

Multi³Hate: Multimodal, Multilingual, and Multicultural Hate Speech Detection with Vision–Language Models

Minh Duc Bui[∇] Katharina von der Wense^{∇♣} Anne Lauscher[◇]

[∇]Johannes Gutenberg University Mainz, Germany

[♣]University of Colorado Boulder, USA [◇]University of Hamburg, Germany

{minhducbui, k.vonderwense}@uni-mainz.de

anne.lauscher@uni-hamburg.de

Abstract

Warning: this paper contains content that may be offensive or upsetting

Hate speech moderation on global platforms poses unique challenges due to the multimodal and multilingual nature of content, along with the varying cultural perceptions. How well do current vision-language models (VLMs) navigate these nuances? To investigate this, we create the first multimodal and multilingual parallel hate speech dataset, annotated by a multicultural set of annotators, called **Multi³Hate**. It contains 300 parallel meme samples across 5 languages: English, German, Spanish, Hindi, and Mandarin. We demonstrate that cultural background significantly affects multimodal hate speech annotation in our dataset. The average pairwise agreement among countries is just 74%, significantly lower than that of randomly selected annotator groups. Our qualitative analysis indicates that the lowest pairwise label agreement—only 67% between the USA and India—can be attributed to cultural factors. We then conduct experiments with 5 large VLMs in a zero-shot setting, finding that these models align more closely with annotations from the US than with those from other cultures, even when the memes and prompts are presented in the dominant language of the other culture. Code and dataset are available at <https://github.com/MinhDucBui/Multi3Hate>.

1 Introduction

Our cultural backgrounds significantly shape our perceptions of the world. For instance, individuals raised in collectivist societies often emphasize group harmony, leading them to interpret events through a relational lens, whereas those from individualist societies may prioritize personal achievements and autonomy, resulting in a perception that focuses on individual characteristics (Triandis, 1995; Nisbett, 2003). Consequently, identical content can be perceived vastly differently depending on cultural background, posing challenges for

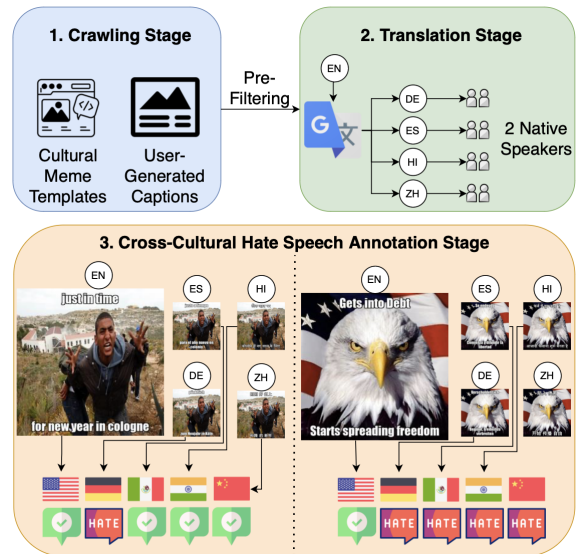


Figure 1: Our dataset creation process is divided into three stages: 1. Crawling Stage; 2. Translation Stage; and 3. Cross-Cultural Hate Speech Annotation Stage. The two examples illustrate the varying ways in which memes are annotated across different cultures.

hate speech moderation models as they must balance diverse perspectives without marginalizing certain cultures while favoring others.

Towards incorporating this important goal, Lee et al. (2024) released the first and only hate speech dataset, annotated by a multicultural set of annotators, revealing that large language models often exhibit bias toward Anglospheric cultures. However, their work leaves critical gaps unaddressed: (1) The dataset is limited to text-based content, excluding multimodal forms of hate; (2) It is restricted to English-language samples, overlooking non-English-speaking cultures. This narrow scope not only hampers the cross-cultural evaluation of multimodal hate speech detection models, providing little guidance for practitioners, but also amplifies the exclusion of non-English-speaking cultures from cross-cultural analysis.

Dataset	Multi-modal	Multi-cultural set of Annotators	Multi-lingual (+Parallel)
HateXplain (Mathew et al., 2021)	✗	✗	✗
XHate-999 (Glavaš et al., 2020)	✗	✗	✓ (+✓)
MMHS150k (Gomez et al., 2020)	✓	✗	✗
Hateful Memes (Kiela et al., 2020)	✓	✗	✗
CrisisHateMM (Bhandari et al., 2023)	✓	✗	✗
MUTE (Hossain et al., 2022)	✓	✗	✓ (+✗)
CREHate (Lee et al., 2024)	✗	✓	✗
Multi³Hate Ours	✓	✓	✓ (+✓)

Table 1: Comparison of hate speech datasets across three dimensions: multimodal, multicultural set of annotators, and multilingual, along with whether they are parallel. Our dataset is the first to be both multimodal and multilingual. Additionally, the multimodal dataset is annotated by a multicultural set of annotators.

To close this gap, we are the first, to the best of our knowledge, to release a parallel **multilingual** and **multimodal** hate speech dataset. Additionally, the dataset is annotated by a **multicultural** set of annotators, as shown in Table 1. Our dataset, **Multi³Hate**, comprises a curated collection of 300 memes—images paired with embedded captions—a prevalent form of multimodal content, presented in five languages: English (*en*), German (*de*), Spanish (*es*), Hindi (*hi*), and Mandarin (*zh*). Each of the 1,500 memes (300×5 languages) is annotated for hate speech in the respective target language by at least five native speakers from the same country. These countries were chosen based on the largest number of native speakers of each target language: USA (US), Germany (DE), Mexico (MX), India (IN), and China (CN) (Instituto Cervantes, 2023; World Population Review, 2024). As in prior research, we use the country of the annotators as a cultural proxy (EVS/WVS, 2022; Koto et al., 2023; Lee et al., 2024).

We demonstrate that cultural background significantly influences multimodal hate speech annotation in our dataset. The average pairwise agreement among countries is only 74%, significantly lower

than that of randomly selected annotator groups. The lowest agreement, at just 67%, occurs between the USA and India. Through qualitative analysis involving multicultural annotators with ties to both countries, we demonstrate that these disagreements can be attributed to cultural factors, such as differing social norms. Consequently, **Multi³Hate enables the analysis of multimodal models for cross-cultural hate speech detection across a range of diverse speaking cultures.**

Furthermore, we conduct experiments using 5 large VLMs in a zero-shot setting. Our experiments with English prompts reveal that these models consistently align more closely with annotations from the US than with those from other cultures, independent of the meme language. Specifically, out of 50 combinations of models, languages, and input variations, 42 demonstrate the highest alignment with US labels. Even when we switch the prompt language to the dominant language of a specific culture, we still observe similarly high alignment to US annotators. We therefore demonstrate that **VLMs align more closely with hate speech annotations from the US** than with those from non-English-speaking cultures, **even when the memes and prompts are presented in the dominant language of the other culture.** This trend poses a risk of marginalizing certain cultures, despite VLMs being used in their native languages, while simultaneously privileging US cultural perspectives.

2 Related Work

Multilingual Hate Speech While several text-based hate speech datasets exist in various languages (Jeong et al., 2022; Mubarak et al., 2022; Yadav et al., 2023; Demus et al., 2022), there has been limited focus on creating a parallel hate speech dataset. The only notable example is Glavaš et al. (2020), which developed a parallel text dataset in six languages.

Moreover, most multimodal hate speech datasets are in English (Suryawanshi et al., 2020; Hossain et al., 2022; Bhandari et al., 2023; Kiela et al., 2020; Gomez et al., 2020), with limited resources available for other languages. Notable exceptions include a Bengali dataset by Karim et al. (2022), an Italian dataset by Miliani et al. (2020), and a Tamil dataset by Suryawanshi et al. (2020). To our knowledge, no parallel multimodal hate speech datasets exist.¹

¹Gold et al. (2021) translated the English captions of the

Category	Topic
Religion	Christianity
	Islam
	Judaism
Nationality	Germany (DE)
	United States (US)
	Mexico (MX)
	China (CN)
	India (IN)
Ethnicity	Asian
	Black
	Middle Eastern
	White
LGBTQ+	Transgender
Political Issues	Law Enforcement
	Feminism

Table 2: Final list of topics across our 5 sociopolitical categories, with each topic featuring 3 image templates. For a comprehensive overview of the topics, associated search keywords, and the final number of samples, please refer to Table 13 in the Appendix.

Cross-cultural Hate Speech Lee et al. (2024) are the first to analyze how cultural background affects hate speech annotations, finding that annotators’ nationality significantly influence their annotation. However, their study is limited to English-speaking cultures due to its exclusively English dataset. Expanding to include non-English-speaking cultures could provide valuable insights for a more inclusive moderation system.

Cross-cultural VLMs Several studies have established benchmarks to probe cultural awareness in VLMs. For instance, researchers have focused on creating culturally diverse image descriptions, visual grounding, and benchmarks for cultural visual question-answering (Liu et al., 2021; Cao et al., 2024; Burda-Lassen et al., 2024; Ye et al., 2024; Karamolegkou et al., 2024; Nayak et al., 2024). However, there has been little to no attention given to cross-cultural multimodal hate speech detection.

3 Dataset Construction

We now describe the pipeline used to create Multi³Hate, as illustrated in Figure 1.

Hateful Meme dataset (Kiela et al., 2020) into German but did not create or release images with the new captions due to licensing restrictions on the original dataset.

²On December 31, 2015, Cologne, Germany, recorded about 1,200 criminal complaints, nearly half for sexual offenses, igniting controversy over the country’s refugee policy (Bosen, 2020).



