

HateXScore: A Metric Suite for Evaluating Reasoning Quality in Hate Speech Explanations

Yujia Hu, Roy Ka-Wei Lee

Singapore University of Technology and Design
yujia_hu@sutd.edu.sg, roy_lee@sutd.edu.sg

Abstract

Hateful speech detection is a key component of content moderation, yet current evaluation frameworks rarely assess why a text is deemed hateful. We introduce HateXScore, a four-component metric suite designed to evaluate the reasoning quality of model explanations. It assesses (i) conclusion explicitness, (ii) faithfulness and causal grounding of quoted spans, (iii) protected group identification (policy-configurable), and (iv) logical consistency among these elements. Evaluated on six diverse hate speech datasets, HateXScore is intended as a diagnostic complement to reveal interpretability failures and annotation inconsistencies that are invisible to standard metrics like Accuracy or F1. Moreover, human evaluation shows strong agreement with HateXScore, validating it as a practical tool for trustworthy and transparent moderation.

Disclaimer: This paper contains sensitive content that may be disturbing to some readers.

1 Introduction

Hateful content remains a persistent challenge, with over 10% online posts estimated to contain hateful content (Fortuna and Nunes, 2018). In response, platforms are increasingly relying on automated approaches for content moderation (Vidgen et al., 2019; Mozafari et al., 2019; Hee et al., 2024). Although classification accuracy has improved, recent regulatory frameworks such as the EU Digital Services Act¹ underscore the growing demand for transparency in automated decisions.

In hate speech detection, such transparency is critical. Misclassifying harmful content risks perpetuating harm, while over-policing benign speech may suppress legitimate expression (ElSherief et al., 2021). Even accurate models can fail silently by relying on spurious features or ignoring implicit

hate. Moderation systems must therefore justify not just what decision was made, but why (Wang et al., 2023).

Existing explanation evaluation methods have assessed whether highlighted tokens align with the predicted label (Arras et al., 2017) or measure overlap with human-annotated rationales (Mathew et al., 2021; ElSherief et al., 2021). However, these approaches do not ask whether the explanation quotes the harmful span, whether removing it changes the prediction, or whether the explanation correctly identifies the targeted group.

These gaps are further compounded by the subjectivity of hate speech annotation (Sap et al., 2019). A model may produce a faithful, well-reasoned explanation that challenges an ambiguous or flawed label, but still be penalized by label-centric metrics. A model that highlights a slur and explains its harm may offer a more faithful interpretation than the label allows. This disconnect highlights the need for an evaluation framework that accounts not just for prediction correctness, but for the underlying reasoning. This issue is amplified when annotation disagreement is present, where LLMs may appear confident despite contested labels (Lu et al., 2025).

To address these challenges, we introduce HateXScore, a metric suite that evaluates explanation quality across four dimensions: explicitness of the conclusion, faithfulness of quoted spans with causal impact, identification of protected groups under configurable policies, and internal consistency among these elements. Figure 1 shows an application of HateXScore to an explanation generated by a large language model (LLM) for a sample from the HateXplain dataset (Mathew et al., 2021). The model’s prediction is deemed incorrect only because it disagrees with the dataset label. In contrast, HateXScore assigns a high score to the generated explanation, indicating that the model provides a coherent justification for its decision. In such cases,

¹https://www.eu-digital-services-act.com/Digital_Services_Act_Article_17.html

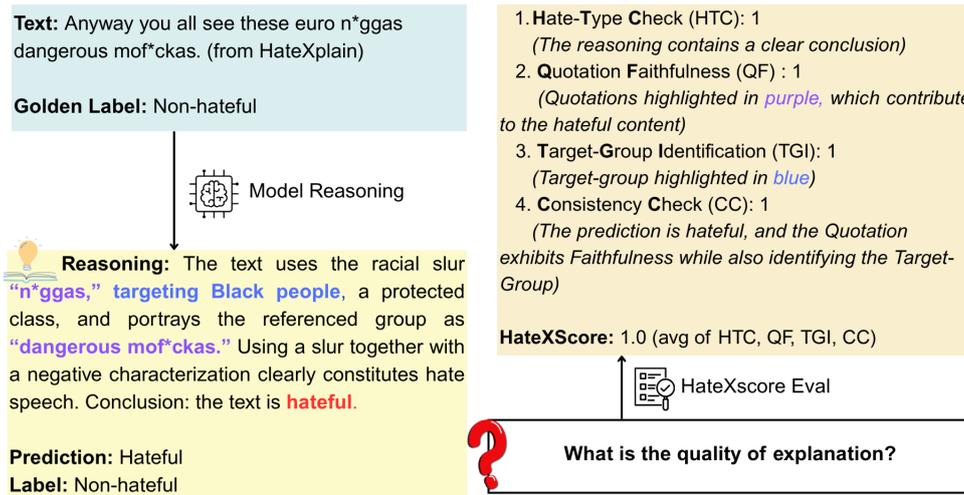


Figure 1: An example of a hate speech detection scored using HateXScore.

human moderators may concur with the model and view the gold label as debatable. This example highlights HateXScore as a diagnostic signal that can surface well-justified explanations under prediction–label mismatch, thereby helping identify potentially unreliable annotations in hate speech datasets. Accordingly, while Accuracy and F1 quantify agreement with dataset labels, HateXScore evaluates the quality of the model’s justification for its own prediction.

We summarize our contribution as follows: (i) We propose HateXScore, a novel suite of metrics to evaluate the quality of model explanations for hate speech detection. (ii) We conducted experiments to evaluate the explanations generated by models on six hate speech datasets, spanning English, Chinese, and Korean, enabling multilingual analysis. Our experiments demonstrated that HateXScore uncovers explanation failures and annotation inconsistencies that standard metrics overlook. (iii) We validate HateXScore through human evaluation and case studies, showing its alignment with human judgment and its capacity to surface hidden annotation disagreements.

2 Related work

2.1 Hate Speech Detection

Hate speech detection has received increasing attention due to the growth of content posted by users on social media (Davidson et al., 2017; Schmidt and Wiegand, 2017; Zampieri et al., 2019; Lin et al., 2021). Early systems employed lexicons and simple classifiers, such as SVMs or logistic regression (Talat and Hovy, 2016), which could detect explicit slurs but struggled with implicit or contextual

hate (Vidgen and Derczynski, 2020).

Recent advances use deep models like BERT (Mozafari et al., 2019) and RoBERTa (Liu et al., 2019), improving accuracy by capturing semantic and syntactic context (Abusaqer et al., 2025; Ghorbanpour et al., 2025). Work has also extended to multilingual (Das et al., 2024; Pamungkas and Patti, 2019; Basile et al., 2019; Xiao et al., 2024; Ng et al., 2024), multimodal (Lee et al., 2021, 2024) and fine-grained hate detection (e.g., racism, sexism, homophobia), reflecting the field’s growing sensitivity to hate’s varied forms. However, most work still relies on Accuracy or F1, giving little insight into model reasoning, motivating explanation-based evaluation.

2.2 Hate Speech Explanation

Efforts to evaluate explanations in hate speech models remain limited. Task fidelity approaches test alignment between highlighted tokens and predicted labels (Arras et al., 2017), while overlap metrics compare model spans with human rationales (Mathew et al., 2021; ElSherief et al., 2021). Data imbalance has also motivated augmentation approaches (Cao and Lee, 2020). While informative, these methods ignore whether models quote the harmful span, link it causally to predictions, or identify the targeted group. Subjectivity in annotations further complicates the evaluation. In HateXplain (Mathew et al., 2021), reclaimed slurs may be labeled non-hateful by some annotators and hateful by others. In ImplicitHate (ElSherief et al., 2021), indirect expressions lead to inconsistent labels. In such cases, a model may provide a well-reasoned explanation that conflicts with the majority label, something conventional metrics are ill-equipped to

capture.

Existing explanation frameworks also treat decision correctness and explanation quality separately (Piot and Parapar, 2024; Nghiem and Daumé III, 2024; Danilevsky et al., 2020). Recent work also shows that linguistically structured and interpretable explanations can speed up human moderation without sacrificing accuracy (Calabrese et al., 2024), reinforcing the importance of interpretable reasoning frameworks such as HateXScore. Approaches such as causal masking (DeYoung et al., 2019) test the importance of the explanation, but are rarely adapted to hate speech, and rarely account for identity group targeting, a core aspect of hateful intent. Recent work highlights reasoning-based explanations and LLMs in hate-speech detection. Roy et al. (2023) explore causal reasoning for hate speech classification, improving interpretability but without focusing on policy alignment or protected group identification. Yang et al. (2023) propose HARE, which emphasizes stepwise reasoning, sharing HateXScore’s focus on explanation structure but lacking cross-verification of explanations with target groups or measurement of overall consistency. Wang et al. (2023) evaluate GPT-generated explanations with human judgments but do not capture the multi-faceted causal and group recognition aspects that HateXScore quantifies. Balkir et al. (2022) offer necessity and sufficiency measures for causal token importance similar to HateXScore’s Quotation Faithfulness sub-metric, but without integrating policy-aware group detection or coherence checks. HateModerate (Zheng et al., 2023) aligns models with moderation policies but does not analyze explanation structure. HateXScore uniquely integrates causal validation, policy-aware group recognition, consistency checks, and explicit conclusions in one interpretable suite. This integration allows a comprehensive, policy-aware assessment of hate speech explanations that goes beyond prior approaches.

3 Metrics Design

HateXScore evaluates the quality of hate speech explanations across four complementary components. Each component targets a distinct aspect of reasoning, including conclusion explicitness, faithfulness and causal grounding of quoted spans, protected group identification (policy-configurable), and logical consistency.

Let T denote the input text, $y \in \{0, 1\}$ the label

(*non-hateful* or *hateful*), and \hat{y} the predicted label. The model also generates a natural-language explanation $E = E(T)$, from which we extract a set of quoted rationales $Q \subseteq T$. Let $p(T) \in [0, 1]$ be the model’s estimated probability that a given text T is hateful. We define an indicator function as follows:

$$\mathbb{I}(\cdot) = \begin{cases} 1, & \text{if the argument is true,} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

3.1 Hate-Type Check (HTC)

HTC assesses whether the explanation explicitly states a conclusion regarding the input text’s hateful or non-hateful nature. It is computed as:

$$HTC(E, \hat{y}) = \mathbb{I}(E \text{ contains a conclusion statement}) \quad (2)$$

This binary check ensures that explanations contain a definitive decision rather than ambiguous or missing judgments. The implementation leverages simple string matching over the explanation to detect presence of conclusive keywords such as “*text is hateful*” or “*text is non-hateful*”.

