini	0.78	0.71	**0.90**	**0.79**
GPT4o-cot	0.74	0.73	0.71	0.72
Llama8b-cot	0.73	0.67	0.84	0.75
Llama70b-cot	0.75	0.72	0.75	0.74
Sonnet-cot	0.73	0.73	0.67	0.70

Table 3: Toxicity detection results. The *D.* prefix stands for distilled.

PLM baselines for toxicity and target social group use BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) with linear classifier layers, and for span prediction we use SpanBERT Joshi et al. (2020) (see appendix B). We include four LLMs: GPT4o (OpenAI et al., 2024), Claude 3.7 Sonnet (Anthropic, 2025), and Llama3.1 (70b and 8b) (Grattafiori et al., 2024) (see Appendix E for the prompt). We use temperature=0, set the seed, and use the default values for others for determinism, although it is not guaranteed (Atil et al., 2025). For toxicity, we also include PerspectiveAPI (Lees et al., 2022), a neural network that provides a probability that an example is toxic. We used the validation set to identify the best probability for the binary class cutoff, which was 0.20.

4.1 Toxicity

The toxicity results in the top of Table 3 show that fine-tuned PLMs perform slightly better than the

Model	Accuracy	Precision	Recall	F1
BERT	**0.76**	**0.68**	**0.82**	**0.72**
RoBERTa	0.73	0.63	0.81	0.70
GPT4o	0.75	0.67	0.78	0.70
GPT4o-TextGrad	0.71	0.57	0.69	0.61
Sonnet	0.74	0.62	0.75	0.67
Sonnet-TextGrad	0.75	0.63	0.68	0.62
Llama70b	0.65	0.55	0.68	0.58
Llama70b-TextGrad	0.69	0.34	0.38	0.34
Llama8b	0.48	0.15	0.15	0.15
Llama8b-TextGrad	0.43	0.10	0.11	0.10
Reasoning and CoT				
o4-mini	0.72	0.58	0.63	0.59
D. Llama70b	0.69	0.24	0.27	0.25
D. Llama8b	0.61	0.17	0.17	0.17
GPT4o-cot	0.75	0.68	0.75	0.70
Llama8b-cot	0.40	0.07	0.06	0.06
Llama70b-cot	0.68	0.50	0.58	0.53
Sonnet-cot	0.75	0.62	0.70	0.65

Table 4: Target group prediction results.

Model	Accuracy	Precision	Recall	F1
BERT	**0.80**	**0.77**	**0.82**	**0.79**
RoBERTa	0.78	0.74	0.81	0.77
GPT4o	0.66	0.64	0.65	0.64
GPT4o-TextGrad	0.64	0.44	0.44	0.44
Sonnet	0.71	0.65	0.70	0.66
Sonnet-TextGrad	0.67	0.62	0.63	0.62
Llama70b	0.57	0.46	0.48	0.44
Llama70b-TextGrad	0.42	0.35	0.30	0.32
Llama8b	0.49	0.07	0.07	0.07
Llama8b-TextGrad	0.53	0.06	0.07	0.06
Reasoning and CoT				
o4-mini	0.67	0.45	0.47	0.45
D. Llama70b	0.64	0.17	0.18	0.17
D. Llama8b	0.56	0.07	0.07	0.07
GPT4o-cot	0.68	0.56	0.61	0.58
Llama8b-cot	0.44	0.05	0.05	0.05
Llama70b-cot	0.57	0.31	0.31	0.31
Sonnet-cot	0.70	0.55	0.60	0.57

Table 5: Higher level target results.

LLMs. Except for Llama8b, LLMs perform similarly. Surprisingly, PerspectiveAPI's accuracy is random. All models have higher recall which is preferable here, where false negatives are worse than false positives.

4.2 Target Social Group

Similar to Zampieri et al. (2023), for two baselines to predict target social group, we train a neural network with a linear classifier layer on top of BERT or RoBERTa encoder. As above, we also test the same SOTA LLMs (cf. appendices F and G for the prompts). Again we find that a fine-tuned PLM slightly outperforms LLMs; cf. the fine-grained target results in Table 4. However, Llama models do not perform as well as GPT4o and Sonnet. The F1 for Llama8b is especially low. Table 5 shows the results for the higher-level target groups. Interestingly, LLMs perform worse at this task than the fine-grained target groups (9% decrease for GPT4o, 3% decrease for Sonnet etc.). Confusion matrices

Model	Precision	Recall	F1
SpanBERT	**0.72**	0.71	**0.70**
GPT4o	0.55	0.79	0.65
GPT4o-TextGrad	0.67	0.35	0.46
Sonnet	0.66	0.45	0.53
Sonnet-TextGrad	0.70	0.32	0.44
Llama70b	0.66	0.22	0.33
Llama70b-TextGrad	0.74	0.11	0.19
Llama8b	0.48	0.48	0.48
Llama8b-TextGrad	0.74	0.10	0.17
Reasoning and Cot			
o4-mini	0.63	0.40	0.49
D. Llama70b	0.45	**0.87**	0.59
D. Llama8b	0.43	0.58	0.40
Sonnet-cot	0.63	0.5	0.56
GPT4o-cot	0.55	0.63	0.59
Llama70b-cot	0.68	0.46	0.54
Llama8b-cot	0.68	0.16	0.26

Table 6: Toxic span prediction results.

for GPT4o and Sonnet (the two top-performing LLMs), show that GPT4o mixes "other" and "ethnicity" predictions with "no target", and Sonnet mixes "no target" with "ethnicity". Additionally, both mix "ethnicity" with "other". This indicates "other" as a high level target is more confusing than the fine-grained targets.

Experiments to assess whether prior knowledge of target social group improves toxicity detection showed no impact (see Table 8). For the RoBERTa/BERT-based models, we either added target information at the text level as "The target social group is <social group>" or at the embedding level. For LLMs, we used two types of prompting: (1) assigning the social group persona to the model, since combining persona with self-correction has been effective (Xu et al., 2024); (2) including the social group in the text of the prompt.

4.3 Toxic Span

The span detection results in Table 6 show that SpanBERT reaches 0.70 F1, outperforming the LLMs, and has balanced recall and precision. GPT4o's F1 approaches SpanBERT's, but the other LLMs do much worse.