Figure 2: Example of a parallel meme. The original English meme reads: “just in time <sep> for new year in cologne”. Only in Germany is this meme perceived as hate speech.²

3.1 Crawling

Image Templates & User Captions To effectively modify captions in memes, we select memes with a simple structure, featuring captions at the top and/or bottom. For this purpose, we crawl a website³ where users can submit captions based on meme image templates provided by other users, collecting both the templates and user-generated captions.

Sociopolitical Categories To ensure our samples are influenced by cultural perceptions, we curate a list of culturally relevant templates by filtering them according to *sociopolitical categories*. These categories were discussed and decided among the authors. Each category is further divided into specific topics based on established criteria, see Appendix A.1 for more details.

For every topic, we generate relevant keywords. As an example, for the topic “Germany”, we create the keyword “german” and match meme templates to these keywords based on their template names. Subsequently, we select the top three meme templates with the highest number of user captions. For details, see Appendix A.2. The final list, which includes a total of 5 categories, 15 topics and 45 image meme templates, is presented in Table 2.

Pre-Filtering We ensure high-quality captions after crawling by implementing three pre-filtering

³<https://memegenerator.net> (Accessed: May, 2024)



Figure 3: We provide examples from each category with hate speech annotations, highlighting cultural variability in perceptions and challenges for annotators in identifying targeted groups and stereotypes.

steps to verify that the captions are: (1) in English, (2) multimodal, and (3) free from wordplay. Memes that can be classified solely based on their captions may lead to underutilization of the images by VLMs. Furthermore, wordplay can introduce translation errors and distort the intended meaning. We provide a detailed description of the pre-filtering implementation in Appendix A.3. After pre-filtering, we have a total of 450 captions distributed across 45 image templates.

3.2 Translation

We start by utilizing the v3 Google Translate API⁴ to generate machine translations of our 450 captions into four target languages: German, Spanish, Hindi, and Mandarin.

Following this, we conduct two rounds of validation with two native speakers of the target language who are also fluent in English. Their task is to verify the accuracy of the translations and make any necessary corrections. Each annotator is provided with a detailed annotation guide, which can be found in Appendix A.4. We then recreate each meme by overlaying the new captions onto the image templates using the Python Pillow package (Clark, 2015), see Figure 2 for one example.

3.3 Cross-Cultural Annotation

Annotator Recruitment We recruit annotators through Prolific⁵, ensuring the following: (1) they are native speakers of the target language; (2) have spent most of their lives in the target country; (3) their nationality aligns with the target country; (4) they identify as monocultural in relation to the target country and (5) they currently reside in the

target country.⁶ We hire 445 annotators across all countries, maintaining a balanced representation of gender. All annotators gave explicit consent, were informed of the risks, and received a fair wage compensation (see 6 Ethics Statement). For a detailed demographic distribution, see Table 9 in Appendix.

Pre-Annotation To ensure our dataset is balanced, we implement a pre-annotation stage, in which the dataset is evenly divided among our five target countries and annotated twice. Subsequently, we adjust the samples of *hate speech* and *non-hate speech* based on the annotation results. For further details, please refer to Appendix A.5.

In total, our final dataset consists of 300 parallel memes across five languages distributed across 45 templates, resulting in 1,500 memes.

Annotation Process Before the annotation process begins, annotators receive a definition of hate speech⁷ along with examples in their native language, see Figure 10 in the Appendix. Each annotator is provided with the survey and samples – also in their native language – and is asked to label each meme (combination of image and embed caption) as *hate speech*, *non-hate speech*, or *I don't know*. For every sample and language, we collect a minimum of five annotations. The final label is determined through majority voting; when there is a tie between *hate speech* and *non-hate speech*, we gather additional annotations until a majority consensus is reached. A detailed description of the survey design and quality checks can be found in Appendix A.6.

⁴<https://cloud.google.com/translate/docs/reference/rest/v3/projects/translateText> (Accessed: June, 2024)

⁵<https://www.prolific.com>

⁶For India and China, we relaxed the residency requirement once we were no longer able to recruit additional participants.

⁷<https://www.un.org/en/hate-speech/understanding-hate-speech/what-is-hate-speech>

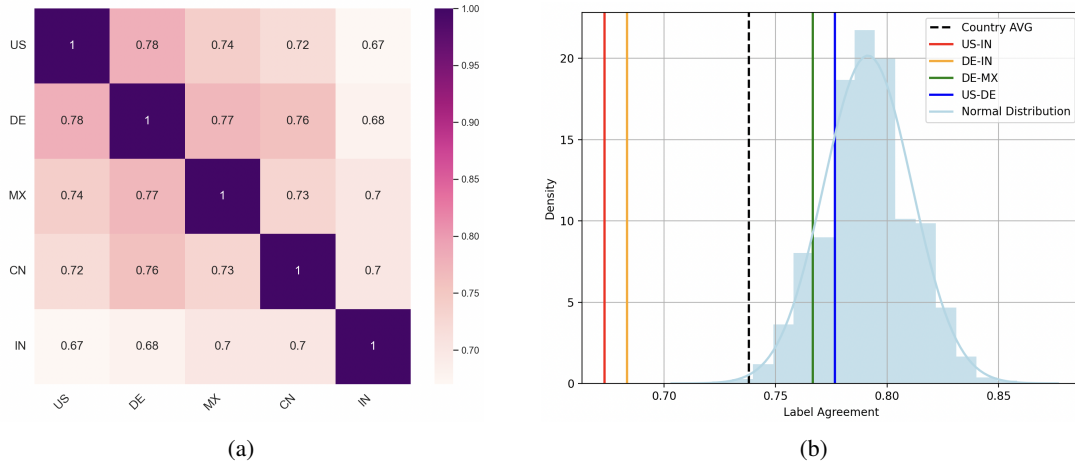


Figure 4: (a) Pairwise label agreement for all countries, ranked by average agreement. (b) A comparison of the top two and bottom two country pairs’ pairwise label agreement, along with the overall average across all countries, against randomly selected annotator groups. The results indicate that the lowest agreement pairs and the overall average differ significantly from random groups

4 Analysis of Annotations

4.1 Dataset Overview

We present examples in Figure 3.

Distribution of Hate Speech We report the proportion of *hate speech* and *non-hate speech* for each culture in Table 3. A significant lower number of samples were classified as *hate speech* by US respondents compared to other cultures. For instance, Chinese annotators labeled approximately 63% of instances as *hate speech*, while US annotators labeled only 51% as such.

Inter-Annotator Agreement (IAA) We measure the IAA across hate speech annotations for each cultural group using Krippendorff’s α coefficient (Krippendorff, 2011). The values obtained are as follows: for the US, $\alpha = 0.4686$; for DE, $\alpha = 0.4537$; for MX, $\alpha = 0.3895$; for IN, $\alpha = 0.4018$; and, for CN, $\alpha = 0.4322$. These values are higher than or comparable to those reported in previous hate speech research (Ross et al., 2016; Lee et al., 2024), demonstrating that there is a consensus on hate speech within each culture and pointing to the general validity of our annotation setup.

4.2 Significance of Culture

To demonstrate that cultural background significantly affects multimodal hate speech annotation in our dataset, we closely follow Lee et al. (2024).

Overall Significance To assess the significance of cultural differences, we apply a chi-squared test

Country	Hate Speech	Non-Hate Speech	Total
US	51%	49%	300
DE	59%	41%	300
MX	55%	45%	300
IN	60%	40%	300
CN	63%	37%	300

Table 3: Proportion of *hate speech* and *non-hate speech* for each country. Chinese annotators labeled the majority of samples as hate speech, whereas US annotators identified the fewest instances as such.

to the hate speech annotations. The results reveal significant disparities ($p < 0.05$) across cultures.

Label Agreement Across Cultures We report the average pairwise label agreement across countries in Figure 4a. The highest agreement is observed between the US and Germany (78%), while the lowest occurs between the US and India (67%).

Additionally, we calculate the proportion of samples with complete or partial agreement across countries: Only 44% of samples show agreement across all countries, four countries agree for 30%, and, for 26%, only three countries agree.

Comparison with Random Annotator Groups

To demonstrate that the label disparity between cultures is not due to random variations among annotators, we create random annotator groups and calculate their agreement. Specifically, for each sample, we randomly select five annotations from across all cultures to form two groups. We then calculate the label agreement between these two random groups, repeating this process 3×10^4 times.




Figure 5: Distribution of disagreements between the USA and India. See Table 14 in the Appendix for detailed information on each category along with examples.

We plot the resulting agreement histogram in Figure 4b. To assess significance, we first confirm that the random group distribution follows a normal distribution using the D’Agostino-Pearson normality test (D’Agostino and Pearson, 1973), with a mean of 0.79 and standard deviation (σ) of 0.019.

We observe that the pairs with the lowest agreement, “US - IN” and “DE - IN”, show significant deviations from the random annotator groups, with differences of -5.97σ and -5.47σ , respectively. Additionally, the overall country average of 74% is significantly lower, by -2.70σ . Upon closer inspection, all country pairs—except for the top three (“DE - MX”, “US - DE”, and “DE - CN”),—exhibit significantly lower agreement compared to the random groups. This analysis demonstrates that an individual’s cultural background significantly influences their perception of multimodal hate speech.

4.3 Analysis of Label Disagreements

Label Agreement Across Categories To further analyze the disagreement between cultures, we examine the sociopolitical categories. Table 4 presents the pairwise agreement across countries for each category. The highest label agreement is observed in the “Religion” category, with an average of 78%, while the “LGBTQ+” category shows the lowest agreement at only 61%, which reflects deeper cultural sensitivities and differing norms. Interestingly, the “US - DE” pair has the highest agreement for every category, while the “US - IN” and “US - CN” pairs exhibit the lowest.

