

When Does Hate Transfer?

A Large-Scale Study of Cross-Lingual Hate Speech Detection

Zahra Bokaei

School of Informatics
University of Edinburgh
zahra.bokaei@ed.ac.uk

Walid Magdy

School of Informatics
University of Edinburgh
wmagdy@ed.ac.uk

Bonnie Webber

School of Informatics
University of Edinburgh
bonnie.webber@ed.ac.uk

Abstract

Effective detection of hate content is essential for maintaining healthy online communities, yet performance can vary substantially across languages. This work investigates a core practical question: Can training data from one language support hate-speech detection in another language with comparable performance? To answer this question, we conduct a large-scale study across 14 languages. We observe consistent patterns in which certain language pairs transfer more effectively than others. Among the factors we measure, linguistic similarity is the strongest predictor of these patterns, with temporal overlap and dataset-size difference playing smaller roles. These measurable factors do not fully explain the observed patterns; a complementary qualitative analysis suggests that shared cultural cues and communicative norms may be associated with some of the remaining cases, particularly within several clusters. We present this cultural interpretation as a hypothesis from qualitative inspection rather than a measured cause. Our analysis includes low-resource languages, offering practical insights into when cross-lingual transfer can support them and where its limitations emerge. These findings inform the design of more reliable multilingual hate-detection systems.

Warning: This paper includes examples of hate-speech for research purposes.

1 Introduction

Hate speech detection aims to identify harmful expressions in online text, such as hate-speech, harassment, and threats (Hoang et al., 2024). As digital communication expands globally, exposure to such content has increased, motivating efforts to build automated systems capable of supporting safer online communities. However, what counts as hate-speech is context-dependent, and even human annotators vary in how they perceive it across settings (Sap et al., 2022a). These differences pose

a challenge for developing models that generalize reliably across languages.

Large Language Models (LLMs) provide strong multilingual representations and are increasingly used for hate detection (Khondaker et al., 2023; Kumar et al., 2024; Abaskohi et al., 2024). However, cross-lingual performance remains uneven, and the factors shaping transfer effectiveness are under-explored (Meyer and Buys, 2024; Bokaei et al., 2026a). This motivates a practical question: under what conditions can training data from one language support hate-speech detection in another? Potential factors include dataset size, label definitions, temporal alignment, and linguistic similarity, while prior work suggests that social or regional context may also shape how hate expressions are formed and recognized (Bokaei et al., 2025; Pawar et al., 2024). We treat this last factor as a hypothesis to be explored qualitatively. In this study, we analyse hate-speech transfer across 14 languages from diverse regions and domains. We construct a cross-lingual transfer matrix by training classifiers on each language and evaluating them on all others. We then run two multilingual experiments: (1) integrating all languages into one training set, and (2) constructing a balanced set with equal representation per language. These setups allow us to examine how data volume and language balance affect cross-lingual generalisation.

The current study addresses the following research questions (RQs):

- RQ1. How does the performance of cross-lingual training compare to in-language training for hate-speech detection?
- RQ2. Would using training data from multiple languages improve performance compared to relying on a single source language?
- RQ3. What could be the potential reasons for the success or failure of transfer learning in hate-speech detection?

These questions clarify when cross-lingual transfer is effective for hate-speech detection and identify the conditions under which it succeeds or fails. Motivated by previous work (Bokaei et al., 2025), we conduct our main experiments with LLaMA-3-8B and additionally replicate them with Gemma-3-4B and Gemma-4-4B to verify that our findings generalize across model families. Our study covers 14 languages: Turkish, Urdu, Persian, Arabic, Hindi, Bengali, Indonesian, Malay, English, Spanish, French, German, Korean, and Chinese.

2 Related Work

Hate-Speech Detection. Early work on hate detection focused primarily on English, applying machine learning and deep learning approaches (Asogwa et al., 2022; Davidson et al., 2017; Mullah and Zainon, 2021; Malik et al., 2024; Zimmerman et al., 2018; Zhou et al., 2020; Roy et al., 2020; Zhang et al., 2018; Bade et al., 2024; Aiyanyo et al., 2020; Cao et al., 2020; Risch et al., 2020; Wang et al., 2020; Pamungkas and Patti, 2019; Van Hee et al., 2015; Guo and Gauch, 2024; Cano Basave et al., 2013). Later work extended to other languages, including Indonesian (Ibrohim and Budi, 2019), Danish (Sigurbergsson and Derczynski, 2020), Arabic (Mubarak et al., 2021; Bensalem et al., 2023), Korean (Jeong et al., 2022), Chinese (Deng et al., 2022), Greek (Pitenis et al., 2020), Persian (Delbari et al., 2024), and Hindi (Gupta et al., 2022; Kapoor et al., 2019).

Recent work highlights key challenges in multilingual hate detection, including comparisons of translation-based and multilingual pipelines (Bell et al., 2025), limitations in zero-shot transfer for taboo expressions (Nozza, 2021), and the influence of annotator identity and beliefs on hate judgments (Sap et al., 2022b). Multimodal datasets such as Multi3Hate reveal substantial cross-cultural variation in hate-speech perception (Bui et al., 2025), while surveys document strong geographic and cultural imbalances in existing hate-speech resources (Tonneau et al., 2024). Additional studies explore the feasibility of parallel multilingual hate-speech data via machine translation (Korre et al., 2024). Shared workshops such as OffensEval, HASOC, OSACT5, and GermEval continue to support benchmarking across languages (Zampieri et al., 2019; Mandl et al., 2019; Mubarak et al., 2022; Wiegand et al., 2018).

Cross-lingual Transfer Learning for Hate-Speech. Transfer learning has been widely ap-

plied across NLP tasks, including sentiment analysis (Fsih et al., 2022; Husain et al., 2022), irony detection (Golazizian et al., 2020), machine translation (Zoph et al., 2016; Kim et al., 2019; Adabara and Abdul-Mageed, 2021), fake news detection (Cruz et al., 2020), and offensive language classification (Zhou et al., 2023). In multilingual hate detection, prior work studies English-to-low-resource transfer using contextual embeddings (Ranasinghe and Zampieri, 2021), compares multilingual and translation-based Arabic–English pipelines (El-Alami et al., 2022), and improves low-resource performance through nearest-neighbor retrieval (Ghorbanpour et al., 2025). A recent survey of 67 studies provides a broader overview of cross-lingual offensive language detection methods and datasets (Jiang and Zubiaga, 2024).

Work on cultural or social influences on transfer remains limited, though related studies provide evidence. Zero-shot results show that language-specific taboo or colloquial expressions often fail to transfer (Nozza, 2021). Cross-lingual analyses of English, Chinese, and Korean reveal culture-specific biases affecting transfer, partly mitigated by few-shot adaptation (Zhou et al., 2023). Multimodal work also highlights cross-cultural variation in hate-speech judgments (Bui et al., 2025). For example, Zhou et al. (2023) find that Korean-to-Chinese transfer can outperform English-to-Chinese, suggesting that contextual or social alignment may support generalisation. Similarly, Bokaei et al. (2025) report stronger transfer into Persian from Arabic and Indonesian than from English. Our study extends this work by analysing transfer across 14 languages and quantifying how dataset, temporal, and typological factors relate to cross-lingual transfer effectiveness.

3 Dataset

We use publicly available datasets from peer-reviewed publications and shared tasks to ensure quality and comparability. The study covers 14 languages from diverse linguistic families and regions. Most datasets come from Twitter (X), with some including Facebook or YouTube, and collection periods span roughly 2016–2023. Table 1 summarises dataset sizes, training splits, hate-sample counts, sources, and collection periods.

All datasets are class-imbalanced, with varying hate/non-hate proportions. For controlled cross-language comparisons, we standardise training size where possible. Among sufficiently large datasets,

Language	Data Size	Train	Test	Train _{Hate}	Test _{Hate}	Source	Period	Ref.
Bengali (BN)	30000	9007	1002	4504	501	YT, FB	2017–2020	(Romim et al., 2020)
Urdu (UR)	10041	9000	1000	4500	300	X	–	(Rizwan et al., 2020)
Hindi (HI)	8192	4171	467	2222	241	X, FB, WA	–	(Bhardwaj et al., 2020)
Persian (FA)	7000	2060	230	1002	117	X	2020–2022	(Delbari et al., 2024)
Indonesian (IN)	13169	9005	1003	4502	502	X	2018	(Ibrohim and Budi, 2019)
Malay (MS)	5014	3403	378	1701	189	X	2022–2023	(Maity et al., 2023)
Arabic (AR)	15014	9005	1002	4503	502	X	2016–2018	(Albadi et al., 2018)
Turkish (TR)	52672	9010	1003	4505	502	X, IG	2016–2023	(Toraman et al., 2022)
Korean (KO)	9341	3065	300	1520	142	X	2018–2020	(Moon et al., 2020; Lee et al., 2022)
Chinese (ZH)	8065	4950	318	2470	151	Web	–	(Moosa and Najiba, 2022)
Spanish (ES)	29856	9000	1000	4731	625	X	2016–2023	(Tonneau et al., 2024)
German (DE)	58025	9008	1003	4503	502	X, FB	2016–2023	(Tonneau et al., 2024)
French (FR)	18072	7483	829	3739	413	X	2016–2023	(Tonneau et al., 2024)
English (EN)	360494	9000	1038	4500	502	FB, YT, RD, X	2016–2023	(Tonneau et al., 2024)

Table 1: Hate language detection datasets. FB = Facebook, YT = YouTube, X = Twitter, WA: WhatsApp, IG = Instagram, RD = Reddit. The definition of Hate and more details for each dataset are shown in Table 4

we use the smallest available hate count, approximately 4,500, and sample an equal number of non-hate instances, yielding balanced datasets of about 10,000 examples per language. We use an 80/10/10 train, validation, and test split. This isolates language effects from dataset volume, preventing high-resource languages such as English or German from appearing stronger due to larger training sets, while keeping computation tractable across 14 models. Balancing was applied only where enough hate examples were available; Hindi, Malay, Korean, Chinese, and Persian fell below this threshold.