4.4 Automated Prompt Engineering

Automated prompt engineering (APO) has been effective on code optimization, agent planning, mathematical reasoning etc. (Yuksekgonul et al., 2025; Ramnath et al., 2025). We used the SOTA TextGrad model (Yuksekgonul et al., 2025). TextGrad uses a meta (optimizer) LLM to criticize the current prompt and suggest improvements and the new prompt is selected based on the performance on the validation set. We used GPT4o as the meta LLM,

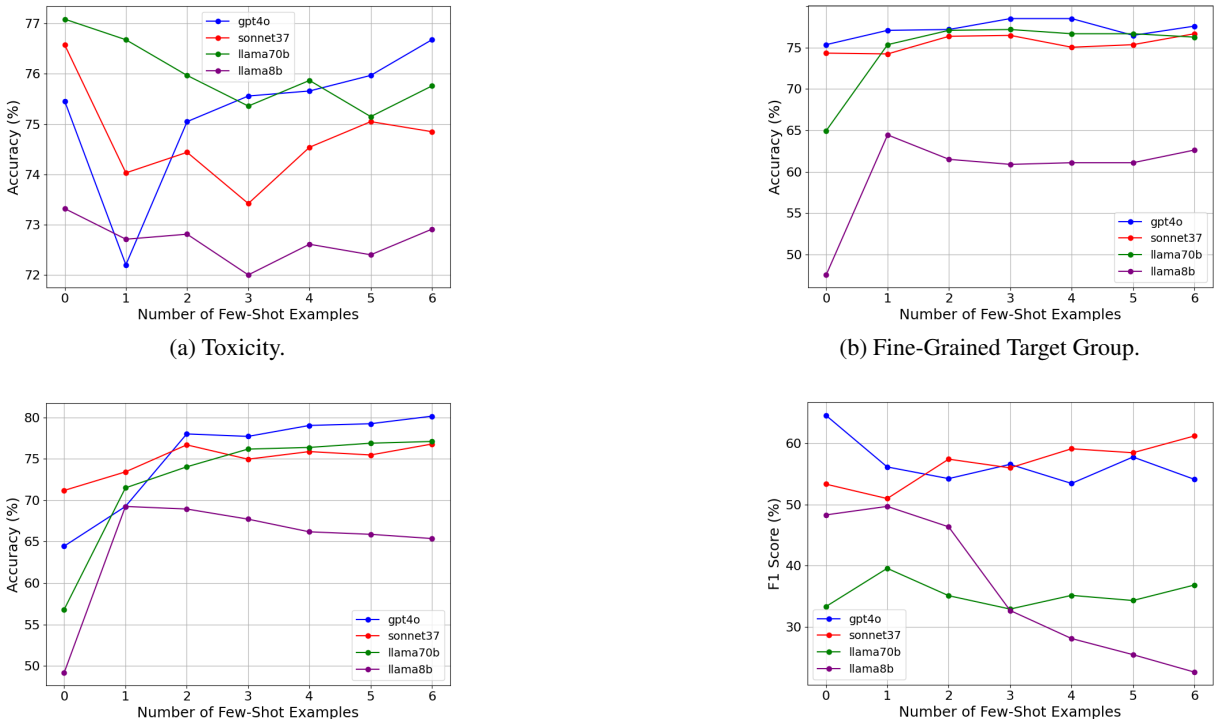

(a) Toxicity.

(b) Fine-Grained Target Group.

(c) High-Level Target Group.

(d) Toxic Span.

Figure 2: Few-Shot Comparison for Zero to 6-Shot across Four Tasks with Four Models.

and optimized the prompts for 15 steps. Our training and validation data have 100 and 200 examples, respectively. Tables 3-6 show that TextGrad optimizations do not produce improvements. One potential reason for this is that we have already done manual prompt engineering, so they are close to optimal. Another reason might be that TextGrad is not very effective on subjective tasks. We leave the detailed investigation of APO for future work.

4.5 Reasoning Models

Although reasoning helps in science or logic (Jaech et al., 2024; DeepSeek-AI, 2025; Zhang et al., 2024; Wei et al., 2023), Tables 3-6 show no improvement on detection of toxicity or target group, and mixed results on span prediction. We experimented with chain-of-thought (CoT) (Wei et al., 2023) and the reasoning models o4-mini and R1 Distilled Llama70b/8b. CoT helps Llama70b and Sonnet, but not Llama8b and Gpt4o. Reasoning usually increases recall, which indicates that models predict more tokens less precisely.

4.6 In-Context Learning

Few-shot learning improves LLM performance in many tasks, such as question answering (Brown et al., 2020). To select few-shot examples, previous work suggested that similarity to the test examples is effective (Paraschiv et al., 2023; Liu et al., 2022; Zebaze et al., 2024), which was borne out in our early experiments. We use Linq-Embed-Mistral (Kim et al., 2024a) to create embeddings, and retrieve similar examples from our training data. Figure 2 shows results for zero- to 6-shot with 4 non-reasoning models. For toxicity detection, examples only help GPT4o, and only with at least three examples. For toxic span prediction, one example is enough to increase the performance for Llama models, whereas Sonnet requires at least two examples. For target group prediction, few-shot learning helps all models, and the improvements are more dramatic, e.g., from 0.48 to 0.64 for Llama8b with one example. However, apart from prediction of target group, LLMs still do not surpass PLMs.