3.2 Quotation Faithfulness (QF)

QF jointly assesses (1) whether a non-trivial span from T is quoted, and (2) whether this span causally influences the prediction.

First, the explanation’s quoted spans $Q = q_1, q_2, \dots$ are extracted by overlapping normalized tokens between the explanation and the input text T . If the explanation trivially quotes the entire input text ($Q = \{T\}$), QF is set to 0 to prevent trivial or uninformative explanations. Then, We mask all spans in Q from T (with “[MASK]”, using both exact and fuzzy matching via `re`² and `fuzzysearch`³), and compute model probabilities on both original and masked texts (p_{orig}, p_{mask}). Then:

$$QF(T, Q, \hat{y}) = \begin{cases} 0, & \text{if } Q = \emptyset \text{ or } \{T\} \\ |p_{orig} - p_{mask}|, & \hat{y} = 1 \\ 1 - |p_{orig} - p_{mask}|, & \hat{y} = 0 \end{cases} \quad (3)$$

3.3 Target-Group Identification (TGI)

TGI checks whether the explanation correctly recognizes a protected or sensitive group (e.g., race, religion, gender) referenced in the hateful text. Our implementation includes a comprehensive, multi-lingual inventory of protected groups sourced from

²<https://docs.python.org/3/library/re.html>

³<https://github.com/taleinat/fuzzysearch>

the United Nations’ official definitions of hate targets, covering race, religion, gender, disability, and more, with lists available in multiple languages as released by the UN. In addition, several commonly used corporate policy lists (e.g., Meta, Twitter, YouTube, etc) are pre-integrated and selectable by the user. Users may also supply their own custom group lists in the form of a dictionary, enabling full configurability for specific research or regulatory requirements. For all experiments in this paper, we default to the United Nations Targets of Hate as the reference list. The full list can be found in Appendix A.6.

We then extract all n -grams (up to trigrams) from the explanation, using language-appropriate tokenization (spaCy, jieba or KoNLPy), and matching them in lemmatized form against the protected group list. We choose up to trigrams because shorter n -grams (unigrams and bigrams) capture single-word or phrase-level mentions, while trigrams allow detection of more specific multi-word expressions (e.g., “*African American people*”, “*transgender women community*”). For hateful samples, we further validate that the explanation ties the identified group to hateful context, using NER and part-of-speech patterns to check for direct relationships (e.g., verbs, adjectives). If a group is mentioned without connection to hateful content, TGI is set to 0.

Concretely, we define a predefined set G of protected group terms based on company policies. The detection involves extracting lemmatized n -grams (up to trigrams) set e from the explanation E and matching them against similarly processed target group lists G . Then

$$\text{TGI}(e, G) = \begin{cases} 1, & \text{if } e \cap G \neq \emptyset \\ 0, & \text{otherwise} \end{cases}. \quad (4)$$

3.4 Consistency Check (CC)

CC evaluates whether the model’s reasoning is logically coherent with respect to its decision, the causal evidence (QF), and the identification of protected groups (TGI). We treat CC as a diagnostic indicator rather than an independent contribution to explanation quality. Formally, CC verifies that a hateful prediction is supported by both (i) a meaningful quoted span with a strong causal effect and (ii) correct target-group identification. Conversely, a non-hateful prediction should generally lack such causal evidence or target reference. We generalize the criterion using a variable threshold $\tau \in [0, 1]$,

which is user-configurable:

$$\text{CC}(\hat{y}, \text{QF}, \text{TGI}) = \begin{cases} 1, & \hat{y} = \text{hateful}, \text{QF} \geq \tau, \text{TGI} = 1 \\ 0, & \hat{y} = \text{hateful}, \text{otherwise} \\ 1, & \hat{y} = \text{non-hate}, \text{QF} < \tau, \text{TGI} = 0 \\ 0, & \hat{y} = \text{non-hate}, \text{otherwise} \end{cases}. \quad (5)$$

To examine robustness, we perform sensitivity analysis by varying τ in the range $[0.1, 0.9]$ with increments of 0.1 in Appendix A.4. For each threshold, we recompute CC and the resulting HateXScore, then visualize the relationship between τ and the average metric across datasets and models. As shown in the sensitivity curves, CC and HateXScore both decrease smoothly as τ increases, reflecting that higher thresholds impose stricter consistency requirements. The ranking of models stays similar across $\tau \in [0.1, 0.3]$. Our goal is to reward even partially faithful quotations and to capture moderate alignment between explanation and decision, we adopt a relatively lenient cutoff of $\tau^* = 0.3$ for all main experiments. Complete threshold-sweep results are reported in Appendix A.4.

3.5 HateXScore: Overall Reasoning Quality

We compute the final HateXScore score as:

$$\text{HateXScore} = \frac{(\text{HTC} + \text{QF} + \text{TGI} + \text{CC})}{4} \quad (6)$$

HateXScore ranges from 0 to 1, integrating conclusion presence with quotation faithfulness, targeted group recognition, and logical consistency. Together, these dimensions provide a principled and interpretable metric for evaluating model reasoning in hate speech detection.

We use an unweighted average as a transparent default aggregation. The four components capture complementary requirements for a usable moderation rationale, and we therefore treat them as jointly necessary rather than optional. In our analysis, we report and interpret the four sub-metrics alongside the aggregate score. Practitioners may re-weight components to reflect application-specific policies, such as, emphasizing TGI for compliance. In this paper, we focus on equal weights for interpretability and ease of adoption.

3.6 Preprocessing and Tokenization

We implement all components in an automated pipeline, leveraging the spaCy⁴ library for lemmatization and named entity recognition (NER), regular expressions for span matching, and support for

⁴<https://spacy.io/>

flexible protected-group inventories for different platform or policy requirements. To ensure accurate and language-appropriate analysis, our pipeline supports flexible selection of tokenizers for text normalization and span extraction. For English and other European languages, we rely on the spaCy library for tokenization, lemmatization, and named entity recognition (NER). For Chinese-language content, as in the ToxiCN dataset, we use the jieba⁵ tokenizer, which is better suited to the segmentation requirements of Chinese text. For Korean, we used KoNLPy⁶ as tokenizer. The codebase allows users to specify the tokenizer most appropriate for their target language, and seamlessly integrates additional tokenizers for further language expansion.

4 Experiment Setup

4.1 Datasets

We evaluate models on six widely used hate speech datasets: HateXplain (Mathew et al., 2021), HateCheck (Röttger et al., 2020), Latent Hatred (ElShierief et al., 2021), HASOC (Mandl et al., 2020), ToxiCN (Lu et al., 2023) and KOLD (Jeong et al., 2022). These datasets span diverse linguistic contexts, target groups, and forms of hate, including both explicit and implicit expressions. Notably, ToxiCN focuses on Chinese-language content and KOLD on Korean content, enabling evaluation in non-English settings. There are 1,370 (HateXplain), 2,864 (Latent Hatred), 1,592 (HASOC), 3,728 (HateCheck), 2,411 (ToxiCN) and 4,000 (KOLD) samples. See Appendix A.1 for full dataset details.

4.2 Models

To examine how model architecture and scale influence explanation quality, we evaluate seven large language models from different families: GPT-4o (Achiam et al., 2023), LLaMA-8B (Dubey et al., 2024), Mistral-7B (Jiang et al., 2023), Qwen-7B (Bai et al., 2023), and the Gemma series (Team et al., 2024a,b) (Gemma-2B, Gemma-9B, Gemma-27B). Each model is prompted to produce a binary hatefulness label and a corresponding free-text explanation. All models were prompted using a standardized instruction, shown in Appendix A.7. Outputs are truncated at 512 tokens, and no task-specific fine-tuning is applied. This setup reflects real-world, 0-shot use of generative LLMs in moderation workflows.

⁵<https://github.com/fxsjy/jieba>

⁶<http://konlpy.org/ko/latest>

While our main experiments adopt the common 0-shot deployment setting, HateXScore is not tied to any specific prompting regime. We also provide a 2-shot robustness check in Appendix A.5, showing that the qualitative trends are consistent under light few-shot prompting.

4.3 Evaluation Metrics

We evaluate models using both traditional classification metrics and HateXScore.

4.3.1 Classification Metrics

Accuracy and macro-F1 are reported to capture classification performance. However, these metrics do not capture the quality of the models' explanation on their decisions.

4.3.2 HateXScore

We compute all four sub-metrics defined in Section 3 for every model-generated explanation. The final score, ranging from 0 to 1, reflects a model's ability to generate coherent, faithful, and complete justifications for its decisions.

4.4 Human Evaluation

To evaluate how well HateXScore aligns with human judgments, we conducted a human annotation study focused on English and Chinese datasets (HateXplain, Latent Hatred and ToxiCN) due to the availability of qualified bilingual annotators. We randomly sampled 600 GPT-4o-generated explanations (300 hateful, 300 non-hateful) from across our evaluation sets. Each instance included the input text, model prediction and explanation.

Three expert annotators who are fluent in English and Chinese, with experience in hate speech moderation, were asked to assess:

- Whether the explanation quotes relevant text and if the quoted text influences the prediction. A score of 1 was assigned when the explanation explicitly quotes a text segment that directly affects the hateful or non-hateful classification and clearly articulates the causal link. A score of 0 was given when the quoted content was missing or irrelevant.
- Whether the explanation identifies the correct targeted group. If the explanation mentions a protected group that matches the harmful context in the input, TGI = 1. When no group is identified or the reference is incorrect, TGI = 0. Annotators referenced the protected-group categories listed in Appendix A.6.

- Whether they think the text is hateful or non-hateful (0 = non-hateful, 1 = hateful)

We only manually annotate QF and TGI as they require semantic judgment. HTC is verified deterministically from the required conclusion line in the explanation, and CC is computed deterministically from \hat{y} , QF and TGI given τ . As a result, separately collecting human labels for HTC or CC would not yield an independent signal.

We report aggregated scores and inter-annotator agreement in Section 5.3. In Section 5.3, to directly assess alignment between HateXScore and human judgments, we threshold the model’s continuous QF/TGI scores using the same τ as in the metric definition, treat the binarized model outputs as an additional rater, and report Fleiss’ κ between the model and human annotators. This evaluation serves to validate HateXScore as a faithful proxy for human-aligned explanation quality.