Annotators’ Disagreement Analysis We conduct a qualitative analysis to examine *why* cultures differ in their hate speech annotations, focusing on the pair with the highest disagreement: the USA

Category	AVG	Highest	Lowest
Religion	78%	US-DE: 83%	US-IN: 74%
Nationality	69%	US-DE: 72%	US-IN: 62%
Ethnicity	77%	US-DE: 81%	US-IN: 70%
LGBTQ+	61%	US-DE: 73%	US-CN: 42%
Political Issues	75%	US-DE: 77%	US-CN: 61%

Table 4: We present the culture average pairwise agreement for each sociopolitical category, highlighting the culture pairs with the highest and lowest agreement.

Inp.	GT					
	US	DE	MX	IN	CN	
Uni- vs. Multimodal Setting						
Multi-	75.8* _{±2.1}	72.2* _{±2.6}	69.2* _{±3.2}	63.1* _{±2.2}	68.7* _{±2.8}	
Uni-	65.4 _{±3.9}	58.3 _{±4.4}	60.8 _{±3.5}	56.7 _{±4.0}	55.0 _{±4.5}	
Scale						
>70B	67.6 _{±5.3}	62.2 _{±5.7}	64.3 _{±4.6}	59.1 _{±4.3}	59.2 _{±5.6}	
<10B	62.0 _{±3.9}	60.1 _{±5.1}	61.7 _{±4.6}	58.0 _{±3.4}	58.1 _{±5.4}	

Table 5: **Upper table:** We compare the average performance of the best model in the unimodal setting versus the multimodal setting. **Lower table:** We compare the average performance of large models (>70B) with that of smaller models of the same model family (<10B).

and India. We recruit 7 annotators who are bilingual in Hindi and English, born in one of the two countries, currently residing in the other, and self-identifying as multicultural with ties to both cultures. These annotators are shown memes where the two cultures’ annotations diverge, and we ask them to explain the reasons for their disagreement in free-text form. Using an inductive “bottom-up” approach, one author extracts keywords from each response, summarizing the text, giving us an initial codebook of 37 codes. A hired annotator then independently reassigns these established codes to the samples. We then establish 6 major themes.

As shown in Figure 5, “Sensitivity Around Minority Groups” and “Social Norms & Cultural Values” account for 53.6%, while “Historical & Political Context” and “Non-Existing Stereotypes” contribute 31%. Together, these four themes, totaling 84.6%, likely *reflect cultural differences*. Ideally, we aim to minimize the proportion of “Language Error”, which accounts for only 5.2%. However, 10.3% fall under “Annotation Ambiguity”, which may stem from annotation noise or reflect individual annotators’ personal preferences. In conclusion, our cross-cultural disagreements can largely be attributed to cultural differences.

GT Inp.	US	DE	MX	IN	CN
GPT-4o					
<i>en</i> : IMG / +CAPT	75.8* _{±2.1} / 75.5** _{±1.2}	72.2 _{±2.6} / 71.7 _{±1.6}	69.2 _{±3.2} / 69.3 _{±2.5}	<u>63.1</u> _{±2.2} / <u>63.2</u> _{±1.6}	68.7 _{±2.8} / 67.6 _{±2.9}
<i>de</i> : IMG / +CAPT	73.8** _{±1.0} / 74.2** _{±0.5}	71.6 _{±1.6} / 71.3 _{±1.8}	69.0 _{±2.2} / 69.0 _{±1.7}	<u>63.8</u> _{±1.7} / <u>63.6</u> _{±1.1}	67.2 _{±2.0} / 66.7 _{±1.7}
<i>es</i> : IMG / +CAPT	74.7** _{±1.0} / 75.6** _{±1.5}	72.0 _{±1.2} / 72.5 _{±2.0}	70.7 _{±2.3} / 70.5 _{±2.7}	<u>63.7</u> _{±2.1} / <u>63.3</u> _{±2.8}	65.6 _{±2.4} / 65.4 _{±3.2}
<i>hi</i> : IMG / +CAPT	69.7* _{±1.4} / 71.3** _{±1.4}	68.1 _{±1.7} / 68.1 _{±1.9}	68.7 _{±2.8} / 68.3 _{±2.3}	<u>65.4</u> _{±2.5} / <u>63.9</u> _{±3.3}	68.2 _{±3.6} / 66.5 _{±2.9}
<i>zh</i> : IMG / +CAPT	71.3** _{±0.7} / 72.9** _{±1.4}	66.1 _{±2.0} / 67.1 _{±2.6}	68.6 _{±1.2} / 68.6 _{±1.4}	<u>63.0</u> _{±2.5} / <u>63.9</u> _{±2.5}	68.1 _{±1.3} / 69.5 _{±2.6}
Gemini 1.5 Pro					
<i>en</i> : IMG / +CAPT	70.9* _{±2.1} / 70.9* _{±2.1}	69.7 _{±2.1} / 69.7 _{±2.1}	68.6 _{±1.0} / 68.6 _{±1.0}	<u>65.0</u> _{±0.7} / <u>65.0</u> _{±0.7}	66.7 _{±2.7} / 66.7 _{±2.7}
<i>de</i> : IMG / +CAPT	69.5 _{±1.3} / 70.9* _{±1.7}	70.7* _{±1.6} / 70.9 _{±2.1}	68.1 _{±1.2} / 68.2 _{±2.3}	<u>67.1</u> _{±2.2} / <u>66.8</u> _{±3.3}	70.1 _{±3.9} / 68.1 _{±4.4}
<i>es</i> : IMG / +CAPT	69.5* _{±2.0} / 70.8* _{±3.2}	69.2 _{±1.7} / 68.8 _{±4.3}	68.7 _{±1.5} / 66.7 _{±3.1}	<u>65.4</u> _{±1.6} / <u>63.7</u> _{±2.8}	69.0 _{±2.8} / 65.9 _{±5.0}
<i>hi</i> : IMG / +CAPT	61.4 _{±5.1} / 63.5 _{±3.1}	61.4 _{±8.2} / 65.4 _{±4.4}	63.9 _{±4.3} / 65.7 _{±3.7}	62.0 _{±7.5} / <u>61.2</u> _{±4.5}	<u>57.8</u> _{±14.2} / 66.0 _{±6.7}
<i>zh</i> : IMG / +CAPT	60.8 _{±2.5} / 63.7 _{±4.6}	62.6 _{±4.2} / 63.4 _{±6.0}	66.0 _{±2.8} / 65.2 _{±5.4}	<u>60.4</u> _{±6.0} / <u>60.7</u> _{±6.6}	63.1 _{±6.3} / 62.8 _{±7.4}
Qwen2-VL 72B					
<i>en</i> : IMG / +CAPT	71.5** _{±3.9} / 70.8* _{±4.8}	62.3 _{±3.5} / 62.4 _{±3.9}	65.5 _{±3.7} / 65.4 _{±3.9}	59.1 _{±3.9} / <u>58.0</u> _{±4.1}	<u>58.9</u> _{±4.4} / 58.2 _{±4.7}
<i>de</i> : IMG / +CAPT	68.7** _{±0.8} / 70.1** _{±2.2}	64.2 _{±2.2} / 65.3 _{±2.6}	66.6 _{±1.8} / 66.4 _{±2.4}	<u>60.1</u> _{±1.4} / <u>59.2</u> _{±2.2}	61.4 _{±2.4} / 61.3 _{±2.5}
<i>es</i> : IMG / +CAPT	70.8** _{±2.1} / 71.2** _{±2.3}	62.5 _{±2.2} / 63.4 _{±3.1}	65.4 _{±1.5} / 66.1 _{±2.6}	<u>59.3</u> _{±3.2} / 59.3 _{±4.2}	59.5 _{±2.8} / <u>58.5</u> _{±3.2}
<i>hi</i> : IMG / +CAPT	62.9* _{±2.0} / 64.5* _{±3.2}	58.2 _{±3.6} / 58.4 _{±4.0}	61.7 _{±4.3} / 61.8 _{±5.2}	55.8 _{±3.2} / 56.1 _{±3.0}	<u>54.9</u> _{±4.3} / <u>54.8</u> _{±3.8}
<i>zh</i> : IMG / +CAPT	66.1 _{±2.0} / 66.7 _{±2.2}	<u>58.3</u> _{±2.2} / 58.9 _{±3.2}	63.8 _{±2.3} / 63.6 _{±2.6}	58.6 _{±3.8} / <u>57.2</u> _{±4.2}	60.6 _{±3.8} / 59.9 _{±4.4}
LLaVA OneVision 73B					
<i>en</i> : IMG / +CAPT	71.2* _{±2.4} / 68.4** _{±1.4}	69.1 _{±2.1} / 61.6 _{±2.2}	69.3 _{±2.7} / 62.9 _{±1.8}	<u>64.3</u> _{±2.5} / <u>57.4</u> _{±1.5}	66.2 _{±2.8} / 58.2 _{±2.0}
<i>de</i> : IMG / +CAPT	60.9* _{±1.5} / 65.6** _{±1.2}	58.8 _{±1.6} / 62.1 _{±2.0}	60.8 _{±1.6} / 62.8 _{±1.4}	<u>57.9</u> _{±2.1} / 59.0 _{±1.4}	59.5 _{±2.2} / <u>57.6</u> _{±1.8}
<i>es</i> : IMG / +CAPT	62.9 _{±1.0} / 64.8** _{±1.0}	63.3 _{±1.7} / 57.6 _{±1.8}	65.8** _{±1.5} / 59.4 _{±1.4}	59.8 _{±1.2} / 55.8 _{±2.6}	<u>57.8</u> _{±1.9} / <u>54.1</u> _{±1.6}
<i>hi</i> : IMG / +CAPT	58.2 _{±1.4} / 64.1** _{±0.3}	57.8 _{±0.7} / 61.8 _{±1.2}	61.5** _{±1.2} / 63.3 _{±0.1}	<u>52.3</u> _{±1.6} / 59.4 _{±1.9}	55.5 _{±2.2} / <u>58.9</u> _{±1.9}
<i>zh</i> : IMG / +CAPT	55.7 _{±0.7} / 65.3** _{±2.0}	52.4 _{±2.4} / 60.3 _{±2.8}	55.8* _{±1.4} / 60.2 _{±2.2}	<u>49.1</u> _{±2.9} / 58.3 _{±2.4}	51.7 _{±2.8} / <u>55.9</u> _{±3.0}
InternVL2 76B					
<i>en</i> : IMG / +CAPT	60.1* _{±3.0} / 65.1* _{±5.0}	55.1 _{±4.5} / 58.8 _{±6.0}	58.2 _{±4.1} / 59.9 _{±4.7}	53.8 _{±3.6} / 55.9 _{±5.4}	<u>52.5</u> _{±4.6} / <u>54.1</u> _{±5.4}
<i>de</i> : IMG / +CAPT	57.1* _{±3.6} / 63.1* _{±3.8}	52.6 _{±5.4} / 57.7 _{±6.3}	54.0 _{±5.3} / 58.8 _{±4.3}	<u>50.8</u> _{±5.2} / 54.8 _{±5.3}	50.9 _{±5.7} / <u>53.3</u> _{±5.4}
<i>es</i> : IMG / +CAPT	56.6* _{±3.3} / 62.1* _{±5.0}	52.8 _{±2.6} / 56.6 _{±5.4}	56.4 _{±3.5} / 59.2 _{±4.6}	52.6 _{±2.4} / 53.7 _{±5.3}	<u>50.2</u> _{±3.3} / <u>51.8</u> _{±5.5}
<i>hi</i> : IMG / +CAPT	48.4 _{±1.8} / 59.1* _{±3.0}	42.4 _{±2.0} / 53.8 _{±4.9}	46.4 _{±2.0} / 56.6 _{±3.1}	43.2 _{±2.3} / 53.2 _{±5.4}	<u>40.9</u> _{±2.9} / <u>49.8</u> _{±4.9}
<i>zh</i> : IMG / +CAPT	54.3 _{±2.1} / 59.4 _{±4.7}	49.2 _{±4.8} / 54.5 _{±4.6}	52.7 _{±4.8} / 56.9 _{±3.3}	49.4 _{±4.4} / 53.1 _{±3.7}	<u>47.4</u> _{±6.5} / <u>52.2</u> _{±4.7}