4.7 Error Analysis

The best results reported above (accuracy; F1) are from the RoBERTa classifier for toxicity (0.79; 0.79), the BERT classifier for fine-grained target (0.76; 0.72) and course-grained target (0.80; 0.79); SpanBERT for toxic spans (F1 of 0.70). These methods were used to annotate all remaining examples in TRuST, apart from the human annotation. In general, LLMs do better at toxicity detection than the other two tasks, and reasoning/CoT helps very little. Here we look into 1) differences across

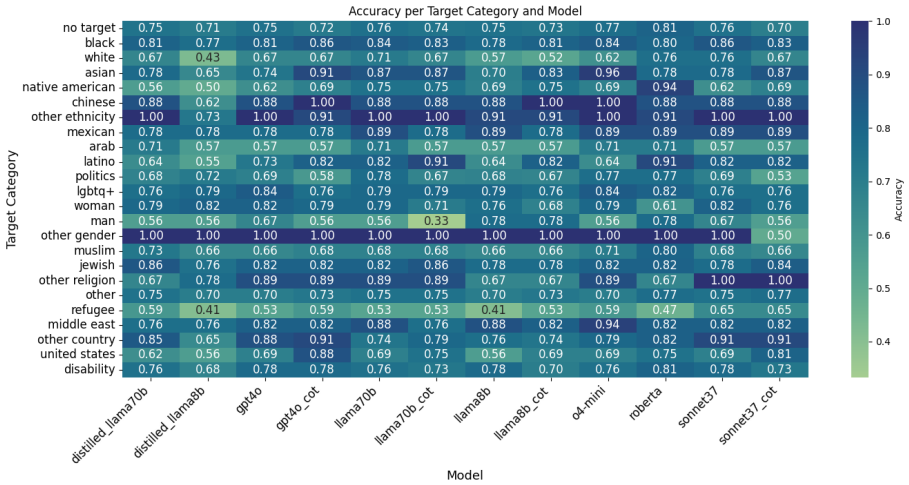

Figure 3: Toxicity detection accuracy by target group

targets for toxicity and target prediction; 2) whether span detection is better when the span is the whole sentence; 3) human vs. machine-generated text.

Differences across social groups. We broke down the accuracy results on toxicity and target group detection by target group. Figure 3 is a heatmap for toxicity detection, which shows wide variation across target groups, ranging from 41% to 100%. Apart from the categories for "other gender" and "other ethnicity," where most models have very high performance, all models show non-uniform performance across targets. In some cases, pairs of models correlate in their performance (e.g. Spearman correlations are 0.89 for GPT4o and Sonnet; 0.87 for o4-mini and Llama70b; 0.83 for Sonnet and Llama70b), but no model performance correlates well with data support.

We also found large differences in accuracy across target social groups (see Figure 6 in Appendix L). For example, Llama8b has a 42% difference in accuracy for "black" versus "white." The models struggle in categories such as "other country" or "other gender," possibly because these categories are necessarily more heterogenous.

Span detection for sentences. As mentioned in Section 3.1, toxic spans sometimes consist of the entire sentence (35% of the test set). Most models, including SpanBERT, have much higher performance on full-sentence toxic spans. For example, SpanBERT has a score of 0.67 on subsentence spans, whereas has a score of 0.76 on full sentence. Distilled Llama models are far worse at detecting toxic subsentence spans. (Cf. Figure 4 in Appendix B, a plot of F1 for full sentence spans on the x-axis by subsentence spans on the y-axis.)

LLM-generated text. Our test data composition is 40% GPT-3 generated (from ToxiGen). Breaking down accuracy by human versus LLM origin shows 20% greater accuracy on the GPT-3 generated text.

5 Conclusion

We believe TRuST can support research into evaluating and mitigating LLM toxicity. About 38% of the 298k examples are toxic, and have been labeled for a larger number of fine-grained target social groups (including "no target") than in previous work, as well as for toxic spans either by reliable humans, or by high-performing models. Over 11K examples were manually annotated by a diverse group of annotators whose inter-annotator agreements on the three labeling tasks were higher than reported in previous work. We benchmarked 13 methods on the three toxicity prediction tasks. Interestingly, models performed worse on higher-level target prediction than the fine-grained one, possibly because the “other” higher-level target is very heterogeneous. Unsurprisingly, the PLM classifiers outperformed the LLMs. Notably, neither CoT nor reasoning models led to higher LLM improved results, underscoring that current reasoning approaches fail to effectively handle the specification of toxicity used here, across diverse targets.

Our best performing models were used to label the remaining 287K examples. These were a RoBERTa PLM for toxicity, a BERT PLM for target group, and SpanBERT for toxic span. We believe TRuST can foster new research on LLM toxicity, including unlearning (Chen and Yang, 2023; Liu et al., 2024), which to our knowledge has not yet been applied to toxicity.

6 Limitations

The work presented here carries out only a preliminary investigation of baseline methods for automatic identification of toxicity, target social group and toxic span detection. The LLMs methods did not explore sophisticated prompt engineering. Although the size of the dataset is competitive, it is not sufficiently large to have separate annotations for some important subgroups. Although we attempted to recruit a pool of annotators that was socially diverse, this was limited due to lack of funds to pay more than six annotators.

References

Saad Almohaimeed, Saleh Almohaimeed, Ashfaq Ali Shafin, Bogdan Carbunar, and Ladislau Bölöni. 2023. THOS: A benchmark dataset for targeted hate and offensive speech. *arXiv preprint arXiv:2311.06446*.

Anthropic. 2025. Claude 3.7 sonnet. Large Language Model.

Berk Atil, Sarp Aykent, Alexa Chittams, Lisheng Fu, Rebecca J. Passonneau, Evan Radcliffe, Guru Rajan Rajagopal, Adam Sloan, Tomasz Tudrej, Ferhan Ture, Zhe Wu, Lixinyu Xu, and Breck Baldwin. 2025. Non-determinism of "deterministic" llm settings. *Preprint*, arXiv:2408.04667.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Jiaao Chen and Diyi Yang. 2023. Unlearn what you want to forget: Efficient unlearning for llms. *Preprint*, arXiv:2310.20150.

Thomas Davidson, Dana Warmsley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media*, volume 11.1, pages 512–515.

Adelson de Araujo, Pantelis M. Papadopoulos, Susan McKenney, and Ton de Jong. 2025. Investigating the impact of a collaborative conversational agent on dialogue productivity and knowledge acquisition. *International Journal of Artificial Intelligence in Education*.

DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *Preprint*, arXiv:2501.12948.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.

Rebecca Dorn, Lee Kezar, Fred Morstatter, and Kristina Lerman. 2024. Harmful speech detection by language models exhibits gender-queer dialect bias. In *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–12.

Xiaoni Duan, Zhuoyan Li, Chien-Ju Ho, and Ming Yin. 2025. Exploring the cost-effectiveness of perspective taking in crowdsourcing subjective assessment: A case study of toxicity detection. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2359–2372, Albuquerque, New Mexico. Association for Computational Linguistics.