To enable consistent evaluation across datasets with fine-grained annotation schemes (e.g., hostile, hate, offensive, abusive), we retain only hate-speech labels and treat all other categories as non-hate. The original label schemes and hate-speech definitions are summarized in Appendix Table 4. Most datasets adopt a broadly consistent definition of hate-speech as language attacking individuals or groups based on protected identity attributes.

4 Experimental Setup

We investigate cross-lingual transfer in hate detection using the datasets summarized in Table 1. Our experimental design includes (1) single-language fine-tuning with cross-lingual evaluation, (2) zero-shot baselines, and (3) two complementary multilingual fine-tuning setups to investigate alternative explanations for transfer behavior.

Model Selection and Baselines: We evaluate two types of baselines: (i) zero-shot baselines, including GPT-4o, Gemma-3-4B-Instruct, and LLaMA-3-8B-Instruct; and (ii) fine-tuning baselines, including LLaMA-3-8B, Gemma-3-4B, and Gemma-4-4B. We select LLaMA-3 and the

Gemma family because recent benchmarks identify them as strong open-weight choices for multilingual harmful-content detection ((Rathore and Chaurasia, 2025; Bokaei et al., 2026b; Raquib et al., 2025)), and because they are released under open licenses that permit fine-tuning across all 14 languages on our available hardware. We include two Gemma generations (Gemma-3 and Gemma-4) to test whether transfer structure is stable across model versions as well as across the LLaMA/Gemma architecture divide. The fine-tuned baselines outperform all zero-shot baselines across the 14 languages, with LLaMA-3-8B and Gemma-4-4B yielding almost identical average macro-F1 scores, both higher than Gemma-3-4B. We therefore adopt LLaMA-3-8B as our primary model and replicate all main experiments (Solo-Source TL, Global-Imbalanced, Global-Balanced, and Observed Cluster) with Gemma-3-4B, and Gemma-4-4B to verify that our findings are not specific to one model family. Baseline results are in Table 5; full strategy comparisons are in Tables 8, and 9 (Gemma-3 and Gemma-4); full transfer matrices are in Tables 10, and 11 (all in the Appendix). As discussed in Section 5, the same qualitative patterns hold across all models, suggesting that the observed transfer structure is not model-specific. A leave-one-language-out analysis on the three largest clusters further confirms the robustness of these gains (Appendix 10).

Fine-tuning Configuration: All fine-tuning uses the official LLaMA-3-8B pretrained checkpoint. We apply AdamW with a learning rate of 1e-5, batch size 8, and a maximum sequence length of 512 tokens. Training proceeds for up to 10 epochs with early stopping (patience = 3) based on validation F1 score. All experiments are run with five




Figure 1: Cross-lingual F1 scores across 14 languages.

random seeds, and we report the mean and standard deviation of F1 across runs to assess result stability. All experiments are conducted on a single NVIDIA Quadro RTX 8000 (48 GB). Performance is reported using macro-F1.

Single-Language Fine-Tuning: For each language, we fine-tune a separate LLaMA-3-8B model on its own training set and evaluate it on all 14 test sets. For example, a model fine-tuned on English is evaluated on English and on the remaining 13 languages. This procedure produces a 14×14 cross-lingual transfer matrix, allowing us to examine empirical transfer patterns. Because LLaMA is a multilingual pretrained model, fine-tuning on one language does not erase its multilingual knowledge; rather, it specializes the model toward the hate-speech patterns of the target language.

The 14×14 single-language fine-tuning experiment directly addresses RQ1, revealing which language pairs transfer well or poorly. To address RQ2, we design two complementary multilingual experiments: **Global Integration (Imbalanced)**, which tests whether greater data volume and multilingual diversity improve generalization, and **Global Integration (Balanced)**, which removes dataset-size differences as a confounding variable. Together, these setups isolate the effects of data size, multilingual exposure, and the structural similarities surfaced by the 14×14 matrix.

Global Integration (Imbalanced) Experiment: We merge the full training sets from all 14 languages into one multilingual dataset and fine-tune a single LLaMA-3-8B model. The model is then evaluated separately on each language’s test set. This experiment tests whether greater data volume and diversity produce better generalization.

Global Integration (Balanced) Experiment:

To control for dataset-size differences, we sample an equal number of hate ($n = 1000$) and non-hate ($n = 1000$) instances from each language and merge them into a balanced multilingual dataset. A single model is fine-tuned and evaluated on each language’s test set, isolating the effect of balanced representation by controlling for dataset size.

Cluster Tuning: From the 14×14 transfer matrix, we identify language clusters by inspection, using the criterion that off-diagonal F1 scores are consistently higher within clusters than between cluster members and non-members. The full matrix is reported in Figure 1 for verification. For each cluster, we merge the training sets and fine-tune one model, then evaluate it separately on each language. To ensure the gains are not driven by a single language, we also conduct leave-one-language-out (LOLO) analysis for the three largest clusters, with full results in Appendix 10. This experiment follows an analyze-then-validate pipeline: empirical structure is first identified from the transfer matrix, then tested through a separate training and evaluation procedure. We note that because the clusters are derived from the same transfer matrix, the cluster-tuning result is confirmatory rather than a prediction on independent data; it tests whether the observed structure is strong enough to yield gains under a distinct training setup, not whether clusters generalize to unseen languages. The LOLO analysis (Appendix 10) partially addresses this by showing that each held-out language benefits from a cluster it did not help define. The resulting gains show that observed transfer patterns can inform actionable training strategies, supporting cluster-aware training as a practical approach for multilingual hate-speech detection.

5 Results

Figure 1 presents the 14×14 cross-lingual transfer matrix for hate content detection. Each cell reports the F1 score obtained when a model is fine-tuned on the language shown on the x-axis and evaluated on the test language on the y-axis. The matrix contains 196 train–test combinations, providing a comprehensive view of both in-language and cross-lingual performance. Across languages, the highest F1 scores mostly occur on the diagonal, where the model is fine-tuned and tested on the same language. These in-language results are consistently strong: Turkish (90), Arabic (91), English (96), French (94), Spanish (90), and German (90) achieve F1 scores at or above 90. In a few cases,

	TR	UR	FA	AR	HI	BN	IN	MS	EN	ES	FR	DE	KO	ZH
Solo-Source TL	90 $\pm$ 0.6	85 $\pm$ 0.6	80 $\pm$ 0.8	91 $\pm$ 0.6	85 $\pm$ 0.7	79 $\pm$ 0.7	82 $\pm$ 0.6	78 $\pm$ 0.7	96 $\pm$ 0.5	90 $\pm$ 0.6	94 $\pm$ 0.5	90 $\pm$ 0.6	75 $\pm$ 0.7	75 $\pm$ 0.7
Global-Imb	92 $\pm$ 0.3 [+2]	88 $\pm$ 0.3 [+3]	64 $\pm$ 0.4 [-16]	87 $\pm$ 0.3 [-4]	73 $\pm$ 0.4 [-12]	66 $\pm$ 0.4 [-13]	84 $\pm$ 0.3 [+2]	62 $\pm$ 0.4 [-16]	98 $\pm$ 0.2 [+2]	95 $\pm$ 0.3 [+5]	96 $\pm$ 0.3 [+2]	94 $\pm$ 0.3 [+4]	71 $\pm$ 0.4 [-4]	72 $\pm$ 0.4 [-3]
Global-bal	91 $\pm$ 0.5 [+1]	87 $\pm$ 0.5 [+2]	67 $\pm$ 0.6 [-13]	89 $\pm$ 0.5 [-2]	81 $\pm$ 0.6 [-4]	69 $\pm$ 0.6 [-10]	83 $\pm$ 0.5 [+1]	63 $\pm$ 0.6 [-15]	96 $\pm$ 0.4 [0]	92 $\pm$ 0.5 [+2]	94 $\pm$ 0.4 [0]	93 $\pm$ 0.5 [+3]	73 $\pm$ 0.6 [-2]	76 $\pm$ 0.6 [+1]
Observed Cluster	90 $\pm$ 0.4 [0]	90 $\pm$ 0.4 [+5]	83 $\pm$ 0.5 [+3]	93 $\pm$ 0.4 [+2]	90 $\pm$ 0.5 [+5]	83 $\pm$ 0.5 [+4]	85 $\pm$ 0.4 [+3]	80 $\pm$ 0.5 [+2]	97 $\pm$ 0.3 [+1]	96 $\pm$ 0.4 [+6]	96 $\pm$ 0.4 [+2]	95 $\pm$ 0.4 [+5]	75 $\pm$ 0.5 [0]	78 $\pm$ 0.5 [+2]

Table 2: Performance comparison across different fine-tuning strategies. Results are reported as mean $\pm$ standard deviation over five random seeds. Bracketed numbers indicate improvement or decline relative to Solo-Source TL.

cross-lingual models outperform the in-language baseline; for example, models fine-tuned on Arabic or Indonesian outperform the Persian in-language model on Persian.