Table 6: The performance of our large VLMs across different meme languages while keeping the prompt in English. We report results using only the meme image as input (IMG) and also when including the image caption in the prompt (+CAPT). **Bold** text indicates the best performance across cultures; underlined text denotes the worst performance. An asterisk (*) indicates statistical significance compared to the lowest cultural performance, and a double asterisk (**) indicates significance compared to the second-highest cultural performance.

5 Experiments

5.1 Experimental Setup

Zero-Shot Setup We evaluate VLMs using a zero-shot approach to detect hate speech. The task is framed as a multiple-choice format, where the model must select between two answers: (a) *hate speech* and (b) *non-hate speech*. We implement three different prompt variations, each altering the order of answers (a) and (b). In total, we generate six prompts, maintaining English as the prompts’ language unless otherwise specified. Additionally, we experiment with two input variations: (1) using only the image (IMG) and (2) incorporating the image caption (+CAPT) into the prompt, see Appendix B.2 for detailed prompts.

Evaluation We present the average accuracy across all prompt variations, along with the standard deviation. To determine whether the observed differences are statistically significant, we apply the Wilcoxon rank-sum test (Wilcoxon, 1945), a non-parametric test that assesses whether one distribution tends to have higher values than another, without assuming normality.

Models We evaluate several models, including GPT-4o⁸ (OpenAI et al., 2024), Gemini 1.5 Pro⁹ (Georgiev et al., 2024), Qwen2-VL (Wang et al., 2024), LLaVA OneVision (Li et al., 2024), and InternVL2 (Chen et al., 2023, 2024). For more details, see Appendix B.1.

⁸API Version: gpt-4o-2024-05-13

⁹API Version: gemini-1.5-pro-001

5.2 Dataset Sanity Check

We start by demonstrating the desired multimodality and evaluating the impact of different model scaling on our dataset. The aggregated results are presented in Table 5, while detailed model performances in Table 10 in the Appendix.

Multimodality To demonstrate multimodality, we compare models that utilize images as input with those that rely solely on captions. We present the top-performing models for English input in both settings, based on average accuracy.

The top-performing multimodal model achieves an accuracy of 75.8% with US labels, compared to 65.4% for the best unimodal model. The significant higher accuracy of the multimodal models underscores the strength of our dataset in supporting multimodal analysis.

Scale We compare the average performance of models within the same family, contrasting those with fewer than 10B parameters against those with more than 70B. On average, larger models exhibit better performance across all cultural labels, with the greatest improvement of 5.5% seen on US labels. Our subsequent analysis focuses exclusively on large VLMs (models with over 70B parameters).

5.3 Prompting in English

In this section, we report experiments with VLMs in a zero-shot setting, using English as the prompt language. The results are presented in Table 6.

Strong Alignment with US Culture Label

Across input and language variations, we observe that nearly all models perform best on US labels: out of 50 variations, 42 achieve their highest performance on US labels. In 39 cases, the performance difference compared to the worst-performing cultural label is statistically significant, and in 18 cases, the difference is significant even compared to the second-highest performing label.

For example, GPT-4o consistently performs best with US labels across all languages and input variants, achieving the highest accuracy on our dataset at 75.8% for English. The model shows a significant difference from the second-highest cultural label in 8 out of 10 variations.

Low Alignment with Indian Culture Label

In contrast, the alignment between the model and hate speech annotations from Indian annotators is notably low, ranking among the bottom in accuracy

	US	DE	MX	IN	CN
Multilingual Prompts: GPT-4o					
<i>de</i>	75.0 ^{**} _{±1.0}	71.6 _{±1.9}	69.4 _{±2.0}	<u>63.7</u> _{±1.9}	67.2 _{±1.5}
- Δ	+0.8	+0.3	+0.4	+0.1	+0.5
<i>es</i>	75.0 ^{**} _{±1.3}	73.8 _{±1.1}	70.3 _{±1.8}	<u>64.1</u> _{±1.1}	67.3 _{±2.6}
- Δ	-0.6	+1.3	-0.2	+0.4	+1.9
<i>hi</i>	72.8 ^{**} _{±0.9}	70.0 _{±1.0}	71.2 _{±1.4}	<u>64.9</u> _{±1.6}	67.4 _{±1.5}
- Δ	+1.5*	+1.9*	+2.9*	+1.0	+0.9
<i>zh</i>	72.4 ^{**} _{±1.3}	66.4 _{±2.5}	69.3 _{±2.1}	<u>63.7</u> _{±2.8}	70.2 _{±3.2}
- Δ	-0.5	-0.7	+0.7	-0.2	+0.7
Multilingual Prompts: Qwen2-VL 72B					
<i>de</i>	69.9 ^{**} _{±2.9}	64.6 _{±3.8}	65.8 _{±3.2}	<u>58.8</u> _{±3.8}	61.3 _{±4.5}
- Δ	-0.2	-0.7	-0.6	+0.4	+0.0
<i>es</i>	71.6 [*] _{±3.1}	63.6 _{±3.4}	65.3 _{±3.0}	58.9 _{±3.7}	<u>57.8</u> _{±4.1}
- Δ	+0.4	-0.2	-0.8	-0.4	-0.7
<i>hi</i>	68.2 [*] _{±2.7}	64.4 _{±4.9}	67.1 _{±4.6}	<u>61.2</u> _{±5.0}	61.3 _{±6.2}
- Δ	+3.7*	+6.1*	+5.3	+5.1	+6.5
<i>zh</i>	67.1 [*] _{±2.3}	62.7 _{±2.6}	65.8 _{±3.0}	<u>59.7</u> _{±3.4}	61.6 _{±3.3}
- Δ	+0.4	+3.8	+2.2	+2.5	+1.7

Table 7: Evaluation of adjusting the prompt language to match the dominant language of the respective culture. Δ shows the difference between the multilingual prompt (+CAPT) and English prompt (+CAPT). The asterisk (*) in the Δ row shows significant difference.

across 30 out of 50 variants. Similarly, annotations from Chinese annotators also show low alignment, with 19 variants reflecting the lowest accuracy.

Comparison: IMG vs. +CAPT

Adding captions into the prompt improves performance in all languages except English, suggesting weaker OCR capabilities in VLMs for non-English text. For instance, LLaVA OneVision’s accuracy on Hindi with US labels rises from 58.2% to 64.5% with captions.