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

Beyza Ermis, Luiza Pozzobon, Sara Hooker, and Patrick Lewis. 2024. From one to many: Expanding the scope of toxicity mitigation in language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15041–15058, Bangkok, Thailand. Association for Computational Linguistics.

Eve Fleisig, Rediet Abebe, and Dan Klein. 2023a. When the majority is wrong: Modeling annotator disagreement for subjective tasks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6715–6726, Singapore. Association for Computational Linguistics.

Eve Fleisig, Aubrie Amstutz, Chad Atalla, Su Lin Blodgett, Hal Daumé III, Alexandra Olteanu, Emily Sheng, Dan Vann, and Hanna Wallach. 2023b. FairPrism: Evaluating fairness-related harms in text generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6231–6251, Toronto, Canada. Association for Computational Linguistics.

Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. RealToxicityPrompts: Evaluating neural toxic degeneration in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, Online. Association for Computational Linguistics.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. The llama 3 herd of models. *Preprint*, arXiv:2407.21783.

Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. ToxiGen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3309–3326, Dublin, Ireland. Association for Computational Linguistics.

Xinlei He, Savvas Zannettou, Yun Shen, and Yang Zhang. 2024. You only prompt once: On the capabilities of prompt learning on large language models to tackle toxic content. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 770–787. IEEE.

Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.

Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the association for computational linguistics*, 8:64–77.

Sakshi Joshi, Anindita Mukherjee, Usha A Jogalekar, Renuka Agrawal, Abhishek Anand, and Santhosh Phanitalpak Gandhala. 2025. Enhancing toxic comment classification with multi-label capabilities: Leveraging bert and roberta models. In *2025 International Conference on Emerging Trends in Industry 4.0 Technologies (ICETI4T)*, pages 1–6. IEEE.

Julia Kansok-Dusche, Cindy Ballaschk, Norman Krause, Anke Zeißig, Lisanne Seemann-Herz, Sebastian Wachs, and Ludwig Bilz. 2023. A systematic review on hate speech among children and adolescents: Definitions, prevalence, and overlap with related phenomena. *Trauma, violence, & abuse*, 24(4):2598–2615.

Urja Khurana, Eric Nalisnick, and Antske Fokkens. 2025. DefVerify: Do hate speech models reflect their dataset‘s definition? In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 4341–4358, Abu Dhabi, UAE. Association for Computational Linguistics.

Urja Khurana, Ivar Vermeulen, Eric Nalisnick, Marloes Van Noorloos, and Antske Fokkens. 2022. Hate speech criteria: A modular approach to task-specific hate speech definitions. In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 176–191, Seattle, Washington (Hybrid). Association for Computational Linguistics.

Junseong Kim, Seolhwa Lee, Sangmo Gu Jihoon Kwon, Yejin Kim, Minkyung Cho, Jy yong Sohn, and Chanyeol Choi. 2024a. Linq-embed-mistral:elevating text retrieval with improved gpt data through task-specific control and quality refinement. Linq AI Research Blog.

Minbeom Kim, Jahyun Koo, Hwanhee Lee, Joonsuk Park, Hwaran Lee, and Kyomin Jung. 2024b. LifeTox: Unveiling implicit toxicity in life advice. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 688–698, Mexico City, Mexico. Association for Computational Linguistics.

Sohyun An Kim, Lauren Baczewski, Maria Pizzano, Connie Kasari, and Alexandra Sturm. 2023. Discrimination and harassment experiences of autistic college students and their neurotypical peers: Risk and protective factors. *Journal of Autism and Developmental Disorders*, 53(12):4521–4534.

Diederik P. Kingma and Jimmy Ba. 2017. Adam: A method for stochastic optimization. *Preprint*, arXiv:1412.6980.

Klaus Krippendorff. 2013. *Content analysis: An introduction to its methodology*, third edition. Sage, Thousand Oaks, CA.

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. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, pages 299–318.

Alyssa Lees, Vinh Q. Tran, Yi Tay, Jeffrey Sorensen, Jai Gupta, Donald Metzler, and Lucy Vasserman. 2022. A New Generation of Perspective API: Efficient Multilingual Character-level Transformers. *Preprint*, arXiv:2202.11176.

Xiaochen Li, Zheng Xin Yong, and Stephen Bach. 2024. Preference tuning for toxicity mitigation generalizes across languages. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13422–13440, Miami, Florida, USA. Association for Computational Linguistics.

Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. 2021. DExperts: Decoding-time controlled text generation with experts and anti-experts. In *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)*, pages 6691–6706, Online. Association for Computational Linguistics.

Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. What makes good in-context examples for GPT-3? In

Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures, pages 100–114, Dublin, Ireland and Online. Association for Computational Linguistics.

Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Yuguang Yao, Chris Yuhao Liu, Xiaojun Xu, Hang Li, Kush R. Varshney, Mohit Bansal, Sanmi Koyejo, and Yang Liu. 2024. Rethinking machine unlearning for large language models. *Preprint*, arXiv:2402.08787.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *Preprint*, arXiv:1907.11692.

Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021. HateXplain: A benchmark dataset for explainable hate speech detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35.17, pages 14867–14875.

Aida Mostafazadeh Davani, Mark Diaz, Dylan K Baker, and Vinodkumar Prabhakaran. 2024. D3CODE: Disentangling disagreements in data across cultures on offensiveness detection and evaluation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18511–18526, Miami, Florida, USA. Association for Computational Linguistics.

John T Nockleby. 2000. Hate speech. *Encyclopedia of the American Constitution (2nd ed., edited by Leonard W. Levy, Kenneth L. Karst et al., New York: Macmillan, 2000)*, pages 1277–1279.

OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, and 401 others. 2024. Gpt-4o system card. *Preprint*, arXiv:2410.21276.

Andrei Paraschiv, Teodora Andreea Ion, and Mihai Dascalu. 2023. Offensive text span detection in Romanian comments using large language models. *Information*, 15(1):8.

Rebecca J. Passonneau. 2006. Measuring Agreement on Set-valued Items (MASI) for Semantic and Pragmatic Annotation. In *LREC*, pages 831–836.

John Pavlopoulos, Leo Laugier, Alexandros Xenos, Jeffrey Sorensen, and Ion Androutsopoulos. 2022. From the detection of toxic spans in online discussions to the analysis of toxic-to-civil transfer. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3721–3734, Dublin, Ireland. Association for Computational Linguistics.