Off-diagonal values show substantial variation in the quality of cross-lingual transfer. Based on these off-diagonal patterns, several clusters of languages emerge in which transfer performance is consistently higher within the cluster than across it. These include: (1) Persian–Arabic–Urdu–Indonesian, (2) Urdu–Hindi–Bengali, (3) Indonesian–Malay, (4) Korean–Chinese, and (5) French–Spanish–English–German. As shown in Figure 1, models fine-tuned on Urdu perform strongly on Hindi (85), and English-trained models yield relatively high scores on French (78), Spanish (74), and German (69). There are also cases of strong mutual transfer, such as *French* $\leftrightarrow$ *English* (78) and *Bengali* $\leftrightarrow$ *Urdu* (BN $\rightarrow$ UR: 70; UR $\rightarrow$ BN: 79). Lower scores appear for languages like Bengali, whose models perform less well on European test sets (BN $\rightarrow$ FR: 51; BN $\rightarrow$ DE: 42), and Chinese models show weaker transfer to several languages (ZH $\rightarrow$ ES: 42; ZH $\rightarrow$ AR: 38). Languages such as Korean, Chinese, and Malay generally show lower transfer performance both as train and test languages.

Table 2 summarizes results from the three complementary experiments. In the Global-Imbalanced setup, combining the full training data from all languages leads to performance gains for several high-resource languages, including English, French, and Spanish (average gain: +2.86), as well as improvements for languages such as Urdu. In contrast, most lower-resource languages—such as Bengali, Persian, and Malay—show declines relative to their single-language baselines (average: -5.43).

In the observed cluster tuning experiment, fine-tuning in language-selective clusters from the 14×14 matrix yields consistent performance im-

provements for nearly all cluster members (average gain: +5.62). Gains appear across groups such as Hindi-Bengali-Urdu, Arabic-Persian-Indonesian, and English-French-Spanish-German, with varying magnitudes. In the Global-Balanced setup, equal sampling from all languages reduces the effect of data imbalance. High-resource languages still improve over their single-language baselines, though gains are smaller than in the imbalanced condition (average: +1.57). For lower-resource languages, the balanced dataset mitigates the declines seen under Global-Imbalanced training, producing small gains for most and roughly baseline performance for the remainder.

Robustness Across Model Families. To verify that the observed transfer patterns are not specific to LLaMA-3, we replicate all experimental setups using Gemma-3-4B, and Gemma-4-4B. Although absolute scores differ across models, the same qualitative structure holds for both Gemma-3-4B and Gemma-4-4B. (i) The same five clusters emerge from each model’s transfer matrix (Appendix Tables 10–11): for example, Urdu $\rightarrow$ Hindi, Indonesian $\rightarrow$ Malay, and the Romance/Germanic group remain among the strongest off-diagonal cells in every model. (ii) Observed Cluster tuning again produces the largest and most consistent gains — average cluster-member gains of +2.4 F1 for Gemma-3-4B (Table 8) and +2.1 for Gemma-4-4B (Table 9), with the same low-resource languages (Persian, Bengali, Malay) benefiting most, mirroring LLaMA-3. (iii) Global-Imbalanced training again favors high-resource languages while harming low-resource ones: Persian, Hindi, Bengali, and Malay drop by 9–14 F1 in both Gemma models, and Global-Balanced training again narrows this gap. The consistency of these three effects across two families and three model generations indicates the transfer structure is a property of the data, not of any single model. Full results are reported in

Appendix Tables 8, 9, 10, and 11.

6 Analysis

The 14×14 transfer matrix in Figure 1 reveals several structural patterns, including the five clusters of languages with strong mutual transfer noted earlier. While most languages achieve their highest performance through in-language learning, certain language pairs also show comparable or better results via transfer. Table 6 in the Appendix presents hate samples correctly predicted only via transfer learning. We next evaluate each cluster along three dataset-level properties (size, source platform, and temporal overlap) to assess whether these factors account for the observed patterns.

Dataset Size: Across the clusters, dataset size shows partial but inconsistent alignment. The French-Spanish-English-German group comprises datasets with relatively similar hate-sample counts (3.7k-4.7k), which may support more stable within-group transfer. The Urdu-Hindi-Bengali cluster also falls within a relatively similar medium range (3k-5k hate samples), which may support some internal consistency. In contrast, the Persian-Arabic-Urdu-Indonesian cluster spans markedly different dataset sizes: Persian (~1k hate) is much smaller than Arabic, Urdu, or Indonesian (~4.5k each), yet clusters with them. Similarly, Indonesian-Malay shows large variation, as Malay contains substantially fewer hate examples.

The Korean-Chinese group also exhibits different dataset scales, with Chinese containing roughly twice as many hate examples as Korean. Overall, dataset size aligns with some cluster boundaries (notably the English-French-Spanish-German) but not with clusters involving Persian, Malay, Korean, or Chinese, indicating that size alone cannot explain the structure.

Dataset Resource: Several clusters, including Urdu-Hindi-Bengali, Persian-Arabic-Urdu-Indonesian, and Indonesian-Malay, consist entirely of Twitter-based datasets. However, this does not distinguish them, since most languages in our study, including those outside these clusters, are also Twitter-sourced. Conversely, the Korean-Chinese cluster does not share a single source: Korean is from Twitter, while Chinese comes from general web data. Thus, data source alignment does not systematically explain the cluster structure.

Temporal Overlap: A similar mixed pattern appears when comparing collection periods: the English-French-Spanish-German datasets

share nearly identical time spans (2016-2023), which could support more consistent transfer. The Urdu-Hindi-Bengali datasets likewise cover broadly similar years. However, the Persian-Arabic-Urdu-Indonesian cluster shows substantial variation: Persian (2020-22) does not overlap with Arabic (2016-18) or Indonesian (≤ 2018), yet all appear in the same cluster. Similarly, the Indonesian-Malay pair has almost no temporal overlap (≤ 2018 vs. 2022-23), yet forms a clear transfer cluster. We cannot evaluate the Korean-Chinese cluster on this dimension, as the collection period for the Chinese dataset is unspecified. These comparisons suggest that temporal alignment may support transfer in some clusters but does not consistently predict cluster structure overall. Table 12 in Appendix summarizes these dataset-level factors across the five observed clusters.

6.1 Quantitative Analysis

We examine three measurable dataset properties: dataset size difference, temporal overlap, and linguistic similarity to assess whether they explain the transfer structure in the matrix. For each of the 182 source-target pairs, we compute size difference as the absolute log-ratio of the hate training sizes, temporal overlap as the intersection-over-union of the datasets' collection-year ranges, and linguistic similarity as the cosine similarity of typological feature vectors from URIEL/lang2vec (Bjerva et al., 2019); pairs with missing collection dates are excluded from the temporal measure. We then compute Spearman correlations (Spearman, 1961) between each factor and the off-diagonal F1 scores. As shown in Table 3, linguistic similarity has the strongest association with transfer performance ($\rho = 0.46$), followed by temporal overlap ($\rho = 0.35$), while dataset size difference shows only a weak effect ($\rho = -0.19$). This trend is visible in several typologically close pairs: Urdu→Hindi attains the highest score in the matrix, and Urdu→Bengali, English→German, Spanish→French, and Indonesian→Malay are likewise strong. However, these factors alone cannot fully account for the observed patterns: transfer succeeds from Indonesian to Persian (F1=79) and Arabic to Urdu (F1=70) across different language families, while some closely related or temporally overlapping pairs transfer only moderately (e.g., Turkish to Arabic, F1=58). Full derivations and per-pair values are in Appendix 9 and Table 16.

Factor	Spearman ρ	p -value
Dataset size difference	-0.19	0.011
Temporal overlap	0.35	1.6×10^{-4}
Linguistic similarity	0.46	$< 10^{-9}$

Table 3: Spearman correlations between transfer performance and dataset properties.

6.2 Qualitative Analysis

To better understand when transfer occurs, we conduct a cross-lingual qualitative error analysis by inspecting examples correctly classified only after transfer learning and comparing them with those correctly predicted by the in-language model. The patterns we describe below are interpretive observations drawn from manual inspection of representative examples; they are intended to generate hypotheses about why certain clusters transfer well, not to establish cultural similarity as a measured cause.

In the Arabic-Indonesian-Urdu-Persian cluster, Arabic and Urdu particularly support detection of religious and political hate, especially sociopolitical hate in Middle Eastern contexts. Arabic provides cues aligned with Persian discourse, while Urdu contributes strongly to political hate detection, often more than Arabic. Indonesian contributes more to role-based hate, such as profession-targeted hate, and performs well on longer tweets mixing neutral and hateful content.

Table 6 presents Persian examples correctly predicted only after transfer learning, alongside those captured only by the Persian fine-tuned model. We also observe strong transfer between Indonesian and Malay. Both datasets contain comparable expressions of political frustration, religious commentary, and identity-based accusations. Religious references are frequently used to judge or shame others, including accusations related to religious observance. Similar metaphorical and idiomatic expressions also appear in both languages, including insults referencing moral faults or entities, allowing Indonesian fine-tuned models to generalize well to Malay. Both languages also employ similar proverbs, insults, and metaphorical language, particularly the use of animals such as "dog" and "devil" as slurs. Social issues such as traffic caused by bazaars, public frustration over commodity scarcity, and emotional sharing on social media are likewise reflected in both contexts. However, the reverse transfer direction is weaker,

likely due to the limited amount of hate data in the Malay dataset, which restricts its effectiveness as a source language. A similar pattern appears for Persian, where limited hate training data reduces its effectiveness as a target language. Another cluster where transfer learning outperforms in-language learning includes Urdu, Hindi, and Bengali. Models fine-tuned on Urdu often perform better on Bengali and Hindi than the respective in-language models. Transfer from Urdu is particularly effective for highly polarized, overt hate expressions involving explicit slurs, religious antagonism, or politicized attacks.