5.4 Prompting in Native Language

We experiment with adjusting the prompt language to match the dominant language of the target culture, as shown in Figure 11 in the Appendix. Incorporating captions in the prompt (+CAPT) typically enhances performance, so we focus on this setting with the best-performing open-source and proprietary models. Results are presented in Table 7

The results show mixed effects: e.g., switching to German improves the performance of GPT-4o by 0.3, while for Qwen2, it decreases by 0.7. However, none of the observed changes are statistically significant. Therefore, we conclude that **altering the prompt language to match the dominant language of a specific culture does not have a meaningful impact on aligning models.**

	US	DE	MX	IN	CN
Country Information: GPT-4o					
<i>en</i>	74.3 $_{\pm 1.0}^{**}$	68.0 $_{\pm 1.7}$	66.1 $_{\pm 1.0}$	60.1 $_{\pm 1.2}$	63.4 $_{\pm 2.2}$
- Δ	-1.2	-3.7*	-3.2*	-3.1*	-4.2*
<i>de</i>	73.5 $_{\pm 1.7}^{**}$	69.3 $_{\pm 2.0}$	66.9 $_{\pm 1.2}$	62.7 $_{\pm 1.4}$	64.8 $_{\pm 1.7}$
- Δ	-0.7	-2.0	-2.1*	-0.9	-2.4
<i>es</i>	72.9 $_{\pm 1.4}^{**}$	70.2 $_{\pm 1.4}$	67.7 $_{\pm 1.7}$	61.6 $_{\pm 2.2}$	64.2 $_{\pm 2.5}$
- Δ	-2.7*	-2.3*	-2.8	-1.7	-1.2
<i>hi</i>	67.8 $_{\pm 1.7}^{**}$	65.3 $_{\pm 2.9}$	65.9 $_{\pm 2.2}$	61.2 $_{\pm 2.2}$	62.7 $_{\pm 2.6}$
- Δ	-3.5*	-2.8	-2.4	-2.7	-3.8
<i>zh</i>	69.9 $_{\pm 1.5}^{*}$	65.8 $_{\pm 3.2}$	67.0 $_{\pm 1.3}$	62.4 $_{\pm 3.1}$	68.1 $_{\pm 3.2}$
- Δ	-3.0*	-1.3	-1.6	-1.5	-1.4
Country Information: Qwen2-VL 72B					
<i>en</i>	70.7 $_{\pm 4.3}^{**}$	61.6 $_{\pm 3.6}$	64.1 $_{\pm 2.6}$	57.8 $_{\pm 4.4}$	58.4 $_{\pm 4.3}$
- Δ	-0.1	-0.8	-1.3	-0.2	+0.2
<i>de</i>	68.5 $_{\pm 2.2}^{**}$	62.9 $_{\pm 2.4}$	64.8 $_{\pm 2.0}$	56.8 $_{\pm 2.0}$	59.3 $_{\pm 2.4}$
- Δ	-1.6	-2.4	-1.6	-2.4	-2.0
<i>es</i>	69.9 $_{\pm 2.6}^{**}$	62.4 $_{\pm 3.9}$	64.3 $_{\pm 2.3}$	58.3 $_{\pm 3.0}$	57.6 $_{\pm 3.6}$
- Δ	-1.3	-1.0	-1.8	-1.0	-0.9
<i>hi</i>	62.9 $_{\pm 2.3}^{*}$	58.7 $_{\pm 4.0}$	62.2 $_{\pm 3.9}$	56.1 $_{\pm 3.2}$	54.7 $_{\pm 3.7}$
- Δ	-1.6	+0.3	+0.4	+0.0	-0.1
<i>zh</i>	63.9 $_{\pm 2.0}^{*}$	58.1 $_{\pm 2.7}$	61.9 $_{\pm 2.4}$	54.3 $_{\pm 3.1}$	57.2 $_{\pm 3.4}$
- Δ	-2.8	-0.8	-1.7	-2.9	-2.7

Table 8: Evaluation of injecting the country information. Δ represents the difference between the prompt with country information injection and those without it. The asterisk (*) in the Δ row shows significant difference.

Even when the prompt’s language is altered, the model continues to show high alignment with US labels. All variations significantly outperform the lowest-performing cultural variant with US culture label, and, for GPT-4o, they even significantly surpass the second-highest. This reinforces the idea that the **models are more aligned with US cultural norms, even when prompted in the dominant language of another culture.**

5.5 Adding Country Information

Building on Lee et al. (2024), we align VLMs with the target culture by adding country information to the prompt. We report results only for the +CAPT setup and best models, as shown in Table 8.

Injecting country information generally decreases performance across target cultures, with the exception of Qwen2 in Hindi, which shows no change. For instance, adding “Germany” to the prompt results in a 2.0 and 2.4-point accuracy drop for GPT-4o and Qwen2-VL, respectively, though these decreases are not statistically significant. Therefore, we conclude that **adding country information does not positively impact performance in the target culture.**

6 Conclusion

We present the first multimodal, parallel, multi-lingual hate speech dataset, annotated by a multicultural set of annotators. This dataset contains 300 parallel meme samples across five languages and has been annotated for hate speech across five cultures. We show that cultural factors significantly impact multimodal hate speech annotations in our dataset. Additionally, we use this dataset to highlight that VLMs exhibit a strong cultural bias towards the US, independent on the image and prompt language.

Limitation

While our dataset contains 300 samples across 5 languages—amounting to 1500 memes in total—the relatively small size reflects the challenges of generating high-quality translations and culturally diverse annotations. Expanding such a dataset is resource-intensive, both in terms of cost and labor.

Additionally, the dataset was sourced from a single website, primarily in English, and does not specifically target content from various cultural contexts. Furthermore, by selecting annotators from Prolific and using a single language per country, we introduce a degree of selection bias, as this method may not fully represent the complex cultural landscapes within each country.

We also recognize that equating culture with country is a limitation, as countries are often multicultural and multiethnic. For example, India is home to thousands of ethnic and tribal groups (Thapar et al., 2024), and our approach does not fully capture this diversity.

Finally, while our work highlights cross-cultural differences in the perception of hate speech, understanding the root causes of these disagreements remains an open challenge. Although we offer a qualitative analysis of annotator disagreements, a comprehensive theory-driven analysis is still lacking. Developing a robust theoretical framework to explain these cultural variations could ultimately help the alignment of VLMs with specific cultural nuances, leading to more accurate and culturally sensitive hate speech detection systems.

Ethics Statement

The annotators recruited through Prolific were compensated at a rate of £10.65 per hour, in alignment with the minimum wage in the authors’ country,

ensuring fair payment. Prior to the start of annotations, the project received ethical approval from the lead author’s institution. All annotators were thoroughly informed about the nature of the project, including warnings regarding potentially harmful and offensive content. Each annotator provided explicit consent before beginning their work, ensuring they were fully aware of the content and the purpose of their involvement.

We also acknowledge the potential risks associated with distributing our dataset. To mitigate these risks, we will establish clear terms of use that strictly prohibit any form of malicious exploitation. Additionally, we release the dataset in an anonymized format, ensuring that all user IDs and any personally identifiable information are removed to protect individual privacy.

We use AI assistants, specifically GPT-4o, to help edit sentences in our paper writing. Multi³Hate is licensed under CC BY-NC-ND 4.0.

Acknowledgement

The work of Minh Duc Bui and Katharina von der Wense is funded by the Carl Zeiss Foundation, grant number P2021-02-014 (TOPML project). The work of Anne Lauscher is funded under the Excellence Strategy of the German Federal Government and the Federal States. We thank Sukananya Purkayastha, Pranav A, Yujie Ren, Zhu Luan, Timm Dill, Carlos Galarza, and Delia Rieger for helping with translations and feedback on non-English text. We also thank Kyung Eun Park, Carolin Holtermann and Abteen Ebrahimi for their helpful feedback and discussions.

References

Lucas Beyer. 2024. On the speed of ViTs and CNNs. <http://lb.eyer.be/a/vit-cnn-speed.html>.

Aashish Bhandari, Siddhant B. Shah, Surendrabikram Thapa, Usman Naseem, and Mehwish Nasim. 2023. Crisishatem: Multimodal analysis of directed and undirected hate speech in text-embedded images from russia-ukraine conflict. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1994–2003.

Ralf Bosen. 2020. *New year’s eve in cologne: 5 years after the mass assaults*. DW. Accessed: 2024-10-10.

Olena Burda-Lassen, Aman Chadha, Shashank Goswami, and Vinija Jain. 2024. *How culturally aware are vision-language models?* *Preprint*, arXiv:2405.17475.

Yong Cao, Wenyan Li, Jiaang Li, Yifei Yuan, Antonia Karamolegkou, and Daniel Hershcovich. 2024. *Exploring visual culture awareness in gpt-4v: A comprehensive probing*. *Preprint*, arXiv:2402.06015.

William Cavnar and John Trenkle. 2001. N-gram-based text categorization. *Proceedings of the Third Annual Symposium on Document Analysis and Information Retrieval*.

Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. 2024. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*.

Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. 2023. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*.

Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. *Marked personas: Using natural language prompts to measure stereotypes in language models*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1504–1532, Toronto, Canada. Association for Computational Linguistics.

Alex Clark. 2015. *Pillow (pil fork) documentation*.

Ralph D’Agostino and E. S. Pearson. 1973. *Tests for departure from normality. empirical results for the distributions of b2 and root b1*. *Biometrika*, 60(3):613–622. Accessed 8 Oct. 2024.

Christoph Demus, Jonas Pitz, Mina Schütz, Nadine Probol, Melanie Siegel, and Dirk Labudde. 2022. *Detox: A comprehensive dataset for German offensive language and conversation analysis*. In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 143–153, Seattle, Washington (Hybrid). Association for Computational Linguistics.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, and et al. 2024. *The llama 3 herd of models*. *Preprint*, arXiv:2407.21783.