Agnieszka Pluta, Joanna Mazurek, Jakub Wojciechowski, Tomasz Wolak, Wiktor Soral, and Michał Bilewicz. 2023. Exposure to hate speech deteriorates neurocognitive mechanisms of the ability to understand others' pain. *Scientific Reports*, 13(1):4127.

Luiza Pozzobon, Beyza Ermis, Patrick Lewis, and Sara Hooker. 2023a. Goodtriever: Adaptive toxicity mitigation with retrieval-augmented models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5108–5125, Singapore. Association for Computational Linguistics.

Luiza Pozzobon, Beyza Ermis, Patrick Lewis, and Sara Hooker. 2023b. On the challenges of using black-box APIs for toxicity evaluation in research. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7595–7609, Singapore. Association for Computational Linguistics.

Kiran Ramnath, Kang Zhou, Sheng Guan, Soumya Smruti Mishra, Xuan Qi, Zhengyuan Shen, Shuai Wang, Sangmin Woo, Sullam Jeoung, Yawei Wang, Haozhu Wang, Han Ding, Yuzhe Lu, Zhichao Xu, Yun Zhou, Balasubramaniam Srinivasan, Qiaojing Yan, Yueyan Chen, Haibo Ding, and 2 others. 2025. A systematic survey of automatic prompt optimization techniques. *Preprint*, arXiv:2502.16923.

Anchal Rawat, Santosh Kumar, and Surender Singh Samant. 2024. Hate speech detection in social media: Techniques, recent trends, and future challenges. *Wiley Interdisciplinary Reviews: Computational Statistics*, 16(2):e1648.

Lucas Rosenblatt, Lorena Piedras, and Julia Wilkins. 2022. Critical Perspectives: A Benchmark Revealing Pitfalls in PerspectiveAPI. In *Proceedings of the Second Workshop on NLP for Positive Impact (NLP4PI)*, pages 15–24.

Koustuv Saha, Eshwar Chandrasekharan, and Munmun De Choudhury. 2019. Prevalence and psychological effects of hateful speech in online college communities. In *Proceedings of the 10th ACM conference on web science*, pages 255–264.

Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A. Smith, and Yejin Choi. 2020. Social Bias Frames: Reasoning about Social and Power Implications of Language. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5477–5490, Online. Association for Computational Linguistics.

Timo Schick, Sahana Udupa, and Hinrich Schütze. 2021. Self-Diagnosis and Self-Debiasing: A proposal for reducing corpus-based bias in nlp. *Transactions of the Association for Computational Linguistics*, 9:1408–1424.

Omar Shaikh, Hongxin Zhang, William Held, Michael Bernstein, and Diyi Yang. 2023. On second thought,

let's not think step by step! bias and toxicity in zero-shot reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4454–4470, Toronto, Canada. Association for Computational Linguistics.

Xavier Suau, Pieter Delobelle, Katherine Metcalf, Armand Joulin, Nicholas Apostoloff, Luca Zappella, and Pau Rodríguez. 2024. Whispering experts: Neural interventions for toxicity mitigation in language models. *arXiv preprint arXiv:2407.12824*.

Tsuang Han Tsai, Hsu Sung Ting, Huang Yu Cheng, Hsueh Shan Ting, and Yuan Chen Yao. 2025. Fine-tuning distilbert for toxic comment detection and classification. In *2025 IEEE Gaming, Entertainment, and Media Conference (GEM)*, pages 1–4. IEEE.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-thought prompting elicits reasoning in large language models. *Preprint*, arXiv:2201.11903.

Yuchen Wei, Dennis Pearl, Matthew Beckman, and Rebecca J. Passonneau. 2025. Concept-based rubrics improve llm formative assessment and data synthesis. *Preprint*, arXiv:2504.03877.

Johannes Welbl, Amelia Glaese, Jonathan Uesato, Sumanth Dathathri, John Mellor, Lisa Anne Hendricks, Kirsty Anderson, Pushmeet Kohli, Ben Coppin, and Po-Sen Huang. 2021. Challenges in detoxifying language models. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2447–2469, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Jiaxin Wen, Pei Ke, Hao Sun, Zhexin Zhang, Chengfei Li, Jinfeng Bai, and Minlie Huang. 2023. Unveiling the implicit toxicity in large language models. *arXiv preprint arXiv:2311.17391*.

Rongwu Xu, Zian Zhou, Tianwei Zhang, Zehan Qi, Su Yao, Ke Xu, Wei Xu, and Han Qiu. 2024. Walking in others' shoes: How perspective-taking guides large language models in reducing toxicity and bias. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8341–8368, Miami, Florida, USA. Association for Computational Linguistics.

Yee Hui Yeo, Allistair Clark, Muskaan Mehra, Itai Danovitch, Karen Osilla, Ju Dong Yang, Alexander Kuo, Hyun-Seok Kim, Aarshi Vipani, Yun Wang, Walid Ayoub, Hirsh Trivedi, Jamil S. Samaan, Tiffany Wu, Vijay H. Shah, Omer Liran, and Brennan Spiegel. 2024. The feasibility and usability of an artificial intelligence-enabled conversational agent in virtual reality for patients with alcohol-associated cirrhosis: A multi-methods study. *Journal of Medical Extended Reality*, 1(1):257–270.

Yee Hui Yeo, Yuxin Peng, Muskaan Mehra, Jamil Samaan, Joshua Hakimian, Allistair Clark, Karisma Suchak, Zoe Krut, Taiga Andersson, Susan Persky, Omer Liran, and Brennan Spiegel. 2025. Evaluating for evidence of sociodemographic bias in conversational ai for mental health support. *Cyberpsychology, Behavior, and Social Networking*, 28(1):44–51.

Shuzhen Yu, Alexey Androsov, and Hanbing Yan. 2025. Exploring the prospects of multimodal large language models for automated emotion recognition in education: Insights from gemini. *Computers & Education*, 232:105307.

Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Pan Lu, Zhi Huang, Carlos Guestrin, and James Zou. 2025. Optimizing generative ai by backpropagating language model feedback. *Nature*, 639(8055):609–616.

Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Predicting the type and target of offensive posts in social media. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1415–1420, Minneapolis, Minnesota. Association for Computational Linguistics.