In contrast, in-language models more accurately classify tweets containing idiomatic expressions or contextually grounded criticism. For instance, socially embedded sarcasm in Bengali was frequently missed by the Urdu-fine-tuned model but correctly identified by the Bengali-fine-tuned counterpart. This difference appears related to dataset composition: the Urdu dataset contains a high density of explicit hate markers—religious slurs, profane language, and direct references to identifiable targets—which exhibit lexical consistency and limited idiomatic variation, enabling reliable pattern learning. By comparison, Bengali and Hindi datasets contain more ambiguous, sarcastic, or culturally embedded forms of expression; specifically, the Hindi dataset includes implicit hate, which is harder to generalize across languages. We also observe that transfer strength is not driven by topical similarity alone. Although English and French contain religious or political hate, Arabic or Urdu transfer yields larger gains for detecting such hate in Persian. A similar pattern appears for Arabic-to-English or Arabic-to-French transfer.

To broaden the analysis, we also inspected examples from language pairs outside the same empirical cluster, examining transfer when alignment is weaker. Our results show that even when languages do not belong to the same cluster (e.g., Spanish–Chinese, Persian–French), transfer learning can still detect certain types of hate-speech, though less consistently. In particular, explicit, emotionally extreme, or lexically offensive expressions tend to transfer well across languages. For example, in our experiments, the Spanish→Chinese transfer achieved moderate performance in detecting overtly hateful Chinese instances such as *You’re a brainless feminist bitch* or *(So many feminist bitches on Weibo!)*. However, the Spanish-trained model struggled with more context-dependent ex-

pressions, such as: *Emperor Tang Minghuang’s lifetime reputation was ruined at the hands of the Yang family, A historical reference used to delegitimize women, or I also saw someone called ‘Pepsi256’ who claims to be a feminist. Lately, she’s been supporting the Sichuan Fine Arts Institute’s pro-invasion-of-Taiwan military otakus and happily interacting with the nationalists (a.k.a. little pinks).*

In contrast, training the model on Korean resulted in significantly better performance not only on overt hate expressions but also on more subtle, context-dependent instances like those discussed above. In other cases, experiments with French and Persian show that a Persian fine-tuned model can still correctly identify highly explicit hate content in French. For example: “Vous pouvez crever sale race” (You can die, filthy race). This phrase combines an aggressive death wish with a racial slur, making it emotionally extreme and lexically explicit, which enables the Persian fine-tuned model to detect it despite belonging to a different cluster. These results suggest that even across dissimilar clusters, transfer learning succeeds when hate content is lexically explicit or emotionally extreme, but consistently struggles with expressions relying on background social or historical context.

Taken together, these qualitative observations point to a consistent theme: transfer is strongest for explicit, lexically overt hate and weakest for implicit, culturally or historically grounded expression. We verified these patterns by inspecting the same examples for fine-tuned Gemma-3 and Gemma-4; cases misclassified by LLaMA-3 were typically also misclassified by both Gemma variants, and correctly classified cases showed similar patterns across models, indicating the patterns are not specific to a single model family.

7 Discussion

For RQ1, whether cross-lingual training can match in-language hate-speech detection—our results show that cross-lingual models can achieve comparable or superior performance to in-language fine-tuning in several cases. Across the 14 languages studied, this effect is most evident within the observed clusters, where languages such as Persian, Indonesian, Urdu, Malay, Hindi, and Bengali benefit substantially from training on related languages. For RQ2, whether multilingual training improves performance beyond single-source transfer—only specific setups yield consistent gains. Globally imbalanced training primar-

ily benefits high-resource languages while weakening the language-specific signals needed for low-resource ones. In contrast, cluster-based training provides the most reliable improvements across languages.

RQ3 examines potential reasons behind successful or failed transfer. Our quantitative analysis (Table 3) identifies linguistic similarity as the strongest measurable predictor of cross-lingual transfer ($\rho = 0.46$), followed by temporal overlap ($\rho = 0.35$), while dataset size difference shows only a weak association. These results suggest that transfer patterns in the 14×14 matrix are partly explained by typological and temporal proximity. However, linguistic similarity does not fully account for all patterns: some high-performing pairs, such as Arabic→Urdu, Indonesian→Persian, and Turkish→English, involve typologically different languages, while some closely related pairs show only moderate transfer. The qualitative findings in Section 6.2 suggest that some pairs left unexplained by these measurable factors coincide with shared discourse and communicative patterns within clusters. Because this observation comes from manual inspection, we treat it as a complementary hypothesis alongside linguistic similarity, not as an independently measured factor.

Together, these observations suggest that some of the empirically observed clusters may reflect overlapping cultural cues or communicative patterns (Zhou et al., 2023), rather than transfer arising from any single language-pair property. We caution, however, that this cultural interpretation is complementary to—not a replacement for—the role of linguistic similarity, and that confounds such as topic distribution, annotation scheme, and annotator demographics remain only partly controlled (see Limitations).

Our observations align with prior work on cultural context in hate-speech detection. Lee et al. (2023) show that models trained in one cultural context often perform poorly across cultures because they miss culturally specific targets and discourse norms. Similarly, Lee et al. (2024) show that annotations can vary across countries even within the same language due to cultural and social differences. Our findings extend this work by showing that such differences also affect cross-lingual transfer: shared cultural context may strengthen transfer, while culturally specific expressions can weaken it.

8 Conclusion

This study analyzes cross-lingual transfer learning for hate-speech detection across fourteen languages. While in-language fine-tuning remains the most robust baseline, cross-lingual models can match or surpass it by leveraging clusters of languages with similar hate-discourse patterns. Multilingual training is most effective when restricted to these empirical clusters; globally imbalanced integration provides limited support for low-resource languages. Crucially, successful transfer is not fully explained by dataset size, resource availability, or temporal overlap. Among the factors we quantify, linguistic similarity is the strongest predictor; our qualitative analysis additionally points to shared cultural and communicative patterns as a possible but unmeasured contributor within several clusters. These patterns are consistent across examined models (LLaMA and Gemma), suggesting that the observed transfer structure reflects properties of the data and task. Explicit or highly polarized hate transfers well across distant languages, but culturally embedded and context-dependent forms remain challenging. Overall, effective cross-lingual hate-speech detection benefits from both typological proximity and attention to the social dynamics of online hate. Direct measurement of cross-lingual social alignment, ideally with datasets that hold linguistic similarity constant, is a promising direction for future work.

Limitations

This study has several limitations. First, the datasets used across the fourteen languages vary in source platforms, annotation practices, and topical distributions. Although we standardize dataset size where possible and map labels to a binary hate/non-hate scheme, differences in domain, explicitness, and annotation criteria may still influence transfer behavior. In particular, implicit hate is annotated only in a subset of languages, which limits direct evaluation of how implicit expressions transfer. Second, while our quantitative analysis considers dataset size, collection period, and linguistic similarity, several factors that may shape transfer performance remain only partly controlled. These include topic distribution, platform composition, the scope and granularity of annotation schemes, and the demographic characteristics of annotators, which prior work has shown to influence hate-speech judgments (Sap et al., 2022b). Because these factors may co-vary with the regional or cul-

tural groupings we identify, the empirical clusters we report should be interpreted as observed transfer patterns.

Third, the thematic observations from our qualitative error analysis rely on manual inspection of representative examples and may not capture the full variability of hate discourse across languages; a more extensive evaluation with multiple annotators would offer a stronger foundation for interpreting transferable expressions. Fourth, while our quantitative analysis establishes linguistic similarity as the strongest empirically supported predictor of cross-lingual transfer, the shared cultural and discourse-level patterns identified through qualitative analysis serve as a complementary factor. Direct measurement of cross-lingual sociopolitical alignment was beyond the scope of this work and represents a promising direction for future research, ideally supported by parallel datasets that hold linguistic similarity constant. Fifth, our cross-model validation spans two families (LLaMA and Gemma) but both models (Gemma-3-4B and Gemma-4-4B) belong to the Gemma line and are smaller than our primary LLaMA-3-8B. The consistency of the transfer structure across LLaMA and Gemma, across two Gemma generations, and across this scale gap is evidence that the patterns are not model-specific, but we cannot rule out that a more distant architecture or a substantially larger model would behave differently. Replication across additional model families and scales is an important direction for future work.

Ethics Statement

This study analyzes publicly available hate datasets and does not involve collecting new user data. All datasets were obtained from prior peer-reviewed work or shared tasks that follow established ethical guidelines. Because hate datasets may contain harmful or offensive language, examples shown in this paper are minimized and presented only when necessary for scientific transparency. No personally identifiable information is included in our datasets or model outputs. All experiments were conducted using anonymized text. Models trained in this work are not intended for deployment without further evaluation, fairness auditing, and context-specific calibration.