EVS/WVS. 2022. *Joint evs/wvs 2017-2022 dataset*. GESIS, Cologne. ZA7505 Data file Version 4.0.0.

Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, and et al. 2024. *Gemini 1.5: Unlocking multimodal*

- understanding across millions of tokens of context. *Preprint*, arXiv:2403.05530.
- Goran Glavaš, Mladen Karan, and Ivan Vulić. 2020. **XHate-999: Analyzing and detecting abusive language across domains and languages**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6350–6365, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Darina Gold, Piush Aggarwal, and Torsten Zesch. 2021. **Germemehate: A parallel dataset of german hateful memes translated from english**. In *Multimodal Hate Speech Workshop 2021*, pages 1–6.
- R. Gomez, J. Gibert, L. Gomez, and D. Karatzas. 2020. **Exploring hate speech detection in multimodal publications**. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1459–1467.
- Conrad Hackett, Brian Grim, Marcin Stonawski, Vegard Skirbekk, Michaela Potančoková, and Guy Abel. 2012. *The Global Religious Landscape: A Report on the Size and Distribution of the World's Major Religious Groups as of 2010*. Pew Research Center.
- Eftekhari Hossain, Omar Sharif, and Mohammed Moshui Hoque. 2022. **MUTE: A multimodal dataset for detecting hateful memes**. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 32–39, Online. Association for Computational Linguistics.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. **Lora: Low-rank adaptation of large language models**. *Preprint*, arXiv:2106.09685.
- Instituto Cervantes. 2023. **El español en el mundo. informe 2023**. Pages 7-9. Survey conducted by Instituto Cervantes and Various sources (national statistics agencies).
- 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.
- Antonia Karamolegkou, Phillip Rust, Yong Cao, Ruixiang Cui, Anders Søgaard, and Daniel Hershcovich. 2024. **Vision-language models under cultural and inclusive considerations**. *Preprint*, arXiv:2407.06177.
- Md. Rezaul Karim, Sumon Kanti Dey, Tanhim Islam, Md. Shajalal, and Bharathi Raja Chakravarthi. 2022. **Multimodal hate speech detection from bengali memes and texts**. In *International conference on Speech & Language Technology for Low-resource Languages (SPELLL)*, pages 1–15. SPELLL.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. **The hateful memes challenge: Detecting hate speech in multimodal memes**. In *Advances in Neural Information Processing Systems*, volume 33, pages 2611–2624. Curran Associates, Inc.
- Fajri Koto, Nurul Aisyah, Haonan Li, and Timothy Baldwin. 2023. **Large language models only pass primary school exams in Indonesia: A comprehensive test on IndoMMLU**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12359–12374, Singapore. Association for Computational Linguistics.
- Klaus Krippendorff. 2011. **Computing krippendorff's alpha-reliability**. *Technical Report*.
- Nayeon Lee, Chani Jung, Junho Myung, Jiho Jin, Jose Camacho-Collados, Juho Kim, and Alice Oh. 2024. **Exploring cross-cultural differences in English hate speech annotations: From dataset construction to analysis**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4205–4224, Mexico City, Mexico. Association for Computational Linguistics.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. **Llava-onevision: Easy visual task transfer**. *Preprint*, arXiv:2408.03326.
- Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. 2021. **Visually grounded reasoning across languages and cultures**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10467–10485, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021. **Hatexplain: A benchmark dataset for explainable hate speech detection**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17):14867–14875.
- Martina Miliani, Giulia Giorgi, Ilir Rama, Guido Anselmi, and Gianluca Leboni. 2020. **DANKMEMES @ EVALITA 2020: The Memeing of Life: Memes, Multimodality and Politics**, pages 1–. Accademia University Press.
- Hamdy Mubarak, Sabit Hassan, and Shammur Absar Chowdhury. 2022. **Emojis as anchors to detect arabic offensive language and hate speech**. *Preprint*, arXiv:2201.06723.
- Shravan Nayak, Kanishk Jain, Rabiul Awal, Siva Reddy, Sjoerd van Steenkiste, Lisa Anne Hendricks, Karolina Stańczak, and Aishwarya Agrawal. 2024. **Benchmarking vision language models for cultural understanding**. *Preprint*, arXiv:2407.10920.

- R.E. Nisbett. 2003. *The Geography of Thought: How Asians and Westerners Think Differently— and why*. Nicholas Brealey.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, and et al. 2024. *Gpt-4 technical report*. Preprint, arXiv:2303.08774.
- Pew Research Center. 2024. Cultural issues and the 2024 election. Technical report, Pew Research Center. Accessed June 2024.
- Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurowsky, and Michael Wojatzki. 2016. *Measuring the reliability of hate speech annotations: The case of the european refugee crisis*. ArXiv, abs/1701.08118.
- Nakatani Shuyo. 2010. *Language detection library for java*.
- Shardul Suryawanshi, Bharathi Raja Chakravarthi, Michael Arcan, and Paul Buitelaar. 2020. *Multimodal meme dataset (MultiOFF) for identifying offensive content in image and text*. In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 32–41, Marseille, France. European Language Resources Association (ELRA).
- R. Thapar, A.L. Srivastava, Sanjay Subrahmanyam, Stanley A. Wolpert, T.G. Percival Spear, Joseph E. Schwartzberg, Sanat Pai Raikar, K.R. Dikshit, Muzaffar Alam, Philip B. Calkins, Frank Raymond Allchin, and R. Champakalakshmi. 2024. *India*. *Encyclopedia Britannica*. Accessed: 2024-10-08.
- Harry C. Triandis. 1995. *Individualism and Collectivism*, 1st edition. Routledge.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. *Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution*. Preprint, arXiv:2409.12191.
- Frank Wilcoxon. 1945. *Individual comparisons by ranking methods*. *Biometrics Bulletin*, 1(6):80–83.
- World Population Review. 2024. *German speaking countries 2024*. Accessed: 2024-10-08.
- Ankit Yadav, Shubham Chandel, Sushant Chatufale, and Anil Bandhakavi. 2023. *Lahm : Large annotated dataset for multi-domain and multilingual hate speech identification*. Preprint, arXiv:2304.00913.
- Andre Ye, Sebastin Santy, Jena D. Hwang, Amy X. Zhang, and Ranjay Krishna. 2024. *Computer vision datasets and models exhibit cultural and linguistic diversity in perception*. Preprint, arXiv:2310.14356.

A Dataset Construction Details

A.1 Topic List

We explain, how we further divide each sociopolitical category into smaller topics: (1) **Religion**, divided into the world’s major religions (Hackett et al., 2012); (2) **Nationalities**, aligned with our target countries; (3) **Ethnicity**, structured as outlined in Cheng et al. (2023); (4) **LGBTQ+**, representing the groups denoted by the acronym; and (5) **Political Issues**, identified as cultural issues during the US election, as defined by Pew Research Center (2024).

A.2 Keyword Matching

Following keyword matching, we retain only the templates with at least 10 captions (after pre-filtering). Additionally, each topic must have a minimum of three templates that meet this criterion to ensure a diverse set of templates per topic. Topics that do not meet these requirements are filtered out.

A.3 Pre-Filtering

A.3.1 Filtering for English Captions

To filter out non-English captions, we utilize the implementation by Shuyo (2010), which employs a Naive Bayes approach based on n-grams (Cavnar and Trenkle, 2001).

A.3.2 Filtering for Multimodal Hate Speech

We outline our method for filtering potentially multimodal hate speech samples by comparing two types of classifications: (1) user captions combined with manually created image descriptions and (2) the user captions alone. By comparing the outcomes of these two classifications, we identify content as multimodal hate speech when the first case (caption + image description) is flagged as hate speech, but the second case (caption only) is not.

Experimental Setup We employ zero-shot learning with Llama 3 (Dubey et al., 2024) to detect hate speech. For the image descriptions, two annotators were tasked with generating descriptions for all meme templates. A sample is only classified as hate speech in scenario (1) if both image descriptions+caption are classified as hate speech by the model. The prompt used is reported in Figure 6.

A.3.3 Filtering for Wordplay

To ensure captions are easily translatable and avoid noise from wordplay, two fluent English speakers

	USA	Germany	Mexico	India	China
No. of Annotators	105	103	101	66	70
Gender (%)					
male	53.33	49.51	52.48	53.03	48.57
female	46.67	50.49	47.52	46.97	51.43
non-binary	–	–	–	–	–
Ethnicity (Simplified) (%)					
Asian	1.9	–	–	93.94	97.14
Black	15.24	–	–	–	–
White	82.86	100.0	28.71	–	1.43
Mixed	–	–	41.58	–	1.43
Other	–	–	29.7	6.06	–
Level of Education (%)					
Below High School	–	–	0.99	–	–
High School	24.76	22.33	9.9	3.03	–
College	20.95	16.5	17.82	3.03	7.14
Bachelor	37.14	38.83	60.4	71.21	32.86
Master's Degree	14.29	19.42	9.9	22.73	48.57
Doctorate	2.86	2.91	0.99	–	11.43
Age (%)					
18-19	1.9	4.85	1.98	9.09	1.43
20-29	28.57	64.08	74.26	60.61	60.0
30-39	35.24	21.36	22.77	16.67	30.0
40-49	17.14	8.74	–	9.09	5.71
50-59	13.33	0.97	0.99	4.55	1.43
60-69	3.81	–	–	–	1.43
70-79	–	–	–	–	–
80-89	–	–	–	–	–
Political Orientation (%)					
Liberal/Progressive	35.24	37.86	34.65	25.76	8.57
Moderate Liberal	20.95	32.04	26.73	18.18	22.86
Independent	22.86	15.53	13.86	43.94	37.14
Moderate Conservative	10.48	12.62	14.85	9.09	12.86
Conservative	8.57	0.97	3.96	1.52	–
Other	1.9	0.97	5.94	1.52	18.57
Religion (%)					
None	42.86	70.87	42.57	–	74.29
Christian	46.67	28.16	37.62	6.06	5.71
Buddhism	–	–	–	3.03	8.57
Islam	1.9	–	–	21.21	–
Judaism	–	–	–	–	–
Hinduism	–	–	–	60.61	–
Irreligion	2.86	0.97	8.91	1.52	2.86
Other	5.71	–	10.89	7.58	8.57

Table 9: Demographics of annotators during the hate speech annotation phase.

classify each caption as either non-wordplay or wordplay. Only captions unanimously classified as non-wordplay are selected.