Marcos Zampieri, Skye Morgan, Kai North, Tharindu Ranasinghe, Austin Simmmons, Paridhi Khandelwal, Sara Rosenthal, and Preslav Nakov. 2023. Target-based offensive language identification. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 762–770, Toronto, Canada. Association for Computational Linguistics.

Armel Zebaze, Benoît Sagot, and Rachel Bawden. 2024. In-context example selection via similarity search improves low-resource machine translation. *Preprint*, arXiv:2408.00397.

Xuan Zhang, Chao Du, Tianyu Pang, Qian Liu, Wei Gao, and Min Lin. 2024. Chain of preference optimization: Improving chain-of-thought reasoning in llms. *Advances in Neural Information Processing Systems*, 37:333–356.

Xuhui Zhou, Hao Zhu, Akhila Yerukola, Thomas Davidson, Jena D. Hwang, Swabha Swayamdipta, and Maarten Sap. 2023. COBRA frames: Contextual reasoning about effects and harms of offensive statements. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6294–6315, Toronto, Canada. Association for Computational Linguistics.

A Statistics on the Whole Dataset

Target	H. Count (%)	H. Toxic %	M. Count (%)	M. Toxic %	C. Count (%)	C. Toxic %
No Target	3618 (36.10)	38.92	73211 (25.40)	24.03	76829 (25.76)	24.73
Ethnicity	1763 (17.59)	56.21	68962 (23.93)	48.02	70725 (23.71)	48.22
black	626 (6.25)	75.56	15654 (5.43)	59.93	16280 (5.46)	60.53
white	244 (2.43)	49.18	8691 (3.02)	51.29	8935 (3.00)	51.24
asian	238 (2.37)	49.16	12961 (4.50)	45.36	13199 (4.43)	45.43
native american	143 (1.43)	36.36	8817 (3.06)	34.04	8960 (3.00)	34.07
chinese	139 (1.39)	37.41	10242 (3.55)	41.53	10381 (3.48)	41.47
other ethnicity	108 (1.08)	38.89	5156 (1.79)	35.47	5264 (1.76)	35.54
mexican	95 (0.95)	35.79	6221 (2.16)	42.65	6316 (2.12)	42.54
arab	88 (0.88)	65.91	2267 (0.79)	70.27	2355 (0.79)	70.11
latino	82 (0.82)	52.44	5044 (1.75)	49.15	5126 (1.72)	49.20
Politics	1132 (11.30)	62.54	9624 (3.34)	57.41	10756 (3.61)	57.95
Gender	1000 (9.98)	50.20	34280 (11.89)	42.67	35280 (11.83)	42.88
lgbtq+	456 (4.55)	51.10	15598 (5.41)	41.50	16054 (5.38)	41.77
woman	431 (4.30)	50.58	14839 (5.15)	47.94	15270 (5.12)	48.02
man	102 (1.02)	49.02	5248 (1.82)	43.45	5350 (1.79)	43.55
other gender	11 (0.11)	9.09	1531 (0.53)	14.30	1542 (0.52)	14.27
Religion	966 (9.64)	60.97	30381 (10.54)	50.65	31347 (10.51)	50.97
muslim	468 (4.67)	58.33	12707 (4.41)	44.31	13175 (4.42)	44.81
jewish	407 (4.06)	68.06	14765 (5.12)	58.23	15172 (5.09)	58.49
other religion	91 (0.91)	42.86	5127 (1.78)	31.44	5218 (1.75)	31.64
Other	716 (7.14)	51.96	14771 (5.12)	43.73	15487 (5.19)	44.11
other	412 (4.11)	55.34	2489 (0.86)	52.95	2901 (0.97)	53.29
refugee	160 (1.60)	42.50	4710 (1.63)	39.53	4870 (1.63)	39.63
middle east	144 (1.44)	52.78	7408 (2.57)	42.54	7552 (2.53)	42.73
Country	470 (4.69)	31.49	23189 (8.05)	25.02	23659 (7.93)	25.15
other country	308 (3.07)	30.52	12511 (4.34)	25.01	12819 (4.30)	25.14
united states	162 (1.62)	33.33	9080 (3.15)	25.55	9242 (3.10)	25.69
Disability	357 (3.56)	30.53	24332 (8.44)	27.05	24689 (8.28)	27.11
Total	10222	48.16	288233	37.79	298255	38.14

Table 7: The statistics on the training data including machine annotated data. The second and third columns are human annotated training data (H), the fourth and fifth are machine annotated training data (M), and the last two are the combination of both (C). The results for the higher level categories for human annotation are computed by combining the fine-grained category results. For the machine generated, however, we carry out independent tests of the higher and lower categories. As a result, we find a discrepancy of about 10% of cases.

B Experimental Details

We use AdamW optimizer (Kingma and Ba, 2017) to fine-tune PLMs. We use a batch size of 8. We tried 0.01, 0.0001, 0.00001, 0.05,0.005,0.0005, 0.00005 for learning rate and we chose the best one for each task, we chose the best one based on the validation performance. The best learning rate is 1e-05 for all tasks and embedders. We ran the experiments on a single NVIDIA RTX A6000 GPU and it took 2 hours for each experiment.

C Annotation Instructions

Goal Creating annotated data for toxic language, where, very generally, a sentence is toxic if it has negative stereotyping, hate speech, racism, psychological threat, sexual harassment, abusive language, sexism, discrimination based on sexual orientation, or any other type of language that might hurt or affect a member of some sociodemographic group badly.

Task Annotate these 3 categories:

- Toxicity (binary)
- Social target group (24 categories belonging to 6 higher level groups including other, and no target)
- Toxic Span (words that make the sentence toxic, only for toxic sentences.)

Assumptions

- There is no **speaker** information (who produced the language)
- There is no **audience** information (who the speaker was addressing, not necessarily the target).
- The toxicity label should take into account the target social group (note that the same words might be toxic for one social group target and not for another)
- There is no **context**, so texts should be interpreted as they are, based on easily available readings.
- If you see a politician's name/party name, choose politics.
- If the target is individual, we are not interested in so they should be labeled as no target.
- If there are multiple targets, try to choose the dominating one. In case they are equally targeted, choose other.