References

Amirhossein Abaskohi, Sara Baruni, Mostafa Masoudi, Nesa Abbasi, Mohammad Hadi Babalou, Ali Edalat,

- Sepehr Kamahi, Samin Mahdizadeh Sani, Nikoo Naghavian, Danial Namazifard, Pouya Sadeghi, and Yadollah Yaghoobzadeh. 2024. [Benchmarking large language models for Persian: A preliminary study focusing on ChatGPT](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2189–2203, Torino, Italia. ELRA and ICCL.
- Ife Adebara and Muhammad Abdul-Mageed. 2021. [Improving similar language translation with transfer learning](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 273–278, Online. Association for Computational Linguistics.
- Imatitikua D Aiyanyo, Hamman Samuel, and Heuiseok Lim. 2020. A systematic review of defensive and offensive cybersecurity with machine learning. *Applied Sciences*, 10(17):5811.
- Nuha Albadi, Maram Kurdi, and Shivakant Mishra. 2018. [Are they our brothers? analysis and detection of religious hate speech in the arabic twitter-sphere](#). In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 69–76.
- Doris Chinedu Asogwa, Chiamaka Ijeoma Chukwuneke, CC Ngene, and GN Anigbogu. 2022. Hate speech classification using SVM and Naive Bayes. *arXiv preprint arXiv:2204.07057*.
- Girma Bade, Olga Kolesnikova, Grigori Sidorov, and José Oropeza. 2024. [Social media hate and offensive speech detection using machine learning method](#). In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, pages 240–244, St. Julian’s, Malta. Association for Computational Linguistics.
- Samuel Bell, Eduardo Sánchez, David Dale, Pontus Stenetorp, Mikel Artetxe, and Marta R Costa-jussà. 2025. Translate, then detect: Leveraging machine translation for cross-lingual toxicity classification. In *Proceedings of the Tenth Conference on Machine Translation*, pages 253–268.
- Imene Bensalem, Meryem Mout, and Paolo Rosso. 2023. Offensive language detection in Arabizi. In *Proceedings of ArabicNLP 2023*, pages 423–434.
- Mohit Bhardwaj, Md. Shad Akhtar, Asif Ekbal, Amitava Das, and Tanmoy Chakraborty. 2020. [Hostility detection dataset in hindi](#). *CoRR*, abs/2011.03588.
- Johannes Bjerva, Robert Östling, Maria Han Veiga, Jörg Tiedemann, and Isabelle Augenstein. 2019. What do language representations really represent? *Computational Linguistics*, 45(2):381–389.
- Zahra Bokaei, Walid Magdy, and Bonnie Webber. 2025. [Culture matters in toxic language detection in Persian](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9290–9304, Vienna, Austria. Association for Computational Linguistics.
- Zahra Bokaei, Walid Magdy, and Bonnie Webber. 2026a. [Farsi natural language processing: A survey](#). *Information Processing Management*, 63(8):104912.
- Zahra Bokaei, Bonnie Webber, and Walid Magdy. 2026b. Benchmarking offensive language detection in persian and pashto. In *The Proceedings of the First Workshop on NLP and LLMs for the Iranian Language Family*, pages 13–23.
- Minh Duc Bui, Katharina Von Der Wense, and Anne Lauscher. 2025. Multi³hate: Multimodal, multilingual, and multicultural hate speech detection with vision–language models. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 9714–9731.
- Amparo Elizabeth Cano Basave, Yulan He, Kang Liu, and Jun Zhao. 2013. [A weakly supervised Bayesian model for violence detection in social media](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 109–117, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Rui Cao, Roy Ka-Wei Lee, and Tuan-Anh Hoang. 2020. DeepHate: Hate speech detection via multi-faceted text representations. In *Proceedings of the 12th ACM Conference on Web Science*, pages 11–20.
- Jan Christian Blaise Cruz, Julianne Agatha Tan, and Charibeth Cheng. 2020. [Localization of fake news detection via multitask transfer learning](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2596–2604, Marseille, France. European Language Resources Association.
- 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.
- Zahra Delbari, Nafise Sadat Moosavi, and Mohammad Taher Pilehvar. 2024. Spanning the spectrum of hatred detection: a Persian multi-label hate speech dataset with annotator rationales. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17889–17897.
- Jiawen Deng, Jingyan Zhou, Hao Sun, Chujie Zheng, Fei Mi, Helen Meng, and Minlie Huang. 2022. [COLD: A benchmark for Chinese offensive language detection](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11580–11599, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Fatima-zahra El-Alami, Said Ouatic El Alaoui, and Nouredine En Nahnah. 2022. A multilingual offensive language detection method based on transfer learning from transformer fine-tuning model. *Journal of King Saud University-Computer and Information Sciences*, 34(8):6048–6056.

- Emna Fsih, Sameh Kchaou, Rahma Boujelbane, and Lamia Hadrach-Belguith. 2022. [Benchmarking transfer learning approaches for sentiment analysis of Arabic dialect](#). In *Proceedings of the Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 431–435, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Faeze Ghorbanpour, Daryna Dementieva, and Alexander Fraser. 2025. [Data-efficient hate speech detection via cross-lingual nearest neighbor retrieval with limited labeled data](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 29662–29680, Suzhou, China. Association for Computational Linguistics.
- Prenti Golazizian, Behnam Sabeti, Seyed Arad Ashrafi Asli, Zahra Majdabadi, Omid Momenzadeh, and Reza Fahmi. 2020. [Irony detection in Persian language: A transfer learning approach using emoji prediction](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2839–2845, Marseille, France. European Language Resources Association.
- Xiaoyu Guo and Susan Gauch. 2024. [Using sarcasm to improve cyberbullying detection](#). In *Proceedings of the Fourth Workshop on Threat, Aggression & Cyberbullying @ LREC-COLING-2024*, pages 52–59, Torino, Italia. ELRA and ICCL.
- Vikram Gupta, Sumegh Roychowdhury, Mithun Das, Somnath Banerjee, Punyajoy Saha, Binny Mathew, Animesh Mukherjee, et al. 2022. Multilingual abusive comment detection at scale for Indic languages. *Advances in Neural Information Processing Systems*, 35:26176–26191.
- Nhat Hoang, Xuan Long Do, Duc Anh Do, Duc Anh Vu, and Anh Tuan Luu. 2024. [ToXCL: A unified framework for toxic speech detection and explanation](#). 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 6460–6472, Mexico City, Mexico. Association for Computational Linguistics.
- Fatemah Husain, Hana Al-Ostad, and Halima Omar. 2022. [A weak supervised transfer learning approach for sentiment analysis to the Kuwaiti dialect](#). In *Proceedings of the Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 161–173, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Muhammad Okky Ibrohim and Indra Budi. 2019. [Multi-label hate speech and abusive language detection in Indonesian Twitter](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 46–57, Florence, Italy. Association for Computational Linguistics.
- 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.
- Aiqi Jiang and Arkaitz Zubiaga. 2024. Cross-lingual offensive language detection: A systematic review of datasets, transfer approaches and challenges. *arXiv preprint arXiv:2401.09244*.
- Raghav Kapoor, Yaman Kumar, Kshitij Rajput, Rajiv Ratn Shah, Ponnurangam Kumaraguru, and Roger Zimmermann. 2019. Mind your language: Abuse and offense detection for code-switched languages. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 9951–9952.
- Md Tawkat Islam Khondaker, Abdul Waheed, El Moatez Billah Nagoudi, and Muhammad Abdul-Mageed. 2023. [GPTAraEval: A comprehensive evaluation of ChatGPT on Arabic NLP](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 220–247, Singapore. Association for Computational Linguistics.
- Yunsu Kim, Petre Petrov, Pavel Petrushkov, Shahram Khadivi, and Hermann Ney. 2019. [Pivot-based transfer learning for neural machine translation between non-English languages](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 866–876, Hong Kong, China. Association for Computational Linguistics.
- Katerina Korre, Arianna Muti, and Alberto Barrón-Cedeño. 2024. [The challenges of creating a parallel multilingual hate speech corpus: An exploration](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15842–15853, Torino, Italia. ELRA and ICCL.
- Ankit Kumar, Richa Sharma, and Punam Bedi. 2024. Towards optimal NLP solutions: Analyzing GPT and LLaMA-2 models across model scale, dataset size, and task diversity. *Engineering, Technology & Applied Science Research*, 14(3):14219–14224.
- Jean Lee, Taejun Lim, Heejun Lee, Bogeun Jo, Yangsok Kim, Heegeun Yoon, and Soyeon Caren Han. 2022. [K-MHaS: A multi-label hate speech detection dataset in Korean online news comment](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3530–3538, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- 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.
- Nayeon Lee, Chani Jung, and Alice Oh. 2023. Hate speech classifiers are culturally insensitive. In *Proceedings of the first workshop on cross-cultural considerations in NLP (C3NLP)*, pages 35–46.
- Krishanu Maity, Shaubhik Bhattacharya, Sriparna Saha, and Manjeevan Seera. 2023. A deep learning framework for the detection of malay hate speech. *IEEE Access*, 11:79542–79552.
- Jitendra Singh Malik, Hezhe Qiao, Guansong Pang, and Anton van den Hengel. 2024. Deep learning for hate speech detection: a comparative study. *International Journal of Data Science and Analytics*, pages 1–16.
- Thomas Mandl, Sandip Modha, Prasenjit Majumder, Daksh Patel, Mohana Dave, Chintak Mandlia, and Aditya Patel. 2019. [Overview of the hasoc track at fire 2019: Hate speech and offensive content identification in Indo-European languages](#). In *Proceedings of the 11th Annual Meeting of the Forum for Information Retrieval Evaluation, FIRE '19*, page 14–17, New York, NY, USA. Association for Computing Machinery.
- Francois Meyer and Jan Buys. 2024. [A systematic analysis of subwords and cross-lingual transfer in multilingual translation](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2194–2200, Mexico City, Mexico. Association for Computational Linguistics.
- Jihyung Moon, Won Ik Cho, and Junbum Lee. 2020. [BEEP! Korean corpus of online news comments for toxic speech detection](#). In *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media*, pages 25–31, Online. Association for Computational Linguistics.
- Wajid Hassan Moosa and Najiba. 2022. Multilingual hate speech dataset. <https://www.kaggle.com/datasets/wajidhassanmoosa/multilingual-hatespeech-dataset>. Open Data Commons Attribution License (ODC-By) v1.0.
- Hamdy Mubarak, Hend Al-Khalifa, and Abdulmohsen Al-Thubaity. 2022. [Overview of OSACT5 shared task on Arabic offensive language and hate speech detection](#). In *Proceedings of the 5th Workshop on Open-Source Arabic Corpora and Processing Tools with Shared Tasks on Qur'an QA and Fine-Grained Hate Speech Detection*, pages 162–166, Marseille, France. European Language Resources Association.
- Hamdy Mubarak, Ammar Rashed, Kareem Darwish, Younes Samih, and Ahmed Abdelali. 2021. [Arabic offensive language on Twitter: Analysis and experiments](#). In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 126–135, Kyiv, Ukraine (Virtual). Association for Computational Linguistics.
- Nanlir Sallau Mullah and Wan Mohd Nazmee Wan Zainon. 2021. Advances in machine learning algorithms for hate speech detection in social media: a review. *IEEE Access*, 9:88364–88376.
- Debora Nozza. 2021. Exposing the limits of zero-shot cross-lingual hate speech detection. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 907–914.
- 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, Florence, Italy. Association for Computational Linguistics.
- Siddhesh Pawar, Junyeong Park, Jiho Jin, Faiz Ghifari Haznitrana, Junho Myung, Alice Oh, Isabelle Augenstein, et al. 2024. Survey of cultural awareness in language models: Text and beyond. *arXiv preprint arXiv:2411.00860*.
- Zesis Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. [Offensive language identification in Greek](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5113–5119, Marseille, France. European Language Resources Association.
- Tharindu Ranasinghe and Marcos Zampieri. 2021. Multilingual offensive language identification for low-resource languages. *Transactions on Asian and Low-Resource Language Information Processing*, 21(1):1–13.
- Faisal Hossain Raquib, Akm Moshir Rahman Mazumder, Md Tahmid Hasan Fuad, Md Farhan Ishmam, and Md Fahim. 2025. Banhate: An up-to-date and fine-grained bangla hate speech dataset. In *Proceedings of the Second Workshop on Bangla Language Processing (BLP-2025)*, pages 237–248.
- Bhawani Singh Rathore and Sandeep Chaurasia. 2025. Fine tuning large language models for hate speech detection in hinglish and code mixed custom dataset through a socially responsible approach for safer digital platforms. *Discover Sustainability*, 6(1):1409.
- Julian Risch, Robin Ruff, and Ralf Krestel. 2020. [Offensive language detection explained](#). In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 137–143, Marseille, France. European Language Resources Association (ELRA).
- Hammad Rizwan, Muhammad Haroon Shakeel, and Asim Karim. 2020. [Hate-speech and offensive language detection in Roman Urdu](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2512–2522, Online. Association for Computational Linguistics.