4.5 Implementation Details

For generating explanations from the evaluated language models, we used greedy decoding with a temperature set to 0.0, which deterministically selects the most likely next token at each step, resulting in consistent outputs for closed-source models (GPT-4o). In contrast, for open-source models accessed via the OpenRouter API with two NVIDIA A40 GPUs, we utilized top-p sampling with a cumulative probability threshold $p = 0.9$ and a temperature of 0.7. All experiments were repeated three times, and the average values were reported.

5 Results and Analysis

5.1 Dataset-Level Analysis

We begin by analyzing explanation quality across datasets, grouped by the nature of hate expressed: *explicit* hate (HateXplain, HateCheck), *implicit* hate (Latent Hatred, HASOC) and non-English corpora (ToxicCN, KOLD). This contrast reveals how different types of linguistic aggression interact with sub-metrics in HateXScore, exposing patterns that conventional metrics fail to uncover.

Explicit Hate. Across datasets with explicit hate such as HateXplain and HateCheck, almost all models state a clear decision in their explanations, yet only a subset produce rationales that are both causally grounded and policy aligned. GPT-4o leads on both datasets with strong quotation faithfulness and high consistency, while Qwen 7B and Gemma 9B follow at a short distance. In contrast,

models like Mistral 7B and the smallest Gemma variant often quote irrelevant spans or fail to tie quoted content to the final decision, which suppresses their HateXScore despite competitive accuracy. These outcomes illustrate that explicit hate is easy to label but still demands careful grounding of evidence and target recognition to reach a high explanation score.

Implicit Hate. When hate is implicit or nuanced, as in Latent Hatred and HASOC, explanation quality separates the model families more sharply. HTC remains high because most systems format an explicit conclusion, but the quotation and consistency signals vary widely. On Latent Hatred, Qwen 7B reaches the strongest overall reasoning score among open models through reliable target group identification, yet its quotation faithfulness remains moderate, which indicates that part of the rationale is not fully causal. On HASOC, GPT-4o obtains the highest HateXScore through strong causal evidence yet still shows gaps in target identification, confirming that oblique attacks on identity are challenging even for high end systems. These findings stress that success on implicit hate requires both precise span selection and robust semantic linking between evidence and targeted groups.

Non-English Hate. On KOLD, GPT-4o delivers a very high HateXScore driven by near perfect target group identification and near perfect internal consistency. Gemma 9B reaches a similarly strong overall score with balanced quotation and consistency, while Qwen 7B ranks close behind and clearly ahead of several larger models. On ToxicCN, Qwen 7B achieves competitive overall reasoning with high target recognition and solid quotation, and GPT-4o maintains strong causal grounding but shows lower accuracy than on English tests. These results affirm that the metric suite operates effectively across different languages when appropriate tokenizers and group lexicons are used. In contrast, traditional metrics cannot reflect reasoning quality, whereas HateXScore captures whether models ground their judgments in contextually relevant expressions and culturally specific slurs.

5.2 Model-Level Analysis

We examine how model size and architecture affect explanation quality across datasets, using HateXScore and its sub-metrics. Table 1 reveals that while larger models tend to improve overall prediction accuracy, this improvement does not necessarily translate into better explanations. The Gemma se-

Dataset	Model	HTC	QF	TGI	CC	HateXScore	Accuracy	F1
HateXplain	GPT-4o	1.000	0.677	0.980	0.866	0.881	0.766	0.750
	LLaMA-8B	1.000	<u>0.640</u>	0.879	0.788	<u>0.827</u>	0.635	0.630
	Qwen-7B	1.000	<u>0.607</u>	0.888	<u>0.798</u>	0.823	<u>0.702</u>	<u>0.647</u>
	Mistral-7B	1.000	0.195	0.815	0.320	0.583	0.670	0.543
	Gemma-2b	1.000	0.294	0.676	0.294	0.566	0.310	0.556
	Gemma-9b	1.000	<u>0.447</u>	<u>0.887</u>	0.595	0.732	0.555	0.595
	Gemma-27b	1.000	0.365	<u>0.897</u>	0.630	0.723	0.571	0.599
Latent Hatred	GPT-4o	1.000	0.677	0.776	0.861	0.829	0.723	0.681
	LLaMA-8B	1.000	0.529	0.775	<u>0.809</u>	0.778	<u>0.649</u>	<u>0.623</u>
	Qwen-7B	1.000	<u>0.572</u>	0.836	0.780	<u>0.797</u>	0.587	0.611
	Mistral-7B	1.000	0.216	0.726	0.316	<u>0.565</u>	0.403	0.558
	Gemma-2b	1.000	0.275	0.781	0.410	0.617	0.433	0.570
	Gemma-9b	1.000	0.485	<u>0.819</u>	0.616	0.730	0.473	0.588
	Gemma-27b	1.000	0.443	0.802	0.653	0.725	0.485	0.594
HASOC	GPT-4o	1.000	0.850	0.527	0.949	0.832	0.562	0.320
	LLaMA-8B	0.997	0.604	0.403	<u>0.831</u>	0.709	0.594	0.420
	Qwen-7B	1.000	<u>0.657</u>	<u>0.588</u>	0.741	<u>0.747</u>	<u>0.616</u>	0.636
	Mistral-7B	1.000	0.385	0.452	0.284	0.530	0.572	<u>0.682</u>
	Gemma-2b	1.000	0.383	0.477	0.282	0.536	0.597	0.711
	Gemma-9b	1.000	0.460	0.590	0.683	0.683	0.606	0.592
	Gemma-27b	1.000	0.295	0.551	0.711	0.639	0.646	0.608
HateCheck	GPT-4o	1.000	0.813	0.889	0.900	0.901	<u>0.922</u>	0.947
	LLaMA-8B	0.995	<u>0.751</u>	0.858	<u>0.796</u>	<u>0.850</u>	0.927	<u>0.949</u>
	Qwen-7B	1.000	0.632	0.822	0.745	0.800	0.832	0.869
	Mistral-7B	1.000	0.156	0.807	0.240	0.551	0.853	0.825
	Gemma-2b	1.000	0.096	0.567	0.138	0.450	0.704	0.821
	Gemma-9b	1.000	0.554	<u>0.894</u>	0.685	0.783	0.838	0.895
	Gemma-27b	1.000	0.334	0.907	0.667	0.727	0.851	0.902
ToxicCN	GPT-4o	1.000	0.774	<u>0.923</u>	0.925	0.906	0.686	0.629
	LLaMA-8B	0.538	0.359	<u>0.460</u>	0.549	0.477	0.656	0.647
	Qwen-7B	0.892	<u>0.688</u>	0.824	<u>0.851</u>	<u>0.814</u>	0.677	0.715
	Mistral-7B	0.939	0.298	0.564	0.389	0.548	0.653	0.700
	Gemma-2b	0.917	0.371	0.941	0.636	0.716	0.457	0.730
	Gemma-9b	0.942	0.449	0.701	0.536	0.657	<u>0.698</u>	<u>0.740</u>
	Gemma-27b	<u>0.962</u>	0.489	0.908	0.733	0.733	0.730	0.747
KOLD	GPT-4o	1.000	0.792	<u>0.997</u>	0.980	0.942	0.741	0.657
	LLaMA-8B	0.466	0.352	0.845	0.572	0.559	0.530	0.371
	Qwen-7B	0.731	<u>0.676</u>	0.987	<u>0.890</u>	0.821	0.620	0.417
	Mistral-7B	0.643	<u>0.217</u>	0.357	<u>0.425</u>	0.411	0.552	0.470
	Gemma-2b	0.822	0.145	0.975	0.343	0.571	0.667	0.708
	Gemma-9b	<u>0.939</u>	0.612	0.983	0.803	<u>0.834</u>	<u>0.740</u>	0.771
	Gemma-27b	<u>0.923</u>	0.487	1.000	0.767	0.794	0.725	<u>0.734</u>

Table 1: Summary of HateXScore, sub-metrics and general metrics (Accuracy and F1 score) in different models for the evaluated datasets. Highest and second-highest scores are bold and underline, respectively.

ries (2B, 9B, 27B) illustrates the limits of scaling. Although Accuracy improves steadily with model size, gains in QF and CC remain marginal or even regress. For example, Gemma-2B achieves a QF of 0.294 on HateXplain, which increases only modestly to 0.365 for Gemma-27B. Similar stagnation is observed in Latent Hatred and HASOC, suggesting that scaling amplifies memorization rather than reasoning depth. In contrast, GPT-4o and Qwen-7B maintain high HateXScore values even in low-resource languages, confirming that their explanation quality is driven by reasoning capability and not dataset familiarity.

When analyzing multilingual performance, Ha-

teXScore proves especially revealing. In ToxicCN and KOLD, traditional metrics yield moderate F1 scores, yet HateXScore sharply differentiates models that truly understand cross-lingual hate expressions. GPT-4o achieves 0.942 on KOLD by producing explanations that correctly link racial or gendered expressions to corresponding protected groups, while smaller models with similar Accuracy fail to justify their predictions coherently.

5.3 HateXScore Alignment with Human Evaluation

To ensure a fair comparison between model and human judgments of QF, we first address the scale

Condition	Prefer \hat{y}	Prefer y
HateXScore > 0.5	0.748	0.252
HateXScore ≤ 0.5	0.227	0.773

Table 2: Annotator preferences when the model’s prediction (\hat{y}) conflicts with the label (y), i.e., $y \neq \hat{y}$, stratified by high vs. low HateXScore.

mismatch. The model produces real-valued QF scores in $[0, 1]$, whereas annotators assign binary judgments (0 or 1). We therefore normalize the model’s QF using the same threshold τ introduced in Section 3.4. Specifically, if $QF_{\text{model}} > \tau$, the value is normalized to 1; otherwise, 0. We then compute inter-rater agreement Fleiss’ κ by treating the binarized model outputs as an additional rater alongside the three human annotators. The resulting κ values are 0.783 (for QF) and 0.746 (for TGI), indicating substantial agreement between the model and human raters. For reference, inter-human agreement (only the three annotators) reached $\kappa = 0.758$ for QF and $\kappa = 0.731$ for TGI, showing that the model’s normalized judgments align with human consensus almost as closely as humans align with one another.

We further investigated whether HateXScore helps adjudicate disagreements between model predictions and dataset labels. When the predicted label differed from the label, annotators tended to trust the model’s prediction if its explanation received a high HateXScore. Specifically, as shown in Table 2, annotators preferred the model’s prediction 74.8% of the time when the HateXScore exceeded 0.5. In contrast, when the explanation scored below 0.5, the annotators sided with the dataset label in 77.3% of cases.