Target Groups **Higher Level Target Groups** We have 5 distinct higher level social targets, plus labels for "Other" category (there is a target but it doesn't fall into these categories) and absence of a target ("No Target").

- Ethnicity/Nationality
- Gender
- Religion
- Disability
- Country
- Other
- No Target

Full Set of Target Groups We have more specific targets within higher level categories. For each category, except individual and disability, we also have other. You should use the label "other $< higher_level_target >$" in the case that the group belongs to $< higher_level_target >$ and not among the ones listed below (e.g. for the statement "Turkish people are useless", the target is ethnicity, but there is no Turkish category so it should be other ethnicity)

- Black, white, Asian, Mexican, Native American, Chinese, Latino, and Arab belong to **ethnicity**.
- LGBTQ+, Woman, and Man belong to **gender**.
- Muslim and Jewish belong to **religion**.
- Disability does not have any sub-targets.
- United States is the only country belonging to **country**. You should label as "other country" if it is another country.
- Middle East and refugee belong to **other**.
- Immigrants belongs to **refugee**

Note: Do not forget about the "other" categories within each category, except individual and disability!

Toxicity Label Criteria For assigning positive/negative toxicity; toxic if any of the following categories, where the first five (blue font) are varieties of hate speech, the next two are examples of abusive language other than sexual harrassment, the final category.

- **Negative stereotyping**: Traits, characteristics and generalizations that are negatively attributed towards a group. e.g. Turkish people are selfish
- **Hate speech**: Offensive and discriminatory discourse towards a group or an individual based on characteristics such as race or religion. E.g. All Jews should be killed.
- **Racism**: Usually a bias that includes acceptance of negative stereotypes, and racial discrimination (discrimination; different treatment of different age, gender, racial etc., race in this case) against individuals/groups belonging to a certain race. E.g Arabs should not be offered any jobs.
- **Sexist language**: Discrimination or behaviors that foster stereotypes based on a gender. E.g. Men cannot be successful in teaching
- **Discriminating based on sexual orientation**: E.g. Being lesbian should be banned.
- **Psychological threat**: It is an expression of an intent of harm, such as humiliation, causing stress, or constant criticism. E.g. You always suck
- **Abusive language**: Piece of text that contains inappropriate words such as swear profanity, and disrespectful nicknames. E.g It is no wonder everyone says you're a jerk
- **Sexual harassment**: Unwelcome sexual moves, sexual favor requests or other physical/verbal behaviors of a sexual nature towards someone. E.g. Your breasts are amazing.

Note: there are other situations of language use that can be perceived as toxic that do not fall under our guidelines given our assumptions, e.g., some microaggressions might be considered toxic but to identify them would require other information such as the audience, more of the overall context in which the language occurs, etc.

Toxic Spans You should choose the word(s) that causes the toxicity of the statement (e.g. We need to **ignore** the needs of disabled people).

General Instructions

- We should think as being in the target group.
- You should consider multiple possible interpretations since almost every statement/utterance has multiple "readings"; If there is an easily available understanding that leads to being non-toxic, it should be labeled as non-toxic.
- For the higher level target, you need to check if the opinion is towards some group, individual or no one. There might be more than one group, in that case try to choose the dominant group. Also, if it is for an individual and seems like they belong to a social group, do not try to infer a social group for them, instead label them as "individual"

D Potential Risks

Due to the nature of our work, TRuST has some potential risks of increasing toxic content. However, all text data were already public before our work. We add human and machine annotations to these texts, which should not increase the current potential risk. We hope this work will be used for good, not to make the system more harmful.

E LLM Prompt for Toxicity

System Prompt: You are an expert in linguistic and sociological analysis. Your task is to analyze a text and determine if it is toxic or not. Toxicity here refers to negative stereotyping, hate speech, racism, sexist language, discrimination based on sexual orientation, physical or psychological threat, and abusive language. Note that there are other situations of language use that can be perceived as toxic that do not fall under our guidelines, given our assumptions For each text, at the end, respond in the following format: Toxicity Label: [toxic/non-toxic]

User Prompt: Text: <text>

F LLM Prompt for Target Group

System Prompt: You are an expert in linguistic and sociological analysis. Your task is to determine which social group is targeted by the sentiment or content of each text. For each text, use the provided list of social groups and always select the group that is most directly addressed or implied. If an individual not involved in politics is mentioned, or if no group is targeted, respond with 'Targeted Group: no target'. If multiple groups are targeted with no clear dominance, or if a group is certainly targeted but not clearly identified, respond with 'Targeted Group: other.' Here are the social groups:

- no target
- other country
- united states (notes: must be related to the United States as a whole)
- black
- asian (note: ethnicity, excludes Chinese)
- white
- arab
- latino (note: ethnicity, includes Hispanic and excludes Mexican)
- native american
- chinese
- mexican
- other ethnicity
- jewish
- muslim
- other religion
- politics (note: must reference a political figure or party)
- woman
- man
- lgbtq+
- other gender
- disability
- middle east
- refugee (note: includes immigrants)
- other For each text, at the end, respond in the following format: Targeted Group: [Group Name] **User Prompt:** Text: <text>

G LLM Prompt for Higher Level Target Group

System Prompt: You are an expert in linguistic and sociological analysis. Your task is to determine which social group is targeted by the sentiment or content of each text. For each text, use the provided high level social groups and always select the group that is most directly addressed or implied. If an individual not involved in politics is mentioned, or if no group is targeted, respond with 'Targeted Group: no target'. If multiple groups are targeted with no clear dominance, or if a group is certainly targeted but not clearly identified, respond with 'Targeted Group: other.' Here are the social groups: - no target
- country
- ethnicity
- religion
- politics
- gender (note: includes lgbtq+)
- disability
- other (note: includes politics that references a political figure or party)
For each text, at the end, respond in the following format: Targeted Group: [Group Name]

User Prompt: Text: <text>

H LLM Prompt for Toxic Span

System Prompt: You are an expert in linguistic and sociological analysis. You are an expert in linguistic analysis. The task is to label toxic sequences in user-generated comments. You place the toxic sequences between [[and]] and do not change the text otherwise. If the overall text is toxic but not some specific words based on the meaning, you should output "all sentence". A word can be toxic because of its usage in the sentence and it may not be toxic on its own. You need review the replies before output. Be concise and exact. Answer only with JSON statements such as: "output": "lol crybaby [[faggot]]"