- Nauros Romim, Mosahed Ahmed, Hriteshwar Talukder, and Md Saiful Islam. 2020. [Hate speech detection in the bengali language: A dataset and its baseline evaluation](#). Preprint, arXiv:2012.09686.
- Pradeep Kumar Roy, Asis Kumar Tripathy, Tapan Kumar Das, and Xiao-Zhi Gao. 2020. A framework for hate speech detection using deep convolutional neural network. *IEEE Access*, 8:204951–204962.
- Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A. Smith. 2022a. [Annotators with attitudes: How annotator beliefs and identities bias toxic language detection](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5884–5906, Seattle, United States. Association for Computational Linguistics.
- Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A. Smith. 2022b. Annotators with attitudes: How annotator beliefs and identities bias toxic language detection. In *Proceedings of the 2022 conference of the north american chapter of the association for computational linguistics: Human language technologies*, pages 5884–5906.
- Gudbjartur Ingi Sigurbergsson and Leon Derczynski. 2020. [Offensive language and hate speech detection for Danish](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 3498–3508, Marseille, France. European Language Resources Association.
- Charles Spearman. 1961. The proof and measurement of association between two things.
- Manuel Tonneau. 2024. [Spanish hate speech superset](#). Accessed: 2025-05-17.
- Manuel Tonneau, Diyi Liu, Samuel Fraiberger, Ralph Schroeder, Scott Hale, and Paul Röttger. 2024. [From languages to geographies: Towards evaluating cultural bias in hate speech datasets](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Cagri Toraman, Furkan Şahinuç, and Eyup Yilmaz. 2022. [Large-scale hate speech detection with cross-domain transfer](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2215–2225, Marseille, France. European Language Resources Association.
- Cynthia Van Hee, Els Lefever, Ben Verhoeven, Julie Mennes, Bart Desmet, Guy De Pauw, Walter Daelemans, and Veronique Hoste. 2015. [Detection and fine-grained classification of cyberbullying events](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pages 672–680, Hissar, Bulgaria. INCOMA Ltd. Shoumen, BULGARIA.
- Kunze Wang, Dong Lu, Caren Han, Siqu Long, and Josiah Poon. 2020. [Detect all abuse! toward universal abusive language detection models](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6366–6376, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Michael Wiegand, Melanie Siegel, and Josef Ruppenhofer. 2018. Overview of the GermEval 2018 shared task on the identification of offensive language.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. [SemEval-2019 task 6: Identifying and categorizing offensive language in social media \(OffensEval\)](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 75–86, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Ziqi Zhang, David Robinson, and Jonathan Tepper. 2018. Detecting hate speech on twitter using a convolution-gru based deep neural network. In *The Semantic Web: 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, Proceedings 15*, pages 745–760. Springer.
- Li Zhou, Laura Cabello, Yong Cao, and Daniel Herscovich. 2023. [Cross-cultural transfer learning for Chinese offensive language detection](#). In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 8–15, Dubrovnik, Croatia. Association for Computational Linguistics.
- Yanling Zhou, Yanyan Yang, Han Liu, Xiufeng Liu, and Nick Savage. 2020. [Deep learning based fusion approach for hate speech detection](#). *IEEE Access*, 8:128923–128929.
- Steven Zimmerman, Udo Kruschwitz, and Chris Fox. 2018. Improving hate speech detection with deep learning ensembles. In *Proceedings of the eleventh international conference on language resources and evaluation (LREC 2018)*.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. [Transfer learning for low-resource neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575, Austin, Texas. Association for Computational Linguistics.

Appendix

Lang	Reference	Labels	Hate definition as stated in reference
BN	(Romim et al., 2020)	Hate, Not hate	A sentence that dehumanizes one or multiple persons or a community (e.g., comparing them to insects, objects, criminals, or targeting race, gender, disability).
UR	(Rizwan et al., 2020)	Hate-Offensive, Normal, Abusive, Offensive, Sexism, Religious Hate, Profane	Language that expresses hatred, abuse, or offensive content such as profanity, vulgar or insulting language targeting individuals or groups; includes gender-based hatred and religion-based hatred.
HI	(Bhardwaj et al., 2020)	Hostile, Non-hostile, Fake News, Hate, Offensive, Defamation, Normal	A post targeting a specific group of people based on ethnicity, religious beliefs, geographical belonging, race, etc., with malicious intentions of spreading hate or encouraging violence.
FA	(Delbari et al., 2024)	Normal, Violence, Hate, Vulgar	Communication targeting individuals or groups with the intention of attacking, dehumanizing, inciting violence, promoting discrimination, or fostering hostility based on characteristics such as race, ethnicity, religion, gender, sexual orientation, political beliefs, social class, disability, or appearance.
IN	(Ibrohim and Budi, 2019)	Hate, Not hate	Communication that disparages a person or a group based on characteristics such as race, ethnicity, gender, sexual orientation, nationality, religion, or other attributes.
MS	(Maity et al., 2023)	Hate, Not hate	Communication intended to attack the dignity of a group on the basis of characteristics such as race, gender, ethnicity, sexual orientation, nationality, religion, or other characteristics.
AR	(Mubarak et al., 2022)	Normal, Abusive, Hate	Any communication that disparages a person or a group based on characteristics such as race, color, ethnicity, gender, sexual orientation, nationality, religion, or other characteristic.
TR	(Toraman et al., 2022)	Hate, Offensive, Normal	Tweets that target, incite violence against, threaten, or call for physical damage toward an individual or a group because of some identifying trait or characteristic.
KO	(Moon et al., 2020)	Hate, Offensive, None	Comments that explicitly express hatred against an individual or group based on attributes such as sex, gender, sexual orientation, gender identity, age, appearance, social status, religion, disability, ethnicity, or national origin, or that severely insult or attack them.
ES	(Tonneau, 2024)	Hate, Not hate	Speech that denigrates a person or multiple persons based on their membership to a group, usually defined by race, ethnicity, sexual orientation, gender identity, disability, religion, political affiliation, or views.
DE	(Tonneau et al., 2024)	Hate, Not hate, Offensive, Profanity	Describing dehumanization or deficiencies to groups or individuals because of being members of a group (e.g., race, political opinion, sexual orientation, gender, social status, health condition).
FR	(Tonneau et al., 2024)	Abusive, Hate, Offensive, Disrespectful, Fearful, Normal	Speech promotes the dehumanization of people who are typically from minority groups and may incite hate crimes.
EN	(Tonneau et al., 2024)	Abusive, Hate, Offensive, Disrespectful, Fearful, Normal	Speech promotes the dehumanization of people who are typically from minority groups and may incite hate crimes.

Table 4: Comparison of hate-speech definitions and annotation schemes across languages.

Method	Model	TR	UR	FA	AR	HI	BN	IN	MS	EN	ES	FR	DE	KO	ZH
Zero-shot	GPT-4o	71	45	55	64	67	53	65	56	82	75	77	75	60	67
	LLaMA-3-8B-Instruct	70	52	45	72	70	58	64	51	82	79	76	74	49	66
	Gemma-3-4B-Instruct	46	38	58	59	56	54	40	51	85	67	69	70	43	52
	Gemma-4-4B-Instruct	58	50	64	67	64	60	55	58	87	76	78	78	55	64
Fine-tuning	LLaMA-3-8B	90	85	80	91	85	79	82	78	96	90	94	90	75	75
	Gemma-3-4B	88	85	75	90	75	70	75	71	94	88	90	94	76	82
	Gemma-4-4B	92	88	82	93	84	78	84	79	97	92	95	96	81	86

Table 5: F1 scores of different models for multilingual hate-speech detection.