These results show that HateXScore functions as a diagnostic framework rather than a substitute for human evaluation. It detects well-reasoned disagreements where models challenge ambiguous labels and poorly reasoned agreements where labels are matched without justification. This diagnostic ability is crucial in hate-speech detection, where subjectivity and socio-linguistic variation cause inconsistent annotations.

6 Case Studies

6.1 Controversial Label Example

Even high-quality hate speech datasets sometimes contain ambiguous or mislabeled instances. These edge cases offer critical opportunities to evaluate whether HateXScore remains aligned with intu-

itive human judgment. Table 3 presents an example from Latent Hatred where the post states simply: “*hate white people.*” Despite the clearly offensive and group-targeted language, the label is marked as non-hateful, likely an annotation oversight. The model (Gemma-9B) predicts the text as hateful and provides a well-reasoned explanation that quotes the offensive span, identifies race as the protected group, and shows strong causal grounding, resulting in a high HateXScore of 1.0. This example shows that when a model prediction conflicts with the dataset label, a low HateXScore serves as a cue for human adjudication, prompting deeper review of the instance rather than treating it as simply incorrect under traditional metrics. More detailed examples are shown in Appendix A.2

Text	“ <i>hate white people.</i> ” (from <i>Latent Hatred</i>)
label	Non-hateful
Explanation (Gemma9b)	The text “ hate white people ” is explicitly expressing a negative sentiment towards a specific racial group , which qualifies as hate speech. The phrase directly calls for hatred against individuals based on their race, which is a clear violation of principles promoting respect and equality among different racial and ethnic groups.[...] It is a straightforward expression of racial animosity. Conclusion: hateful .
Metrics	HTC: 1 QF: 1 TGI: 1 CC: 1 HateXScore: 1

Table 3: An Example with Incorrect label. Target Group is highlighted by pink, Quoted Rationales are highlighted by yellow.

6.2 Perturbation Test and Adversarial Cases

To evaluate the robustness of HateXScore, we applied controlled perturbations to model-generated explanations. These edits included replacing the quoted rationale with irrelevant phrases, hallucinating protected groups not mentioned in the text, and omitting causal justifications. Table 4 shows an example involving the post “*Calling white trash is an insult to styrofoam! The Klan is despicable.*” The original explanation correctly highlights “*white trash*”, links it *White people*, and grounds the label decision causally, resulting in a high HateXScore of 1.0. After perturbation, the metric drops sharply. Substituting the rationale with “*white bubble*” (Perturbed 1) eliminates overlap and causal grounding (QF = CC = 0), yielding a score of 0.5. In Perturbed 2, the explanation retains the slur but fails to identify the target group, reducing TGI and CC, and lowering the score to 0.5. These cases show that

Text	<i>Calling white trash is an insult to styrofoam! The Klan is despicable.</i>
label	Non-hateful
Original	The phrase "white trash" is a well-known derogatory slur directed at White people . Using a racial slur to demean a protected group constitutes hate speech. Calling "the Klan" despicable is criticism of an extremist organization. In conclusion, the text is hateful.
Metrics	HTC: 1 QF: 1 TGI: 1 CC: 1 HateXScore: 1
Perturbed 1	The phrase "white bubble" is a well-known derogatory slur directed at White people . Using a racial slur to demean a protected group constitutes hate speech. In conclusion, the text is hateful.
Metrics	HTC: 1 QF: 0 TGI: 1 CC: 0 HateXScore: 0.5
Perturbed 2	The phrase "white trash" is a well-known derogatory slur. Calling "the Klan" despicable is criticism of an extremist organization. In conclusion, the text is hateful.
Metrics	HTC: 1 QF: 1 TGI: 0 CC: 0 HateXScore: 0.5

Table 4: Original vs. Perturbed explanation and resulting sub-metrics. Target Group is highlighted by pink . Quoted Rationales are highlighted by yellow .

HateXScore penalizes explanations that break coherence, misattribute targets, or lose causal clarity, even when the predicted label remains the same.

7 Conclusion

We introduced HateXScore, a suite of metrics for evaluating LLM-generated explanations in hate speech detection. Beyond label accuracy, it provides a principled, multifaceted assessment of explanatory quality. Experiments on six datasets and seven models show that HateXScore reveals interpretability failures and annotation inconsistencies overlooked by standard metrics. Notably, when model predictions diverge from labels, higher HateXScore values align with human judgments, underscoring its effectiveness in flagging potentially ambiguous or controversial annotations for further human adjudication.

Limitations

Despite its advantages, HateXScore faces certain limitations. It relies on automated methods for identifying and masking rationale text, which can be prone to errors, particularly in highly nuanced or figurative contexts. Moreover, evaluating whether

a group is truly attacked remains an inherently subjective judgment, albeit mitigated by structured target-group inventories. In addition, We have been experimenting with English, Chinese and Korean, but more languages are still to be expanded. Finally, HateXScore currently focuses on evaluating static, text-based explanations; future work could explore dynamic or multimodal reasoning (e.g., image-text hate speech) and human–AI co-evaluation frameworks to better capture the evolving nature of online hate.

For instances attacking multiple protected groups, we currently mark TGI as successful if the explanation identifies any clearly targeted group. This simplifies annotation but does not capture target coverage (partial or full), which can overstate performance on multi-target examples. Future work can adopt graded coverage labels or set-based scoring over targeted groups.

Ethical Statement

This research focuses on developing and evaluating HateXScore for detecting and explaining hateful or offensive language. Our primary objective is to highlight gaps in current hate-speech detection systems and propose methods (via HateXScore) that improve both classification and explanation quality. We do not seek to facilitate censorship or suppress legitimate free expression. Instead, by revealing vulnerabilities and inconsistencies in detection pipelines, we hope to support more accurate, transparent, and fair content moderation practices.

Consistent with the approach in hateful-language research, we rely on established datasets (e.g., HateXplain, HateCheck, Latent Hatred, HASOC, ToxiCN and KOLD). No additional personal or private data was collected for this study. As with other public datasets of toxic or hateful content, all resources are for scientific research only. ToxiCN is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0). HateCheck and HateXplain are under cc-by-4.0 license.

Examples throughout this paper necessarily include hateful and offensive terms, which may be disturbing or harmful to readers and annotators alike. We have attempted to minimize repeated exposure to such text by limiting the dataset size for examples and advising annotators on content warnings.

Our methods and code are designed to be trans-

parent, allowing researchers, practitioners, and regulatory bodies to understand, audit, and improve our approach. We welcome inquiries, critiques, and collaborations that aim to further the responsible use of these technologies.

By adhering to these principles, we strive to advance the scientific study of hateful language detection and the interpretability of machine learning models in a manner that respects human dignity, privacy, and well-being.

Acknowledgments

This research project is supported by the National Research Foundation, Singapore under its National Large Language Models Funding Initiative, (AISG Award No: AISG-NMLP-2024-004), Ministry of Education, Singapore, under its MOE Academic Research Fund Tier 2 (Award No: T2EP20222-0036). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the National Research Foundation and Ministry of Education, Singapore.

References

- Mahmoud Abusager, Jamil Saquer, and Hazim Shatnawi. 2025. Efficient hate speech detection: Evaluating 38 models from traditional methods to transformers. In *Proceedings of the 2025 ACM Southeast Conference*, pages 203–214.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Leila Arras, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2017. Explaining recurrent neural network predictions in sentiment analysis. *arXiv preprint arXiv:1706.07206*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, and 1 others. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Esma Balkir, Isar Nejadgholi, Kathleen C Fraser, and Svetlana Kiritchenko. 2022. Necessity and sufficiency for explaining text classifiers: A case study in hate speech detection. *arXiv preprint arXiv:2205.03302*.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In *Proceedings of the 13th international workshop on semantic evaluation*, pages 54–63.
- Agostina Calabrese, Leonardo Neves, Neil Shah, Maarten Bos, Björn Ross, Mirella Lapata, and Francesco Barbieri. 2024. Explainability and hate speech: Structured explanations make social media moderators faster. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 398–408, Bangkok, Thailand. Association for Computational Linguistics.
- Rui Cao and Roy Ka-Wei Lee. 2020. Hategan: Adversarial generative-based data augmentation for hate speech detection. ACL.
- Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen. 2020. A survey of the state of explainable ai for natural language processing. *arXiv preprint arXiv:2010.00711*.
- Susmita Das, Arpita Dutta, Kingshuk Roy, Abir Mondal, and Arnab Mukhopadhyay. 2024. A survey on automatic online hate speech detection in low-resource languages. *arXiv preprint arXiv:2411.19017*.
- Thomas Davidson, Dana Warmusley, 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, pages 512–515.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C Wallace. 2019. Eraser: A benchmark to evaluate rationalized nlp models. *arXiv preprint arXiv:1911.03429*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- 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. *arXiv preprint arXiv:2109.05322*.
- Paula Fortuna and Sérgio Nunes. 2018. A survey on automatic detection of hate speech in text. *Acm Computing Surveys (Csur)*, 51(4):1–30.
- Faeze Ghorbanpour, Daryna Dementieva, and Alexander Fraser. 2025. Can prompting llms unlock hate speech detection across languages? a zero-shot and few-shot study. *arXiv preprint arXiv:2505.06149*.
- Ming Shan Hee, Shivam Sharma, Rui Cao, Palash Nandi, Preslav Nakov, Tanmoy Chakraborty, and Roy Ka-Wei Lee. 2024. Recent advances in online hate speech moderation: Multimodality and the role of