A.4 Translation Stage Details

To translate the captions by the Google Translate API, each caption is inputted with a separator “//” to clearly distinguish between the top and bottom text.

Furthermore, each human translator is provided with a detailed annotation guide outlining the criteria for what constitutes a correct translation and how corrections should be made. The annotation guidelines are shown in Figure 7.

A.5 Pre-Annotation

Before beginning the main annotation process, we conduct a pre-annotation stage to balance the dataset. For this phase, we create a parallel multilingual meme dataset consisting of 450 samples. The dataset is divided into five equal parts, with each

Classify the following meme into 'hatespeech' or 'no hatespeech'. The definition of hatespeech is: <definition> Be culturally aware and sensitive. Meme Image: <image description> Text: <caption>

Figure 6: Prompt for classification of multimodal hate speech in the pre-filtering stage.

Correct Translation (1):

- The translation should accurately convey the same meaning as the original text.
- It does not have to be grammatically perfect (memes are rarely grammatically correct), but it should be easily understandable in the target language.
- Ensure that the “<sep>” marker is placed in a position that maintains the same semantic meaning as in the original context (if possible).
- It should maintain the appropriate context and style for the given content.
- Proper cultural references and idiomatic expressions should be correctly translated or adapted as needed.

False Translation (0):

- The translation significantly distorts or changes the meaning of the original text.
- It contains major grammatical errors that hinder understanding.
- It fails to convey the intended context, tone, or style of the original text.
- Key terms or names are mistranslated or omitted.
- Important nuances or details from the original text are lost or incorrectly translated.

Figure 7: Annotation guidelines for translators.

part assigned to annotators from a different cultural background. Each sample is annotated twice, with a total of 50 annotators involved—10 from each cultural group.


To achieve balance, we adjust the dataset so that 40% of the samples are labeled as hate speech, 40% as non-hate speech, and 20% where there was a tie between annotators. This process results in the final set of 300 samples.

A.6 Hate Speech Survey Design

We use Google Forms¹⁰ to design and distribute our surveys. To create surveys in each target language, we use the v3 Google Translate API to translate the

¹⁰<https://www.google.de/intl/de/forms/about>

Does this image contain hate speech? (Question 10 of 20, Filename: 6167601.jpg)



Please provide your feedback for Question 10 (6167601.jpg): *

Hate Speech
 Non-Hate Speech
 I Don't Know

Figure 8: An example of an explicit attention check used in our survey.

surveys, which are then reviewed and corrected by native speakers for accuracy. We then create fixed random parallel batches, which are then assigned to each annotator.

Each batch includes four attention checks: one explicit check, where annotators are required to select a specific pre-defined answer (see Figure 8), and three implicit checks. These implicit checks consist of samples presented as examples at the beginning of the survey, accompanied by explanations of why the samples are classified as non-hate speech or hate speech based on the given definition.

We only retain annotations where the explicit attention check is answered correctly, and at least two out of the three implicit checks are passed. After collecting five annotations per sample, we review the results for any ties that need resolution and create new batches accordingly.

A.7 Terms of Use

Our research is conducted in the public interest under the GDPR, fulfilling the conditions for substantial public interest as academic research. We were unable to locate any Terms of Service on <https://memegenerator.net>, and the contact information provided on the website appears to be outdated and non-functional. To ensure we respect the platform’s rights, we are publishing Multi³Hate under the CC BY-NC-ND 4.0 license.

A.8 Time Required for Dataset Development

Estimating the effort required to create such a dataset is challenging due to the multiple, often unforeseen, refinement stages involved. For instance, unexpected challenges—such as translating wordplay—necessitated filtering them out and revisiting

<PREFIX>:
Definition of Hate Speech:

Any kind of communication [...], that attacks or uses pejorative or discriminatory language with reference to a person or a group on the basis of who they are, in other words, based on their religion, ethnicity, nationality, race, colour, descent, gender or other identity factor.

<PROMPT #1>:

Answer if this meme is hate or not (in <country>) with a single alphabet letter among given answer choices a and b.

<PROMPT #2>:

Given a meme from online, answer if it is 'Hate Speech' or 'Non-Hate Speech' (in <country>). Answer with only one single alphabet between 'a' and 'b', with no other description.

<PROMPT #3>:

Would you consider this meme as a hate speech (in <country>)? Give an answer in a single alphabet only.

<POSTFIX>:

Meme: <image>
(Caption inside the meme image: <caption>)
Choose:
a: Hate
b: Non-hate

Figure 9: All three prompt variations: The order of options (a) and (b) is switched to create a total of six variations. Brackets are optional, allowing for insertion of the caption (+CAPT Setting) or country information as described in Section 5.5.

the translation process, significantly increasing the time and effort required. Overall, the entire dataset creation process took approximately four months from start to finish.

B Experiments Details

B.1 Model Details

Table 12 lists the models and their sizes, all run on three H100 GPUs. Each large VLM processes five languages in around 1.5 hours. To support better text extraction from memes, images are resized to 512x512 pixels (Beyer, 2024). For all models, we generate deterministic outputs and limit generation to 40 new tokens. For the Gemini 1.5 Pro model,

GT Inp.	US	DE	MX	IN	CN
Unimodal					
GPT-4o	65.4 ± 3.9	58.3 ± 4.4	60.8 ± 3.5	56.7 ± 4.0	55.0 ± 4.5
Gemini 1.5 Pro	63.5 ± 3.9	59.2 ± 5.5	60.1 ± 5.2	56.8 ± 5.1	54.6 ± 4.4
InternVL2 76B	59.2 ± 3.1	52.1 ± 3.0	60.8 ± 2.6	56.8 ± 4.1	55.0 ± 3.4
LLaVA OneVision 73B	61.8 ± 1.5	54.3 ± 1.3	56.5 ± 1.1	51.7 ± 1.1	49.7 ± 1.8
Qwen2-VL 72B	61.9 ± 2.5	53.7 ± 2.9	56.2 ± 2.1	51.5 ± 2.8	49.4 ± 2.9
< 10B Models					
Qwen2-VL 7B	67.4 ± 2.0	67.2 ± 4.1	68.2 ± 2.6	62.3 ± 3.4	65.3 ± 3.9
InternVL2 8B	58.2 ± 2.8	57.4 ± 4.0	58.7 ± 4.9	57.6 ± 4.8	56.7 ± 4.2
Llava OneVision 7B	60.3 ± 5.5	55.6 ± 7.3	58.3 ± 6.9	54.1 ± 6.9	52.3 ± 7.4

Table 10: Unimodal setting: Models only get the caption as the input. The best value in each column is bolded. Sorted by average accuracy.

we disable all safety settings to minimize rejected responses.

To derive binary classifications from the answers, we implement a custom keyword extraction. We relax the constraints on possible answers significantly, moving beyond a binary choice of “a” or “b”. For instance, “non-hate” or “hate-speech” is also recognized as a valid response in our analysis. However, answers that are nonsensical are counted as incorrect.

B.2 Prompts

We design prompts similar to those in Lee et al. (2024) and present the various prompt formulations in Figure 9. Additionally, the multilingual version of these prompts is shown in Figure 11.

C Supervised Baseline

We experiment with a small supervised VLM baseline, training the Qwen2-VL 7B separately on each cultural label (i.e., culture-specific models) on only one prompt variation, using a 3-fold cross-validation approach due to the relatively small dataset size. We fine-tune our model using LoRA (Hu et al., 2021), modifying only the Query and Value matrices, with a rank of 8 and an alpha value of 16. We employ a learning rate of $2e - 4$ with a constant learning rate schedule, training for three epochs with a batch size of 16. Additionally, we refrain from hyperparameter tuning to avoid overfitting on our validation folds, as our limited data prevents the creation of a separate test set. The results are presented in Table 11.

Overall, we observe high improvements across all countries, with the most notable gain of 7.1 points in Indian annotations. Moreover, the highest

GT Inp.	US	DE	MX	IN	CN
Zero-Shot	67.4	67.2	68.2	62.3	65.3
Fine-Tuned - US	71.5	69.8	70.9	65.9	66.9
Fine-Tuned - DE	70.1	74.1	71.9	67.5	70.9
Fine-Tuned - MX	69.2	73.2	72.3	68.1	70.4
Fine-Tuned - IN	69.5	72.1	71.0	69.4	71.8
Fine-Tuned - CN	70.1	72.3	70.2	68.8	71.9

Table 11: We compare the zero-shot performance of Qwen2-VL 7B with a version fine-tuned separately for each cultural label using supervised learning and 3-fold cross-validation. We bold the best performance on each cultural label.

level of agreement with each culture-specific annotation is achieved by the model fine-tuned on the respective cultural labels. This suggests that each country’s cultural perception can be effectively improved through supervised fine-tuning. We leave a more thorough analysis of the results and the exploration of advanced fine-tuning approaches for future work.