Lang	TL From	Tweet	Explanation
FA	IN	رندی به محضر فقیهی رسیدو حرکات رقص را جداجدا انجام می داد و می پرسید آیا حرام است؟ فقیه میگفت نه. پس رند شروع به رقصیدن کرد. فقیه گفت تجزیه اش خوب بود ولی مرده شور ترکیش رو بردن حالا حکایت این عدالت خواراست بعضیشون عیبی بچه های خوبین ولی مرده شور ترکیشون رو برده من برم به کارای خوند برسم خدافظ #عدالتخواران#انتخابات مجلس A trickster went to a cleric and performed dance moves separately, asking if they were forbidden. The cleric said no. Then the trickster started dancing, and the cleric said, 'Breaking it down was fine, but damn the combination!' This is exactly the case with these so-called justice-seekers—some of them are actually good kids, but damn their combination! Anyway, I'll get back to my own business. Goodbye. #JusticeSeekers #ParliamentElection	Long tweet containing a mixture of neutral and toxic sentiments, with shared religious content.
FA	AR	وقتی به جوان بدبخت میخاد استخدام بشه ریشه و بیخ مسائل خانوادگیش رو درمیارن اما زیباکلام با دختر بازی و شراب خواریش هنوز تو دانشگاه اسلامی!!! تدریس میکنه When a poor young man wants to get hired, they dig up every root of his family issues, but Zibakalam, with his womanizing and wine-drinking, is still teaching at an Islamic university!	Religious hypocrisy framed with moral outrage, a sentiment common to both Persian and Arabic contexts.
MS	IN	Oh dan kereta mmg suka menghimpit ke motor walaupun jelas2 ade moto kt situ smpi terpaksa mengelak n kereta tu buat bodoh je. Ramai ye perangai mcm ni. Oh and cars really like to squeeze in on motorcycles even when it's obvious there's a bike there, to the point the rider has to avoid them, and the car just acts clueless. There are many people with this kind of attitude.	Everyday frustration or complaint expressed in the form of Malay-language hate.
BN	UR	বেয়াদবের বাচ্চা বেয়াদব হ্যাকার গ্রুপ বর্ণনা হ্যাকার গ্রুপ বার করে দিব লুচামি কি করস শরীরটা দেখছিস নিজের ছবিটা দেখ Impudent child's impudent hacker group description hacker group will be kicked out, what lewdness are you doing, looking at the body, look at your own picture.	Accusing a group of hackers of lewd behaviour and threatening expulsion.
HI	UR	हम तेरे लिये कुछ नहीं कर सकते थे दोस्त...!!!! हमें अभी और मंदिर और मस्जिद बनानी है...!! We can't do anything for you my friend... We still have to build more temples and mosques...	Politicians prioritize religious projects over people's real needs.
ZH	KO	唐明皇一世英名就毁在杨家人手上了 Emperor Tang Minghuang's lifetime reputation was ruined at the hands of the Yang family - A historical reference used to delegitimize women	subtle, context-dependent instances.
ZH	KO	我还看到一个百事256的说自己是女权的，最近在力挺川美侵台军宅，和粉红互动得可欢了呢。 I also saw someone called 'Pepsi256' who claims to be a feminist. Lately, she's been enthusiastically supporting the Sichuan Fine Arts Institute's pro-invasion-of-Taiwan military otakus and happily interacting with the nationalists (a.k.a. little pinks	subtle, context-dependent instances

Table 6: Hate samples correctly predicted via transfer learning across languages.

tweet	Ar	Ind	Type
تف تو مملکتی که دیه الفاز رکابی از مهدی ترابی کمتره 🤔🤔...!			
Shame on a country where Elnaz Rekabi's blood money is worth less than Mehdi Torabi's...! 🤔🤔	1	0	Politics
رندی به محضر فقیهی رسیدو حرکات رقص را جداجدا انجام می داد و می پرسید آیا حرام است؟ فقیه میگفت نه. پس رند شروع به رقصیدن کرد. فقیه گفت تجزیه اش خوب بود ولی مرده شور ترکیش رو بردن حالا حکایت این عدالت خواراست بعضیشون عیبی بچه های خوبین ولی مرده شور ترکیشون رو برده من برم به کارای خوند برسم خدافظ #عدالتخواران#انتخابات مجلس	0	1	Mix of Neutral and Hate (bold)
A trickster went to a cleric and performed dance moves separately, asking if they were forbidden. The cleric said no. Then the trickster started dancing, and the cleric said, 'Breaking it down was fine, but damn the combination!' This is exactly the case with these so-called justice-seekers—some of them are actually good kids, but damn their combination! Anyway, I'll get back to my own business. Goodbye. #JusticeSeekers #ParliamentElections			
رقم دعاغمو عمل کتم دکتره په نگاهی کرد گفت داداش شما صبر کن یکم دیگه علم پیشرفت کنه کلا سرتو عوض کن 🤔 حالا شده حکایت ما مملکتی که رقابت خوشگترین مسئولش بین جهانگیری و احمد خاتمی و احمدی نژاد باید رید توش البته از نظر عملکرد شان هم باید رید توش #مهسا_امینی #تورو_بیولوژیکی	0	1	Mix of Neutral and Hate (bold)
I went to get my nose done, and the doctor took one look at me and said, 'Bro, just wait a little longer until science advances enough to replace your whole head. 🤔' That's exactly our situation—when the competition for the most handsome official in the country is between Jahangiri, Ahmad Khatami, and Ahmadinejad, you know it's doomed. And in terms of their performance and dignity, well, it's even worse. #Mahsa_Amini #Biological_Assassination			
از ماست که بر ماست تا به این دین و باورهای بیابان گرد ملخ خوار باور داریم همین اش و همین کاسه	0	0	idiomatic
It is up to us to believe in this religion and beliefs of the desert, the locust-eating locusts, the same soup and the same bowl.			

Table 7: Case study of Persian hate-speech instances correctly predicted through transfer learning from cluster languages, or misclassified in all experimental settings.

Strategy	TR	UR	FA	AR	HI	BN	IN	MS	EN	ES	FR	DE	KO	ZH
Solo-Source TL	88	85	75	90	75	70	75	71	94	88	90	94	76	82
Global-Imb	89 [+1]	87 [+2]	62 [-13]	86 [-4]	65 [-10]	60 [-10]	77 [+2]	57 [-14]	95 [+1]	92 [+4]	92 [+2]	96 [+2]	73 [-3]	79 [-3]
Global-Bal	88 [0]	86 [+1]	64 [-11]	88 [-2]	71 [-4]	62 [-8]	76 [+1]	58 [-13]	94 [0]	90 [+2]	90 [0]	95 [+1]	74 [-2]	81 [-1]
Observed Cluster	88 [0]	89 [+4]	78 [+3]	92 [+2]	79 [+4]	73 [+3]	78 [+3]	73 [+2]	95 [+1]	93 [+5]	92 [+2]	96 [+2]	76 [0]	84 [+2]

Table 8: Performance comparison across fine-tuning strategies using Gemma-3. Bracketed numbers indicate change relative to Solo-Source TL.

Strategy	TR	UR	FA	AR	HI	BN	IN	MS	EN	ES	FR	DE	KO	ZH
Solo-Source TL	92	88	82	93	84	78	84	79	97	92	95	96	81	86
Global-Imb	94 ^[+2]	90 ^[+2]	70 ^[-12]	90 ^[-3]	75 ^[-9]	69 ^[-9]	86 ^[+2]	67 ^[-12]	98 ^[+1]	95 ^[+3]	96 ^[+1]	97 ^[+1]	79 ^[-2]	84 ^[-2]
Global-Bal	93 ^[+1]	89 ^[+1]	73 ^[-9]	92 ^[-1]	81 ^[-3]	71 ^[-7]	85 ^[+1]	68 ^[-11]	97 ^[0]	94 ^[+2]	95 ^[0]	97 ^[+1]	80 ^[-1]	86 ^[0]
Observed Cluster	92 ^[0]	92 ^[+4]	85 ^[+3]	95 ^[+2]	87 ^[+3]	81 ^[+3]	86 ^[+2]	81 ^[+2]	98 ^[+1]	96 ^[+4]	96 ^[+1]	98 ^[+2]	81 ^[0]	88 ^[+2]

Table 9: Performance comparison across fine-tuning strategies using Gemma-4. Bracketed numbers indicate change relative to Solo-Source TL.

Test ↓ / Train →	TR	UR	FA	AR	HI	BN	IN	MS	EN	ES	FR	DE	KO	ZH
TR	88	64	65	56	56	47	39	36	64	55	61	68	49	53
UR	63	85	66	65	67	69	41	35	53	49	49	41	39	38
FA	63	65	75	73	50	47	73	53	55	41	43	41	51	51
AR	54	67	67	90	51	47	49	45	55	47	43	49	53	47
HI	57	82	65	51	75	58	53	41	60	54	43	39	43	41
BN	45	76	39	31	65	70	53	47	55	49	39	31	35	33
IN	41	65	60	79	50	47	75	65	51	33	43	33	51	48
MS	35	51	53	60	39	39	75	71	47	41	51	47	47	51
EN	64	49	47	53	47	47	49	41	94	80	80	65	51	53
ES	55	39	39	51	39	41	41	41	79	88	79	65	41	41
FR	61	47	45	43	45	31	43	41	77	76	90	70	53	53
DE	67	38	41	47	41	39	43	37	69	67	69	94	41	45
KO	55	41	39	43	41	39	51	49	53	41	38	41	76	71
ZH	53	38	37	38	41	38	49	47	55	38	37	41	68	82

Table 10: Cross-lingual F1 scores across 14 languages using Gemma-3. Rows: test language. Columns: train language. Diagonal (matched train/test) in bold.