- large models. *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4407–4419.
- Younghoon Jeong, Juhyun Oh, Jongwon Lee, Jaimeen Ahn, Jihyung Moon, Sungjoon Park, and Alice Oh. 2022. **KOLD: Korean offensive language dataset**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10818–10833, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, and 1 others. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Dong-Ho Lee, Hyundong Cho, Woojeong Jin, Jihyung Moon, Sungjoon Park, Paul Röttger, Jay Pujara, and Roy Ka-Wei Lee. 2024. Improving covert toxicity detection by retrieving and generating references. In *Proceedings of the 8th Workshop on Online Abuse and Harms (WOAH 2024)*, pages 266–274.
- Roy Ka-Wei Lee, Rui Cao, Ziqing Fan, Jing Jiang, and Wen-Haw Chong. 2021. Disentangling hate in online memes. In *Proceedings of the 29th ACM international conference on multimedia*, pages 5138–5147.
- Ken-Yu Lin, Roy Ka-Wei Lee, Wei Gao, and Wen-Chih Peng. 2021. Early prediction of hate speech propagation. In *2021 International Conference on Data Mining Workshops (ICDMW)*, pages 967–974. IEEE.
- 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. *arXiv preprint arXiv:1907.11692*.
- Junyu Lu, Kai Ma, Kaichun Wang, Kelaiti Xiao, Roy Ka-Wei Lee, Bo Xu, Liang Yang, and Hongfei Lin. 2025. Is llm an overconfident judge? unveiling the capabilities of llms in detecting offensive language with annotation disagreement. *arXiv preprint arXiv:2502.06207*.
- Junyu Lu, Bo Xu, Xiaokun Zhang, Changrong Min, Liang Yang, and Hongfei Lin. 2023. **Facilitating fine-grained detection of Chinese toxic language: Hierarchical taxonomy, resources, and benchmarks**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16235–16250, Toronto, Canada. Association for Computational Linguistics.
- Thomas Mandl, Sandip Modha, Anand Kumar M, and Bharathi Raja Chakravarthi. 2020. Overview of the hasoc track at fire 2020: Hate speech and offensive language identification in tamil, malayalam, hindi, english and german. In *Proceedings of the 12th annual meeting of the forum for information retrieval evaluation*, pages 29–32.
- 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, pages 14867–14875.
- Marzieh Mozafari, Reza Farahbakhsh, and Noel Crespi. 2019. A bert-based transfer learning approach for hate speech detection in online social media. In *International conference on complex networks and their applications*, pages 928–940. Springer.
- Ri Chi Ng, Nirmalendu Prakash, Ming Shan Hee, Kenny Tsu Wei Choo, and Roy Ka-Wei Lee. 2024. Sghatecheck: Functional tests for detecting hate speech in low-resource languages of singapore. *arXiv preprint arXiv:2405.01842*.
- Huy Nghiem and Hal Daumé III. 2024. Hatecot: An explanation-enhanced dataset for generalizable offensive speech detection via large language models. *arXiv preprint arXiv:2403.11456*.
- Endang Wahyu Pamungkas and Viviana Patti. 2019. Cross-domain and cross-lingual abusive language detection: A hybrid approach with deep learning and a multilingual lexicon. In *Proceedings of the 57th annual meeting of the association for computational linguistics: Student research workshop*, pages 363–370.
- Paloma Piot and Javier Parapar. 2024. Decoding hate: Exploring language models’ reactions to hate speech. *arXiv preprint arXiv:2410.00775*.
- Paul Röttger, Bertram Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet B Pierrehumbert. 2020. Hatecheck: Functional tests for hate speech detection models. *arXiv preprint arXiv:2012.15606*.
- Sarthak Roy, Ashish Harshvardhan, Animesh Mukherjee, and Punyajoy Saha. 2023. Probing llms for hate speech detection: strengths and vulnerabilities. *arXiv preprint arXiv:2310.12860*.
- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. 2019. Social bias frames: Reasoning about social and power implications of language. *arXiv preprint arXiv:1911.03891*.
- Anna Schmidt and Michael Wiegand. 2017. A survey on hate speech detection using natural language processing. In *Proceedings of the fifth international workshop on natural language processing for social media*, pages 1–10.
- Zeerak Talat and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of the NAACL student research workshop*, pages 88–93.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, and 1 others. 2024a. Gemma: Open

models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, and 1 others. 2024b. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.

Bertie Vidgen and Leon Derczynski. 2020. Directions in abusive language training data, a systematic review: Garbage in, garbage out. *Plos one*, 15(12):e0243300.

Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. Challenges and frontiers in abusive content detection. In *Proceedings of the third workshop on abusive language online*. Association for Computational Linguistics.

Han Wang, Ming Shan Hee, Md Rabiul Awal, Kenny Tsu Wei Choo, and Roy Ka-Wei Lee. 2023. Evaluating gpt-3 generated explanations for hateful content moderation. *arXiv preprint arXiv:2305.17680*.

Yunze Xiao, Yujia Hu, Kenny Tsu Wei Choo, and Roy Ka-wei Lee. 2024. Toxicloackn: Evaluating robustness of offensive language detection in chinese with cloaking perturbations. *arXiv preprint arXiv:2406.12223*.

Yongjin Yang, Joonkee Kim, Yujin Kim, Namgyu Ho, James Thorne, and Se-Young Yun. 2023. Hare: Explainable hate speech detection with step-by-step reasoning. *arXiv preprint arXiv:2311.00321*.

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. *arXiv preprint arXiv:1902.09666*.

Jiangrui Zheng, Xueqing Liu, Guanqun Yang, Mirazul Haque, Xing Qian, Ravishka Rathnasuriya, Wei Yang, and Girish Budhrani. 2023. Hatemoderate: Testing hate speech detectors against content moderation policies. *arXiv preprint arXiv:2307.12418*.

A Appendix

A.1 Dataset Statistics

Table 5 shows the statistical distributions of the datasets used in our experiments. Minimal preprocessing was applied (e.g., lowercasing), and each model used its default tokenizer to ensure compatibility.

A.2 Case Study Examples

To illustrate how HateXScore can reveal nuanced explanation behavior beyond simple classification metrics, we focus on several borderline examples

Dataset	# Hateful	#Non-hateful	# Total
HateXplain	688	682	1,370
Latent Hatred	1,092	1,772	2,864
HASOC	807	785	1,592
HateCheck	2,563	1,165	3,728
ToxicCN	1,274	1,137	2,411
KOLD	2,000	2,000	4,000

Table 5: Statistical information for datasets used in our experiments

Text: [The contents from datasets]
Reasoning: [The analysis from models]
Questions:
1. What is the score of Quotation Faithfulness (QF)? (0-1)
Answer:
2. What is the score of Target-Group Identification (TGI)? (0-1)
Answer:
3. Decision on Hatefulness (0 = non-hateful, 1 = hateful)
Answer:

Table 6: The Instructions Given To Annotators

from different datasets, as showed in Table 7, illustrating how classification correctness (HTC) and explanation quality (HateXScore) may sometimes conflict. These cases underscore the importance of considering both the label outcome and the underlying rationale, especially in cases with ambiguous or high disagreement.

A.2.1 Misclassified Label, Robust Explanation

The first example, “it ain’t even noon yet and niggas lying.” is labeled non-hateful in HateXplain. Nonetheless, the model predicts hateful, citing the racial slur and explaining its derogatory connotation. Although the models’ prediction disagrees with the dataset annotation, the explanation itself (QF=1, TGI=1) clearly identifies a protected group and shows a significant drop in probability when the slur is masked (QF close to 1). Consequently, the overall HateXScore remains high (0.80). Such a case casts doubt on the “non-hateful” label: it is plausible that some annotators considered the usage of “niggas” as in-group slang or contextually benign, whereas the model treats it as definitively hateful. Closer inspection suggested that the annotators themselves might have disagreed on whether the language was truly hateful, underscoring the possibility that the model’s explanation exceeded the majority-vote label in clarity or reliability.

A.2.2 Correct Label, Weak Explanation

In the third example (“cut immigration by at least 75%! raise does not do enough!”), the label is hateful, presumably reflecting annotators’ consen-

Text	Label	Pred.	Reasoning	Metrics
“It ain’t even noon yet and niggas lying.”	Non-hate	Hateful	To analyze the provided text, we need to consider both the explicit and implicit elements of hate speech. 1. Explicit hate speech: The term “niggas” is a racially charged term that has historically been used as a derogatory term for black individuals. Its use can be considered explicit hate speech, as it directly targets a racial group with a term that carries a significant amount of historical and social weight. The context in which it is used here suggests a negative connotation, as it is paired with the phrase “lying”, which implies deceitfulness. 2. Implicit hate speech: Implicitly, the use of such a term can perpetuate stereotypes about black individuals, suggesting that they are inherently dishonest or untrustworthy. This reinforces negative societal views and can contribute to a culture of racism and discrimination. The phrase “it ain’t even noon yet” could imply that the speaker is surprised or frustrated by the behavior of the individuals they are referring to, which can further imply a judgment based on race. 3. Context and tone: The overall tone of the text appears to express frustration or disbelief, which can be indicative of a broader negative sentiment towards the group being referenced. The casual nature of the language also suggests a normalization of the derogatory term, which can contribute to a culture of hate. In conclusion, the text contains explicit hate speech through the use of a racially derogatory term and implicitly reinforces negative stereotypes about a racial group. Conclusion: The text is hateful.	HTC: 1.0 QF: 1.0 TGI: 1.0 CC: 1.0 HateXScore: 1.0
“hate white people.”	Non-hate	Hateful	The text “hate white people” is explicitly expressing a negative sentiment towards a specific racial group, which qualifies as hate speech. The phrase directly calls for hatred against individuals based on their race, which is a clear violation of principles promoting respect and equality among different racial and ethnic groups. Upon analyzing the text, we find that it lacks ambiguity or irony. It straightforwardly states a desire to hate a particular group. This explicit declaration of hatred is harmful as it promotes division and animosity based on race, which can lead to further discrimination and social conflict. Given the clear and direct nature of the statement, it does not require interpretation or context to understand its hateful intent. It is a straightforward expression of racial animosity. Conclusion: hateful	HTC: 1.0 QF: 1.0 TGI: 1.0 CC: 1.0 HateXScore: 1.0
“Cut immigration by at least 75%! Raise does not do enough !”	Hateful	Non-hate	Analysis: The text expresses a strong opinion about immigration, advocating for a significant reduction in immigration levels. while it does not contain overtly hateful language or slurs, the statement can imply a negative stance towards immigrants or immigration as a whole. This could reflect a belief that immigrants are undesirable or problematic, which can be seen as a form of implicit bias against immigrant communities. However, the text does not explicitly dehumanize or attack any specific group, nor does it use derogatory language. the phrase “raise does not do enough” is somewhat ambiguous and does not provide clear context about what is being raised or what is insufficient. This could refer to policies or measures related to immigration, but without further context, it is difficult to assess its implications fully. Overall, while the text expresses a strong anti-immigration sentiment, it does not cross the line into hate speech. it reflects a political opinion that may be controversial or divisive, but it does not contain the elements typically associated with hate speech. Conclusion: non-hateful	HTC: 1.0 QF:0.0 TGI: 1.0 CC: 0.0 HateXScore: 0.5

Table 7: Several Borderline Examples where the model prediction contradicts the label. Target Group is highlighted by pink, Quoted Rationales are highlighted by yellow.

sus that targeting “immigration” is a discriminatory statement. The model, however, predicts non-hateful. Its explanation references the phrase “cut immigration,” but fails to clarify why such a demand may be hateful or to specify any protected group under attack (QF=1, but TGI=0). Despite matching the official label incorrectly, the explanation also scores poorly on reasoning components (e.g., CC = 0), resulting in a lower overall HateXScore (0.50). This example demonstrates that even though the model might show acceptable performance on standard metrics in other samples, it can produce incomplete rationales where content is arguably hateful. HateXScore makes that gap in explanation clarity visible.