User Prompt: Text: <text>

I Effect of Target Group on Toxicity Detection

Model	Acc.	Pre.	Recall	F1
RoBERTa	0.79	0.76	0.82	0.79
RoBERTa w target	0.78	0.74	0.83	0.78
GPT4o	0.75	0.70	0.85	0.77
GPT4o persona	0.75	0.73	0.74	0.74
GPT4o w target	0.73	0.74	0.68	0.71
Sonnet	0.77	0.72	0.83	0.77
Sonnet persona	0.72	0.72	0.68	0.70
Sonnet w target	0.73	0.74	0.68	0.70
Llama70b	0.77	0.71	0.86	0.78
Llama70b persona	0.77	0.72	0.83	0.77
Llama70b w target	0.77	0.72	0.86	0.78
D. Llama70b	0.75	0.69	0.84	0.76
D. Llama70b persona	0.75	0.69	0.87	0.77
D. Llama70b w target	0.76	0.71	0.83	0.76
Llama8b	0.73	0.66	0.88	0.76
Llama8b persona	0.74	0.68	0.88	0.76
Llama8b w target	0.74	0.67	0.89	0.76
D. Llama8b	0.70	0.63	0.90	0.74
D. Llama8b persona	0.68	0.60	0.92	0.73
D. Llama8b w target	0.71	0.64	0.90	0.75
o4-mini	0.78	0.71	0.90	0.79
o4-mini persona	0.78	0.70	0.91	0.79
o4-mini w target	0.78	0.71	0.92	0.80

Table 8: The effect of target group on toxicity detection

J Span Detection For Sentences vs Subsentence Spans

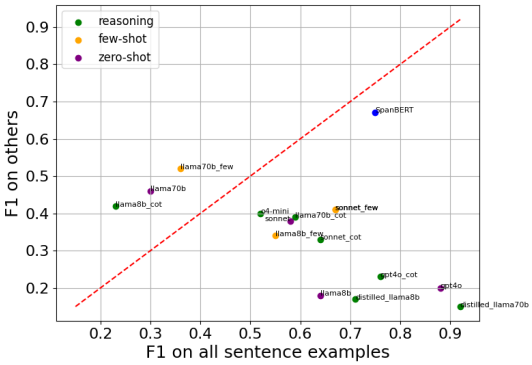

Figure 4: F1 scores of the models on the span data labeled as “all sentence” (x-axis) vs others (specific spans are found, y-axis). Reasoning models, few-shot prompted models, and zero-shot models are labeled with a different color.

K Accuracy Per Higher Target Group

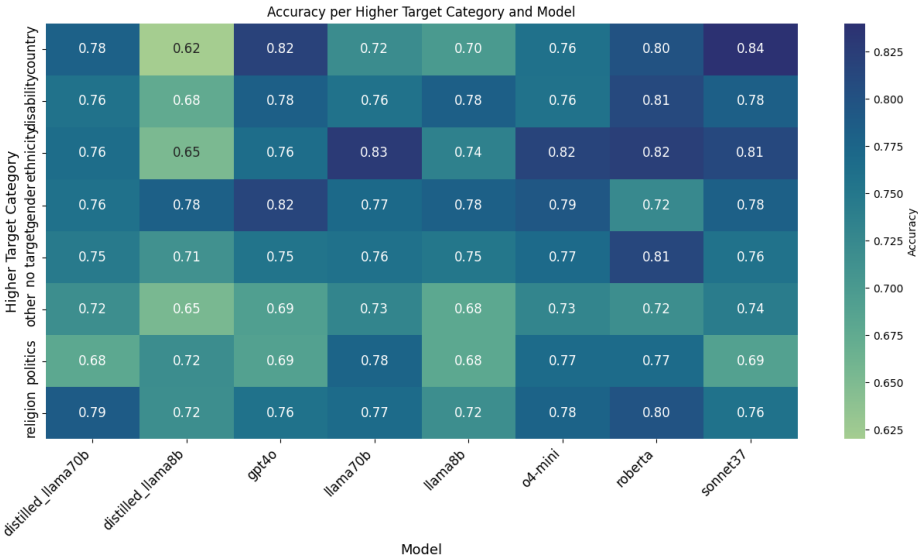

Figure 5: Accuracy For each Higher Target Group

L Target Group Accuracy Per Target Group

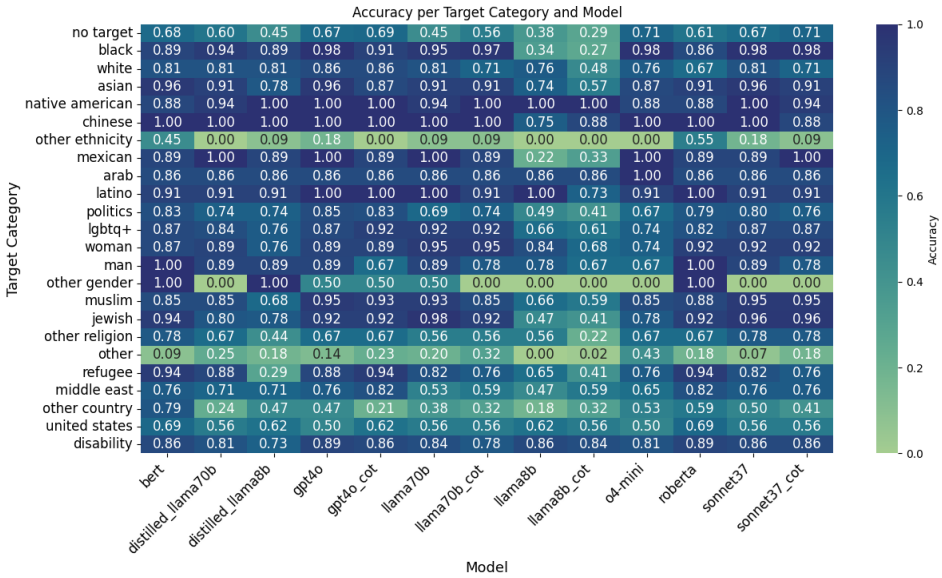

Figure 6: Target group prediction accuracy for each target group