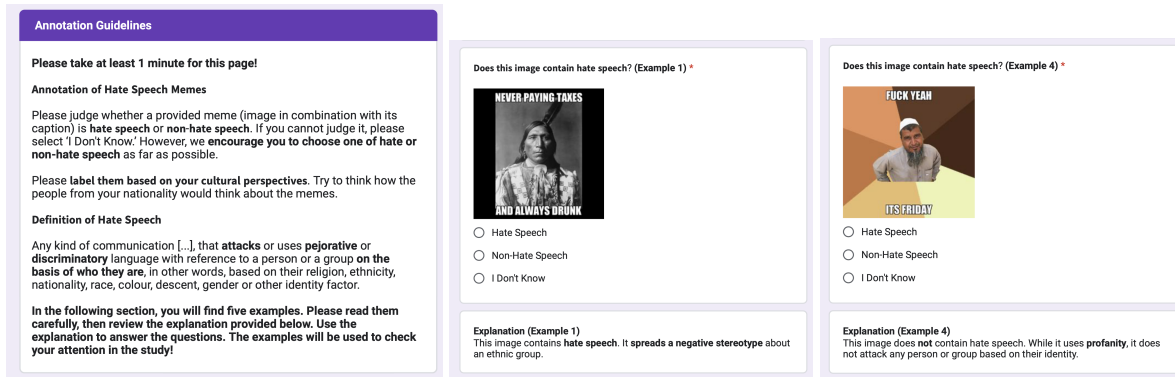


Figure 10: The hate speech guideline interface displayed to annotators before they begin their annotations, along with two out of five example cases.

Lang.	PREFIX	PROMPT #1	PROMPT #2	PROMPT #3	POSTFIX
DE	Definition von Hassrede: Jede Art von Kommunikation [...], die eine abwertende oder diskriminierende Sprache in Bezug auf eine Person oder eine Gruppe aufgrund ihrer Identität angreift oder verwendet, mit anderen Worten aufgrund ihrer Religion, ethnischen Zugehörigkeit, Nationalität, Rasse, Hautfarbe, Abstammung, Geschlecht oder ein anderer Identitätsfaktor.	Beantworte, ob dieses Meme Hass oder nicht ist, mit einem einzelnen Buchstaben der vorgegebenen Antwortmöglichkeiten a und b.	Anhand eines Memes aus dem Internet, beantworte ob es sich um eine "Hassrede" oder eine "Nicht-Hassrede" handelt. Antworten Sie mit nur einem Buchstaben zwischen a "und b "und ohne weitere Beschreibung.	Würdest du dieses Meme als Hassrede betrachten? Gebe deine Antwort nur in einem Buchstaben an.	Meme: <image> Beschriftung im Meme-Bild: <caption> Wähle: a: Hass b: Kein Hass
ES	Definición de discurso de odio: Cualquier tipo de comunicación [...], que ataque o utilice un lenguaje peyorativo o discriminatorio con referencia a una persona o un grupo en función de quiénes son, en otras palabras, en función de su religión, etnia, nacionalidad, raza, color, ascendencia, género u otro factor de identidad.	Responda si este meme es odio o no con una sola letra del alfabeto entre las opciones de respuesta a y b.	Dado un meme de Internet, responda si es un 'discurso de odio' o 'discurso sin odio'. Responda con una sola letra entre 'a' y 'b', sin ninguna otra descripción.	¿Considerarías este meme como un discurso de odio? Responde con una sola letra del alfabeto.	Meme: <image> Título dentro de la imagen del meme: <caption> Elige: a: Odio b: Sin odio
HI	घृणास्पद भाषण की परिभाषा: किसी भी प्रकार का संचार [...], जो किसी व्यक्ति या समूह के संदर्भ में उनके धर्म, जातीयता, राष्ट्रियता, नस्ल, रंग, वंश, लिंग या अन्य पहचान कारक के आधार पर अपमानजनक या भेदभावपूर्ण भाषा का उपयोग करता है।	दिए गए उत्तर विकल्पों क और ख में से किसी एक अक्षर से बताइए कि यह मीम घृणास्पद भाषण है या नहीं।	ऑनलाइन से दिए गए किसी मेमे से, उत्तर दें कि यह 'घृणास्पद भाषण' है या 'गैर-घृणास्पद भाषण' है। 'क' और 'ख' के बीच केवल एक ही अक्षर से उत्तर दें, कोई अन्य विवरण न दें।	क्या आप इस मीम को 'घृणास्पद भाषण' मानेंगे? इसका जवाब सिर्फ एक अक्षर में दीजिए।	मीम: <image> मीम इमेज के अंदर कैप्शन: <caption> चुनें: क: घृणा ख: गैर-घृणास्पद
ZH	仇恨言论的定义: 在任何形式的交流中 [...], 基于个人或群体的身份进行攻击或使用贬义或歧视性语言。换句话说, 基于他们的宗教、民族、国籍、种族、肤色、血统、性别或其他身份因素, 进行攻击或使用贬义或歧视性语言。	请判断该表情包是否包含仇恨, 根据选项用字母 a 或 b 作答	根据网图判断其是否为仇恨言论。用字母 a 或 b 作答, 请勿在答案中输入其他文本	您是否认为该表情包属于仇恨言论? 请用单个字母给出答案	表情包: <image> 表情包所配文字: <caption> 请选择: a: 仇恨 b: 非仇恨

Figure 11: We present all multilingual prompts after removing any spaces, which correspond analogously to the English version shown in Figure 9.

Model	Total
LLaVA-Onevision 7B	8.03B
LLaVA-Onevision 72B	73.2B
Qwen2-VL-7B-Instruct	8.29B
Qwen2-VL-72B-Instruct	73.4B
InternVL2-8B	8.08B
InternVL2-Llama3-76B	76.3B
Gemini 1.5 Pro	?
GPT-4o	?

Table 12: We present the models used in this study, along with their respective total number of parameters (denoted as “|Total”). Each model name is hyperlinked to its corresponding Huggingface repository (when viewed digitally). For Gemini 1.5 Pro and GPT-4o, we use gemini-1.5-pro-001 and gpt-4o-2024-05-13, respectively.

Topic	Keywords	Count
Christianity	christ, jesus, priest	21
Islam	muslim, islam	22
Hinduism	hindu, hinduism	–
Buddhism	buddha, buddhist	–
Folk Religion	folk religion	–
Judaism	jew, judaism	18
Germany	germany, german	18
United States	america, usa, american	21
Mexico	mexico, mexican	20
China	china, chinese	21
India	india, indian	15
Asian	asia, asian	20
Black	black	23
Latine	latino, latine	–
Middle Eastern	middle+eastern, arab	19
White	white	19
Lesbian	lesbian	–
Gay	gay	–
Bisexual	bisexual	–
Transgender	trans, transgender	19
Queer	queer	–
Law Enforcement	police	23
Feminism	feminist	21
Immigration	immigrants	–
Racial Diversity	(already included)	–
LGBTQ+	(already included)	–

Table 13: We conduct a keyword search based on the identified topics and report the final sample count in Multi³Hate. A “–” indicates that the topic did not meet our requirements.

Category	Example	Keywords
Historical & Political Context	<i>In the U.S. it is much more acceptable due to Cold War politics, for individuals to decry Chinese communism as completely evil. This legacy did not affect India [...]</i>	Historical Context of Hindu-Muslim Conflicts Historical Context of Discrimination against Latinos Historical Context of Anti-Semitism Historical Context of Colonialism Historical Context of Colorism Historical Context of Communism Historical Context of Communal Violence Historical Context of Germans Political Tensions
Sensitivity Around Minority Groups	<i>From the US context, this can be seen as mocking an immigrant or a person of color. From the Indian standpoint, this is not considered mocking as brown people are not a marginalized group in India.</i>	Sensitivity towards Immigrants Women's Right Sensitivity to Class Distinctions LGBTQ+ Acceptance and Rights Perceptions of Arab Identity Perception of Indian People Perception of Black Identity Perception of Racial Profiling <i>Attack against Religion as Minority Group</i> <i>Minority Group is Majority Group in other Culture</i>
Social Norms & Cultural Values	<i>In the US culture, parents are expected to follow the society's code of conduct towards the kids. In the Indian context, the father is the patriarch and can discipline the kids. This statement is insulting to the father.</i>	Social Norms Around Nudity Social Norms Around Transportation Social Norms Around Diet Social Norms around (Patriarchal) Family Structure Cultural Norms of Politeness Cultural Norms Around Nudity Cultural Perception of Governance Cultural Perception of Gun Laws Cultural Perception of Police Authority Cultural Perception of Democracy Cultural Perception of War Cultural Perception of Sexual Violence Cultural Perception of Hard Labor Cultural Perception of Freedom of Speech Cultural Sensitivity to Religion Cultural Context of Poverty
Non-Existing Stereotypes	<i>This meme uses the Asian stereotype [...] and hence is offensive in the US. This stereotype is non-existent in India [...]</i>	Non-Existing Stereotypes
Annotation Ambiguity	<i>[...] interviewers are not a protected minority [...]. I would have voted non-hate speech for both cultures.</i>	Hate Speech Annotation Ambiguity
Language Error	<i>The meaning translated to Hindi feels like [...] and can be reinterpreted [...].</i>	Translation Error

Table 14: We present the major themes of disagreement, along with their associated keywords and examples of annotators' comments.