These examples highlight how relying solely on Accuracy or F1 can misrepresent a model’s behavior, particularly in datasets containing ambiguous, context-dependent, or potentially incorrect labels. While HTC offers the conclusion of prediction, the other HateXScore components (QF, TGI, CC) can give credit to a faithful explanation. Conversely, a model might match the label but fail to justify it in terms of the offending snippet or protected group, leading to a subpar HateXScore. Overall, these case studies illustrate two core insights: first, classification correctness alone cannot capture explanation completeness and causal relevance; second, labels in hate speech corpora can themselves be flawed or disputed.

A.3 Instruction for Human Evaluation

Table 6 shows the instructions we gave to the annotator.

A.4 Sensitivity of the Consistency Threshold (τ)

We evaluated the robustness of the Consistency Check (CC) and HateXScore metrics by sweeping τ from 0.1 to 0.9 in steps of 0.1. Figures 2 show the average HateXScore for all models across datasets. As expected, stricter thresholds (larger τ) reduce both CC and HateXScore values, since fewer samples satisfy the consistency condition. The model ranking, however, remains largely stable for $\tau \leq 0.3$, indicating that the metrics are robust in the low-threshold region.

A.5 Few-shot robustness (2-shot prompting)

Our main experiments focus on the widely used 0-shot deployment setting. Since HateXScore is a metric suite applied to model explanations, it is in

Instructions:

Hate speech is defined as any kind of communication in speech, writing, or behaviour that attacks or uses pejorative or discriminatory language with reference to a group or its members on the basis of who they are.

Using the above definition, please decide if the following text is hateful or non-hateful. Provide the explanation for your classification in English/Chinese/Korean, followed by the conclusion (Hateful/Non-hateful). In your analysis, consider both explicit and implicit forms of hate speech. Ensure your reasoning avoids over-reliance on keywords and maintains consistency throughout the analysis.

Text:

{text}

Table 8

principle agnostic to prompting regimes. We therefore conduct a small robustness check to verify that the qualitative trends reported in the main paper remain consistent under light few-shot prompting.

We evaluate the same datasets and models as in Section 4.2, but prepend two demonstrations with one hateful and one non-hateful to the prompt used to elicit predictions and explanations. We keep all other settings identical to the main experiments.

Table 9 reports the 2-shot results. Overall, the relative patterns are consistent with the zero-shot setting: models that achieve higher HateXScore under zero-shot prompting also tend to do so under 2-shot prompting, and the separation between label-based performance (Accuracy/F1) and explanation-quality signals (HateXScore and sub-metrics) remains evident across datasets.

A.6 Target Group List

Table 10 - Table 13 shows the default target group lists from the UN ⁷, Meta ⁸, YouTube ⁹, and Twitter ¹⁰ that we have embedded in our code.

⁷<https://www.un.org/en/hate-speech/impact-and-prevention/targets-of-hate>

⁸<https://transparency.meta.com/en-gb/policies/community-standards/hateful-conduct/>

⁹<https://support.google.com/youtube/answer/2801939>

¹⁰<https://help.x.com/en/rules-and-policies/hateful-conduct-policy>

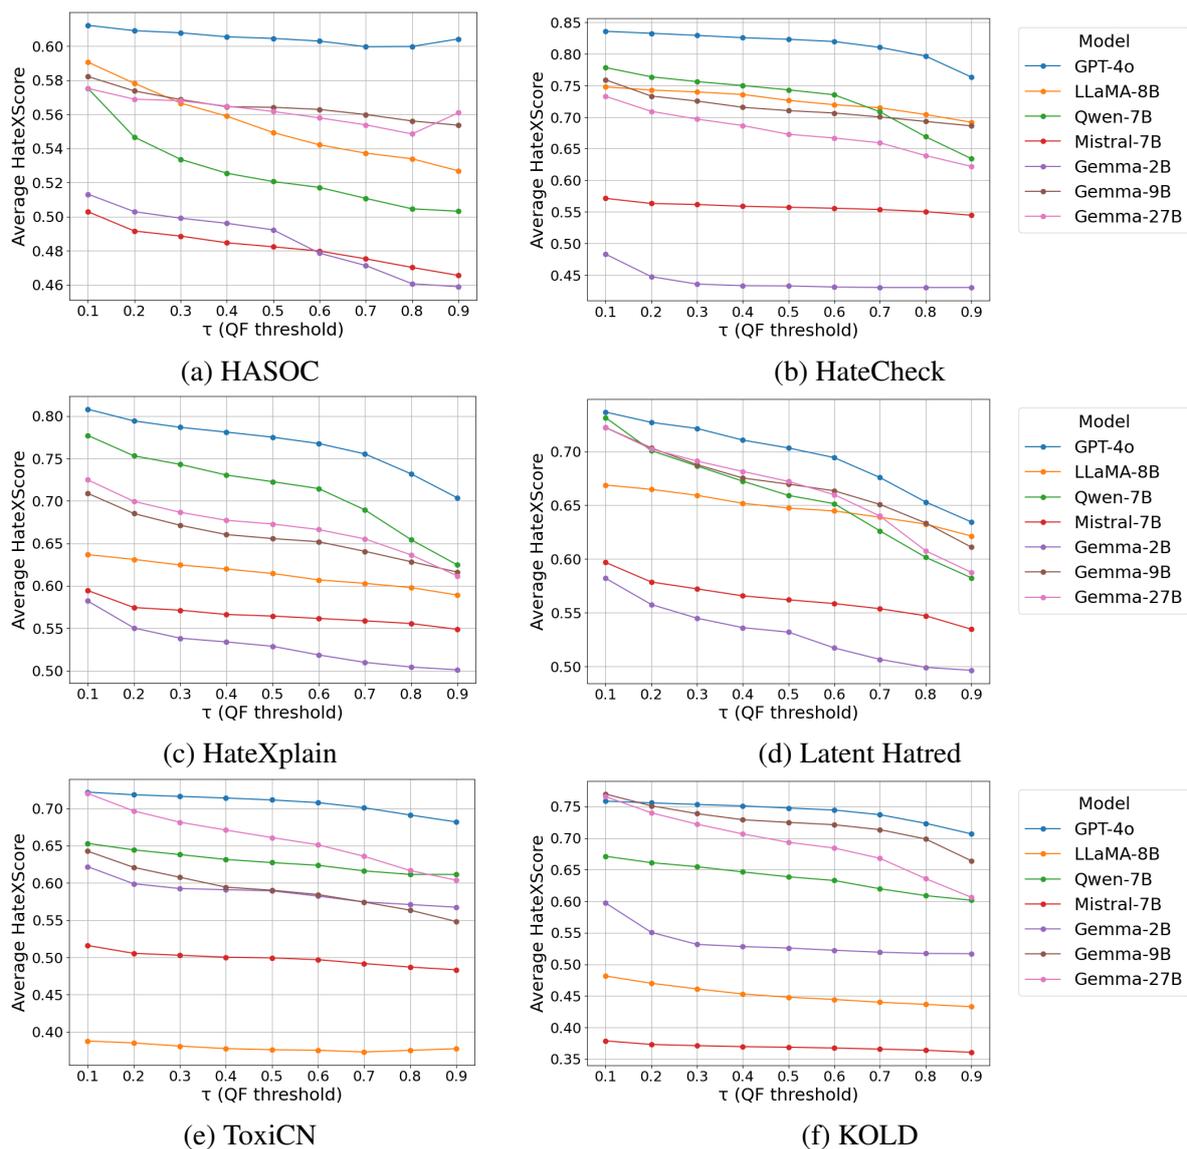


Figure 2: Sensitivity of HateXScore to the Quotation Faithfulness threshold τ across six datasets. Each subplot shows the average HateXScore for all models as τ varies from 0.1 to 0.9. Higher curves indicate stronger consistency between explanation and decision.

A.7 LLMs Prompt Formulations

Table 8 details the prompt templates used for generating model outputs across all experiments. Each prompt includes a standardized definition of hate speech, followed by an instruction for the model to classify the text as hateful or non-hateful while providing a justification.

Dataset	Model	HTC	QF	TGI	CC	HateXScore	Accuracy	F1
HateXplain	GPT-4o	1.000	0.695	0.951	0.914	0.890	0.890	0.867
	LLaMA-8B	0.791	0.292	0.405	0.415	0.477	0.650	0.370
	Qwen-7B	1.000	0.620	0.926	0.841	0.846	0.758	0.737
	Mistral-7B	1.000	0.323	0.913	0.350	0.651	0.544	0.598
	Gemma-2b	1.000	0.288	0.820	0.454	0.640	0.647	0.675
	Gemma-9b	1.000	0.475	0.873	0.609	0.740	0.546	0.564
	Gemma-27b	1.000	0.471	0.882	0.694	0.762	0.713	0.718
Latent Hatred	GPT-4o	1.000	0.693	0.928	0.972	0.898	0.748	0.585
	LLaMA-8B	0.834	0.270	0.475	0.406	0.496	0.623	0.357
	Qwen-7B	1.000	0.643	0.905	0.887	0.858	0.694	0.605
	Mistral-7B	1.000	0.392	0.848	0.493	0.685	0.638	0.646
	Gemma-2b	1.000	0.327	0.778	0.561	0.667	0.632	0.643
	Gemma-9b	1.000	0.522	0.794	0.656	0.742	0.559	0.597
	Gemma-27b	1.000	0.556	0.802	0.793	0.788	0.694	0.668
HASOC	GPT-4o	1.000	0.885	0.851	0.996	0.933	0.504	0.057
	LLaMA-8B	0.810	0.387	0.425	0.565	0.547	0.503	0.139
	Qwen-7B	1.000	0.780	0.863	0.943	0.897	0.536	0.189
	Mistral-7B	0.999	0.335	0.763	0.497	0.659	0.629	0.506
	Gemma-2b	1.000	0.479	0.376	0.818	0.668	0.592	0.426
	Gemma-9b	0.997	0.537	0.475	0.705	0.679	0.565	0.428
	Gemma-27b	1.000	0.604	0.413	0.938	0.739	0.539	0.222
HateCheck	GPT-4o	1.000	0.788	0.968	0.967	0.931	0.961	0.972
	LLaMA-8B	0.807	0.372	0.522	0.457	0.539	0.530	0.527
	Qwen-7B	1.000	0.679	0.933	0.902	0.878	0.867	0.904
	Mistral-7B	1.000	0.464	0.954	0.512	0.734	0.879	0.919
	Gemma-2b	1.000	0.135	0.951	0.220	0.577	0.870	0.913
	Gemma-9b	1.000	0.756	0.892	0.814	0.865	0.873	0.913
	Gemma-27b	1.000	0.525	0.924	0.760	0.802	0.909	0.938
ToxicCN	GPT-4o	1.000	0.774	0.963	0.948	0.921	0.640	0.511
	LLaMA-8B	0.657	0.098	0.515	0.470	0.435	0.495	0.264
	Qwen-7B	0.979	0.723	0.912	0.905	0.880	0.635	0.511
	Mistral-7B	1.000	0.326	0.345	0.498	0.546	0.746	0.762
	Gemma-2b	0.997	0.371	0.358	0.509	0.559	0.679	0.709
	Gemma-9b	0.968	0.393	0.630	0.548	0.635	0.668	0.719
	Gemma-27b	0.987	0.487	0.910	0.766	0.787	0.695	0.701
KOLD	GPT-4o	1.000	0.772	1.000	0.977	0.937	0.677	0.529
	LLaMA-8B	0.502	0.274	0.758	0.471	0.501	0.505	0.231
	Qwen-7B	0.840	0.614	0.993	0.864	0.828	0.572	0.264
	Mistral-7B	1.000	0.410	0.613	0.406	0.608	0.838	0.837
	Gemma-2b	0.998	0.337	0.521	0.592	0.612	0.784	0.798
	Gemma-9b	0.977	0.608	0.849	0.758	0.798	0.788	0.790
	Gemma-27b	0.989	0.586	0.999	0.843	0.854	0.780	0.768

Table 9: 2-shot prompting robustness check. Macro-averaged HTC/QF/TGI/CC/HateXScore and Accuracy/Macro-F1 across datasets (same evaluation pipeline as Table 1)

Policy	Category	Target Group
	National	afghanistan, albania, algeria, andorra, angola, antigua and barbuda, argentina, armenia, australia, austria, azerbaijan, bahamas, bahrain, bangladesh, barbados, belarus, belgium, belize, benin, bhutan, bolivia, bosnia and herzegovina, botswana, brazil, brunei, bulgaria, burkina faso, burundi, cabo verde, cambodia, cameroon, canada, central african republic, chad, chile, china, colombia, comoros, costa rica, croatia, cuba, cyprus, czech republic, democratic republic of the congo, denmark, djibouti, dominica, dominican republic, ecuador, egypt, el salvador, equatorial guinea, eritrea, estonia, eswatini, ethiopia, federated states of micronesia, fiji, finland, france, gabon, gambia, georgia, germany, ghana, greece, grenada, guatemala, guinea, guinea-bissau, guyana, haiti, honduras, hungary, iceland, india, indonesia, iran, iraq, ireland, israel, italy, jamaica, japan, jordan, kazakhstan, kenya, kiribati, kuwait, kyrgyzstan, laos, latvia, lebanon, lesotho, liberia, libya, liechtenstein, lithuania, luxembourg, madagascar, malawi, malaysia, maldives, mali, malta, marshall islands, mauritania, mauritius, mexico, moldova, monaco, mongolia, montenegro, morocco, mozambique, myanmar, namibia, nauru, nepal, netherlands, new zealand, nicaragua, niger, nigeria, north korea, north macedonia, norway, oman, pakistan, palau, panama, papua new guinea, paraguay, peru, philippines, poland, portugal, qatar, romania, russia, rwanda, saint kitts and nevis, saint lucia, saint vincent and the grenadines, samoa, san marino, sao tome and principe, saudi arabia, senegal, serbia, seychelles, sierra leone, singapore, slovakia, slovenia, solomon islands, somalia, south africa, south korea, south sudan, spain, sri lanka, sudan, suriname, sweden, switzerland, syria, tajikistan, tanzania, thailand, timor-leste, togo, tonga, trinidad and tobago, tunisia, turkey, turkmenistan, tuvalu, uganda, ukraine, united arab emirates, united kingdom, united states of america, uruguay, uzbekistan, vanuatu, vatican city, venezuela, vietnam, yemen, zambia, zimbabwe
	Ethnic	black, white, asian, latino, jewish, arab, indian, african, nigger, african american, caucasian, hispanic, asian, native american, pacific islander, middle eastern, north african, indigenous, aboriginal
	Religious	christian, catholic, protestant, orthodox, anglican, baptist, mormon, jehovah's witness, judaism, islam, hinduism, buddhism, sikhism, shinto, taoism, atheist, agnostic, pagan, zoroastrianism, jainism, baha, scientologist, rastafarian, unitarian, falun gong, druze, samaritan, yazidi, ahmadi, alawite, coptic, animist, wiccan, satanist, seventh-day adventist, muslim, jew
UN	Disabilities	disabled, blind, deaf, mute, autistic, down syndrome, schizophrenic, bipolar, mentally ill, wheelchair user, paraplegic, quadriplegic, dwarf, albino, epileptic, diabetic, hiv positive, cancer patient, obese, amputee, retarded
	Sex	lesbian, gay, bisexual, queer, pansexual, asexual, lgbt
	Migrants	refugee, immigrant, migrant, asylum seeker, foreigner, expatriate, stateless
	Others	indigenous peoples, forcibly displaced persons, vocational targets
	Race or ethnicity	black, white, asian, latino, jewish, arab, indian, african, nigger, african american, caucasian, hispanic, asian, native american, pacific islander, middle eastern, north african, indigenous, aboriginal
	National origin	<i>[Same as National and Ethnic of UN]</i>
Meta	Disability or serious disease	disabled, blind, deaf, mute, autistic, down syndrome, schizophrenic, bipolar, mentally ill, wheelchair user, paraplegic, quadriplegic, dwarf, albino, epileptic, diabetic, hiv positive, cancer patient, obese, amputee, retarded
	Religious affiliation	christian, catholic, protestant, orthodox, anglican, baptist, mormon, jehovah's witness, judaism, islam, hinduism, buddhism, sikhism, shinto, taoism, atheist, agnostic, pagan, zoroastrianism, jainism, baha, scientologist, rastafarian, unitarian, falun gong, druze, samaritan, yazidi, ahmadi, alawite, coptic, animist, wiccan, satanist, seventh-day adventist, muslim, jew
	Caste	dalit, brahmin, kshatriya, vaishya, shudra
	Sexual orientation	lesbian, gay, bisexual, queer, pansexual, asexual, lgbt
	Gender identity	women, men, transgender, non-binary, intersex, cisgender, female, male
	Immigration	refugee, immigrant, migrant, asylum seeker, foreigner, expatriate, stateless

Table 10: English Listed target groups of hate speech

Policy	Category	Target Group
Twitter	Race or ethnicity	black, african american, white, caucasian, hispanic, latino, asian, native american, pacific islander, middle eastern, north african, indigenous, aboriginal
	National origin	<i>[Same as National and Ethnic of UN]</i>
	Religious affiliation	christian, catholic, protestant, baptist, methodist, lutheran, presbyterian, orthodox christian, mormon, jehovah's witnesses, jewish, muslim, sunni, shia, sufi, ahmadiyya, druze, ismaili, buddhist, hindu, sikh, jain, zoroastrian, baha'i, taoist, confucian, shinto, pagan, atheist, agnostic, humanist
	Sex	male, female, intersex
	Gender identity	cisgender man, cisgender woman, transgender man, transgender woman, non-binary, genderqueer, agender, bigender, genderfluid, pangender, two-spirit
	Sexual orientation	heterosexual, gay, lesbian, bisexual, pansexual, asexual, demisexual, polysexual, queer
	Age	children, teenager, youth, adult, senior
	Disability or serious disease	physical disability, visual impairment, hearing impairment, deafness, intellectual disability, autism spectrum disorder, down syndrome, epilepsy, cerebral palsy, muscular dystrophy, multiple sclerosis, parkinson's disease, alzheimer's disease, hiv, aids, diabetes, cancer, asthma, heart disease, chronic kidney disease, autoimmune disorders, lupus, crohn's disease, fibromyalgia, sickle cell disease, hemophilia, thalassemia, celiac disease, dwarfism, clinically obese, arthritis, rheumatoid arthritis, hepatic disease, tuberculosis, bipolar disorder, depression, anxiety disorders, borderline personality disorder, schizophrenia, aphasia, dyspraxia
	Age	children, teenager, youth, adult, senior
	Caste, Ethnicity, or Race	dalit, adivasi, brahmin, kshatriya, vaishya, shudra, black or african american, white or caucasian, hispanic or latino, asian, native american, pacific islander, middle eastern, north african, indigenous, aboriginal
YouTube	Disability	physical disability, visual impairment, hearing impairment, deafness, intellectual disability, autism spectrum disorder, epilepsy, cerebral palsy, muscular dystrophy, multiple sclerosis, parkinson's disease, alzheimer's disease, hiv, aids, diabetes, cancer, asthma, heart disease, chronic kidney disease, autoimmune disorders, lupus, crohn's disease, fibromyalgia, sickle cell disease, hemophilia, thalassemia, celiac disease, dwarfism, clinically obese, arthritis, rheumatoid arthritis, hepatic disease, tuberculosis, bipolar disorder, depression, anxiety disorders, borderline personality disorder, schizophrenia, aphasia, dyspraxia
	Immigration Status	citizen, permanent resident, documented immigrant, refugee, asylum seeker, undocumented immigrant
	Nationality	<i>[Same as National and Ethnic of UN]</i>
	Religion	christian, catholic, protestant, baptist, methodist, lutheran, presbyterian, orthodox christian, mormon, jehovah's witnesses, jewish, muslim, sunni, shia, sufi, ahmadiyya, druze, ismaili, buddhist, hindu, sikh, jain, zoroastrian, baha'i, taoist, confucian, shinto, pagan, atheist, agnostic, humanist
	Sex, Gender, or Sexual Orientation	male, female, intersex, cisgender man, cisgender woman, transgender man, transgender woman, non-binary, genderqueer, agender, bigender, genderfluid, pangender, two-spirit, heterosexual, gay, lesbian, bisexual, pansexual, asexual, demisexual, polysexual, queer
	Veteran Status	military veteran, non-veteran
	Victims	victim of sexual assault, victim of domestic violence, victim of crime, victim of harassment, victim of trafficking, victim of bullying, victim of discrimination

Table 11: English Listed target groups of hate speech

Category	Target Group
여성과 소녀	여성, 여자, 소녀, 여아, 여학생, 여성 인권옹호자, 여성인권옹호자, 여성 정치인, 여성 언론인, 여성 활동가
종교적 소수자	종교 소수자, 무슬림, 이슬람교도, 유대인, 유대교도, 시크교도, 힌두교도, 불교도, 바하이 신도, 야지디, 아호마디야 신도, 기독교 소수파, 소수파 기독교인, 소수 종교 신도
인종 민족 국가적 소수자	인종 소수자, 민족 소수자, 국가적 소수자, 국적 소수자, 흑인, 아프리카계, 라틴계, 라티노, 라티나, 라틴엑스, 아시아계, 동남아계, 중국계, 한국계, 일본계, 중동 북아프리카계, 아랍인, 쿠르드인, 로마인, 롬인, 팔레스타인인, 로힝야, 유럽 소수 민족, 유럽의 소수 민족, 원주민, 선주민, 토착민, 아메리카 원주민, 아보리지니, 토레스 해협 섬 주민, 마오리, 사미, 아이누, 인디헤나, 라틴아메리카 원주민, 라틴아메리카 원주민, 아프가니스탄, 알바니아, 알제리, 안도라, 앙골라, 안티과아, 바부다, 아르헨티나, 아르메니아, 오스트레일리아(호주), 오스트리아, 아제르바이잔, 바하마, 바레인, 방글라데시, 바베이도스, 벨라루스, 벨기에, 벨리즈, 베냉, 부탄, 볼리비아, 보스니아 헤르체고비나, 보츠와나, 브라질, 브루나이, 불가리아, 부르키나파소, 부룬디, 카보베르데, 캄보디아, 카메룬, 캐나다, 중앙아프리카공화국, 차드, 칠레, 중국, 콜롬비아, 코모로, 코스타리카, 크로아티아, 쿠바, 키프로스, 체코(체코공화국), 콩고민주공화국, 덴마크, 지부티, 도미니카, 도미니카공화국, 에콰도르, 이집트, 엘살바도르, 적도기니, 에리트레아, 에스토니아, 에스와티니, 에티오피아, 미크로네시아연방, 피지, 핀란드, 프랑스, 가봉, 감비아, 조지아, 독일, 가나, 그리스, 그레나다, 과테말라, 기니, 기니비사우, 가이아나, 아이티, 온두라스, 헝가리, 아이슬란드, 인도, 인도네시아, 이란, 이라크, 아일랜드, 이스라엘, 이탈리아, 자메이카, 일본, 요르단, 카자흐스탄, 케냐, 키리바시, 쿠웨이트, 키르기스스탄, 라오스, 라트비아, 레바논, 레소토, 라이베리아, 리비아, 리히텐슈타인, 리투아니아, 룩셈부르크, 마다가스카르, 말라위, 말레이시아, 몰디브, 말리, 몰타, 마셜제도, 모리타니, 모리셔스, 멕시코, 몰도바, 모나코, 몽골, 몬테네그로, 모로코, 모잠비크, 미얀마, 나미비아, 나우루, 네팔, 네덜란드, 뉴질랜드, 니카라과, 니제르, 나이지리아, 북한, 북마케도니아, 노르웨이, 오만, 파키스탄, 팔라우, 파나마, 파푸아뉴기니, 파라과이, 페루, 필리핀, 폴란드, 포르투갈, 카타르, 루마니아, 러시아, 르완다, 세인트키츠 네비스, 세인트루시아, 세인트빈센트 그레나딘, 사모아, 산마리노, 상투메 프린시페, 사우디아라비아, 세네갈, 세르비아, 세이셸, 시에라리온, 싱가포르, 슬로바키아, 슬로베니아, 솔로몬제도, 소말리아, 남아프리카공화국, 대한민국(남한), 남수단, 스페인, 스리랑카, 수단, 수리남, 스웨덴, 스위스, 시리아, 타지키스탄, 탄자니아, 태국, 동티모르(티모르레스트), 토고, 통가, 트리니다드 토바고, 튀니지, 튀르키예(터키), 투르크메니스탄, 투발루, 우간다, 우크라이나, 아랍에미리트, 영국, 미국, 우루과이, 우즈베키스탄, 바누아투, 바티칸시국, 베네수엘라, 베트남, 예멘, 잠비아, 짐바브웨
언어적 소수자	언어 소수자, 소수 언어 사용자, 이중언어 화자, 사미어 사용자, 아이누어 사용자, 쿠르드어 사용자, 베르베르어 사용자, 아마지그어 사용자, 카탈루냐어 사용자, 카탈란어 사용자, 한국어어 사용자, 수어 사용자
이주민 난민 무국적자	이주민, 이민자, 이주 노동자, 이주노동자, 난민, 난민 신청자, 망명 신청자, 망명 신청자, 국내 실향민, 국내실향민, IDP, 무국적자, 미등록 이주민, 미등록이주민, 이주 배경 청년, 이주배경 청년
LGBTIQ+	성소수자, LGBTIQ+, 레즈비언, 게이, 양성애자, 바이섹슈얼, 팬섹슈얼, 범성애자, 무성애자, 에이섹슈얼, 트랜스젠더, 트랜스 여성, 트랜스여성, 트랜스 남성, 트랜스남성, 논바이너리, 비이분법, 젠더 비순응, 젠더비순응, 인터섹스, 쿼어
장애인	장애인, 지체장애인, 뇌병변장애인, 시각장애인, 청각장애인, 농인, 난청인, 언어장애인, 지적장애인, 발달장애인, 자폐 스펙트럼 당사자, 자폐스펙트럼 당사자, 자폐성 장애인, 학습장애 당사자, 난독증 당사자, 정신장애인, 정신 건강 장애가 있는 사람, 정신건강 장애가 있는 사람, 희귀질환 장애인, 만성질환 장애인
언론인과 인권옹호자	언론인, 기자, 보도진, 편집자, 팩트체커, 여성 언론인, 여성언론인, 인권옹호자, 인권 활동가, 인권활동가, 시민사회 활동가, 시민사회활동가, 법률 지원 활동가, 환경 운동가, 환경운동가, LGBTIQ+ 인권 활동가, LGBTIQ+ 인권활동가, 여성 인권 활동가, 여성 인권활동가, 여성인권 활동가, 여성인권활동가

Table 12: Korean Listed target groups of hate speech

Category	Target Group
种族	非洲人, 非裔美国人, 加勒比非洲裔, 阿拉伯人, 亚美尼亚人, 亚洲人, 亚述人, 澳大利亚土著, 巴尔干人, 孟加拉人, 巴斯克人, 柏柏尔人, 黑人, 巴西人, 英国人, 保加利亚人, 缅甸人, 高加索人, 柬埔寨人, 卡津人, 中国人, 古巴人, 捷克人, 丹麦人, 多米尼加人, 荷兰人, 埃及人, 英格兰人, 爱沙尼亚人, 菲律宾人, 芬兰人, 法国人, 格鲁吉亚人, 德国人, 希腊人, 吉普赛人/罗姆人, 海地人, 汉族, 西班牙裔, 匈牙利人, 冰岛人, 印度人, 土著, 因纽特人, 伊朗人, 伊拉克人, 爱尔兰人, 以色列人, 意大利人, 牙买加人, 日本人, 犹太人, 约旦人, 韩国人, 库尔德人, 老挝人, 拉丁裔, 黎巴嫩人, 马来人, 毛利人, 墨西哥人, 中东人, 蒙古人, 摩洛哥人, 穆斯林, 美洲原住民, 新西兰人, 尼日利亚人, 北欧人, 挪威人, 巴基斯坦人, 巴勒斯坦人, 波斯人, 波兰人, 葡萄牙人, 波多黎各人, 罗马尼亚人, 俄罗斯人, 沙特阿拉伯人, 苏格兰人, 塞尔维亚人, 新加坡人, 索马里人, 南非人, 西班牙人, 斯里兰卡人, 苏丹人, 瑞典人, 瑞士人, 叙利亚人, 台湾人, 泰国人, 藏族人, 土耳其人, 乌克兰人, 越南人, 威尔士人, 也门人, 波斯尼亚人, 白人, 拉丁裔, 非洲人, 黑鬼, 白俄罗斯人, 哈萨克人, 乌兹别克人
宗教	基督徒, 天主教徒, 新教徒, 东正教徒, 圣公会, 浸礼会, 摩门教徒, 耶和华见证人, 犹太教, 伊斯兰教, 印度教, 佛教, 锡克教, 神道教, 道教, 无神论者, 不可知论者, 异教徒, 祆教, 耆那教, 巴哈伊教, 科学教, 拉斯塔法里教徒, 一神论者, 法轮功, 德鲁兹教徒, 撒马利亚人, 雅兹迪教徒, 艾哈迈迪派, 阿拉维派, 科普特教徒, 万物有灵论者, 威卡教徒, 撒旦教徒, 基督复临安息日会
性别	女性, 男性, 跨性别, 非二元, 双性人, 顺性别, 女, 男
性取向	女同性恋, 男同性恋, 双性恋, 酷儿, 泛性恋, 无性恋, LGBT
残疾	残障人士, 盲人, 聋人, 哑巴, 自闭症患者, 唐氏综合症患者, 精神分裂症患者, 躁郁症患者, 精神病患者, 轮椅使用者, 截瘫患者, 四肢瘫痪者, 侏儒, 白化病患者, 癫痫患者, 糖尿病患者, 艾滋病病毒携带者, 癌症患者, 肥胖者, 截肢者
年龄	儿童, 青少年, 青年, 成年人, 老年人
移民	难民, 移民, 迁徙者, 寻求庇护者, 外国人, 侨民, 无国籍者

Table 13: Chinese Listed target groups of hate speech