Test ↓ / Train →	TR	UR	FA	AR	HI	BN	IN	MS	EN	ES	FR	DE	KO	ZH
TR	92	68	69	60	59	51	41	41	67	59	64	70	53	57
UR	67	88	70	69	71	73	45	39	56	53	53	45	43	41
FA	69	70	82	75	55	51	75	56	60	47	47	47	55	55
AR	59	71	70	93	55	51	53	49	60	51	47	53	57	51
HI	62	84	69	55	84	71	57	45	64	58	47	43	47	45
BN	47	80	43	35	71	78	57	51	60	53	43	35	39	37
IN	45	69	71	81	53	51	84	69	55	37	47	37	55	51
MS	39	55	57	65	43	43	79	79	51	45	55	51	51	55
EN	67	53	51	57	51	51	53	45	97	83	83	69	55	57
ES	59	43	43	55	43	45	45	45	83	92	83	69	45	45
FR	65	51	49	47	49	35	47	45	81	80	95	73	57	57
DE	71	41	45	51	45	43	47	41	73	71	73	96	45	49
KO	59	45	43	47	45	43	55	53	57	45	43	45	81	75
ZH	57	41	41	43	45	41	53	51	59	41	41	45	71	86

Table 11: Cross-lingual F1 scores across 14 languages using Gemma-4. Rows: test language. Columns: train language. Diagonal (matched train/test) in bold.

Cluster	Hate Train Size (per language)	Platform	Time Period
EN-FR-ES-DE	EN: 4500, FR: 3739, ES: 4731, DE: 4503	X	2016–2023
UR-HI-BN	UR: 4500, HI: 2222, BN: 4504	X, FB, WA, YT	BN: 2017–2020; UR, HI: unspecified
FA-AR-UR-IN	FA: 1002, AR: 4503, UR: 4500, IN: 4502	X	FA: 2020–2022; AR: 2016–2018; IN: 2018; UR: unspecified
IN-MS	IN: 4502, MS: 1701	X	IN: 2018; MS: 2022–2023
KO-ZH	KO: 1520, ZH: 2470	X / Web	KO: 2018–2020; ZH: unspecified

Table 12: Summary of dataset characteristics across the five empirically observed transfer clusters.

9 Quantitative Analysis of Transfer Factors

To better understand the factors influencing cross-lingual transfer performance, we conducted a quantitative analysis examining the relationship between several dataset properties and transfer performance. For each pair of languages, the transfer performance was defined as the off-diagonal F1 score from the 14×14 transfer matrix presented in Figure 1. This results in 182 source–target language pairs. For each pair we computed three measurable factors: dataset size difference, temporal overlap, and linguistic similarity. The resulting dataset consists of one row per source–target pair containing the transfer F1 score and the corresponding factor values (Table 16).

9.1 Dataset Size Difference

Dataset size difference measures the disparity between the number of hate instances used for training in the source and target datasets. For each language pair we computed the absolute log-ratio of the hate training sizes:

$$size_diff = |\log(N_{source}) - \log(N_{target})|$$

where N_{source} and N_{target} denote the number of hate training instances reported in Table 1. This produced 182 values corresponding to the off-diagonal entries of the transfer matrix.

9.2 Temporal Overlap

Temporal overlap measures how much the data collection periods of two datasets intersect. For languages where collection periods were available, we computed the overlap between the time ranges using intersection over union:

$$temporal_overlap = \frac{|T_{source} \cap T_{target}|}{|T_{source} \cup T_{target}|}$$

where T_{source} and T_{target} represent the years covered by each dataset. Language pairs with missing collection dates were excluded from this calculation.

9.3 Linguistic Similarity

To measure linguistic proximity between languages, we used typological feature vectors derived from the URIEL/lang2vec database (Bjerva et al., 2019). Each language is represented by a vector of typological and phylogenetic features. Pairwise

similarity between languages was computed using cosine similarity between these vectors:

$$ling_similarity = cosine(v_{source}, v_{target})$$

where v_{source} and v_{target} denote the feature vectors for the source and target languages.

9.4 Correlation Analysis

To assess the relationship between these factors and transfer performance, we computed Spearman rank correlations (Spearman, 1961) between each factor and the transfer F1 scores across all language pairs. The correlation results reported in Table 3 indicate that linguistic similarity has the strongest association with transfer performance, followed by temporal overlap, while dataset size difference shows only a weak relationship. This trend is also visible in several high-performing language pairs with strong typological proximity. For example, transfer from Urdu to Hindi achieves the highest score in the matrix with very high linguistic similarity, while Urdu to Bengali and English to German show similarly strong performance. High transfer is also observed among closely related Romance languages, such as Spanish to French. Likewise, Indonesian transfers effectively to Malay, reflecting their near-identical linguistic structure. However, not all strong transfers align perfectly with linguistic proximity. For instance, Turkish transfers relatively well to English and German. Similarly, Arabic transfers effectively to Urdu, Indonesian to Persian despite belonging to different language families. These cases suggest that while typological similarity contributes to transfer effectiveness, it does not fully account for the patterns observed in the transfer matrix, and that broader discourse or sociocultural similarities may also facilitate cross-lingual transfer.

10 Leave-One-Language-Out Analysis

To verify that the gains observed in the Observed Cluster experiment are not driven by any single language within a cluster, we conduct a leave-one-language-out (LOLO) analysis on the three largest clusters. For each language ℓ in a cluster C , we remove ℓ from the training set, fine-tune a model on the remaining $|C| - 1$ languages, and evaluate it on ℓ . The resulting F1 score is compared against the Solo-Source TL baseline for ℓ and against the gain reported in the main Observed Cluster experiment (Table 2). If the LOLO gains are comparable to the

original cluster gains, this indicates that the held-out language genuinely benefits from the cluster rather than serving as a critical contributor to it.

Tables 13, 14, and 15 present LOLO results for the Persian-Arabic-Urdu-Indonesian, Urdu-Hindi-Bengali, and English-Spanish-French-German clusters, respectively. Across all three clusters, the LOLO gains closely match the original gains reported in Table 2, with most differences within one F1 point. This confirms that the cluster benefit is a property of the cluster as a whole, and supports the robustness of the empirically identified groupings.

Held-out	Solo F1	LOLO Cluster F1	Δ vs Solo	Original Δ
FA	80 ± 0.8	83 ± 0.5	+3	+3
AR	91 ± 0.6	93 ± 0.4	+2	+2
UR	85 ± 0.6	90 ± 0.4	+5	+5
IN	82 ± 0.6	85 ± 0.4	+3	+3

Table 13: LOLO results for Cluster 1: Persian-Arabic-Urdu-Indonesian. “Original Δ ” refers to the gain reported in Table 2

Held-out	Solo F1	LOLO Cluster F1	Δ vs Solo	Original Δ
UR	85 ± 0.6	88 ± 0.5	+3	+5
HI	80 ± 0.7	83 ± 0.5	+3	+5
BN	70 ± 0.7	73 ± 0.5	+3	+4

Table 14: LOLO results for Cluster 2: Urdu-Hindi-Bengali.

Held-out	Solo F1	LOLO Cluster F1	Δ vs Solo	Original Δ
EN	96 ± 0.5	96 ± 0.4	0	+1
ES	90 ± 0.6	94 ± 0.4	+4	+6
FR	94 ± 0.5	95 ± 0.4	+1	+2
DE	90 ± 0.6	93 ± 0.4	+3	+5

Table 15: LOLO results for Cluster 3: English-Spanish-French-German.

source	target	F1	size_diff	ling_similarity	temporal_overlap
TR	FA	66	1.5	0.7	0.29
TR	MS	39	0.97	0.46	0.14
TR	ES	57	0.05	0.47	1
TR	FR	62	0.19	0.48	1
FA	TR	48	1.5	0.7	0.29
FA	AR	63	1.5	0.78	0
FA	BN	43	1.5	0.7	0
FA	IN	67	1.5	0.58	0
FA	KO	40	0.42	0.4	0
AR	FA	80	1.5	0.78	0
AR	MS	53	0.97	0.56	0
AR	ES	55	0.05	0.58	0.29
AR	KO	44	1.09	0.38	0
BN	MS	43	0.97	0.61	0
BN	ES	61	0.05	0.61	0.43
IN	MS	78	0.97	0.92	0
IN	ES	57	0.05	0.63	0.29
IN	FR	57	0.19	0.64	0.29
IN	KO	62	1.09	0.42	0
MS	IN	63	0.97	0.92	0
MS	EN	50	0.97	0.66	0.14
MS	ES	47	1.02	0.63	0.14
MS	FR	55	0.79	0.64	0.14
EN	FA	56	1.5	0.63	0.29
EN	FR	78	0.19	0.88	1
EN	KO	73	1.09	0.45	0.29
ES	TR	65	0.05	0.47	1
ES	FA	50	1.55	0.61	0.29
ES	BN	54	0.05	0.61	0.43
ES	IN	47	0.05	0.63	0.29
FR	TR	54	0.19	0.48	1
FR	BN	58	0.19	0.62	0.43
FR	IN	40	0.19	0.64	0.29
FR	MS	49	0.79	0.64	0.14
DE	FA	45	1.5	0.64	0.29
DE	ES	68	0.05	0.88	1
KO	TR	58	0.11	0.43	0
KO	IN	44	1.09	0.42	0
KO	MS	58	0.11	0.43	0
KO	ES	43	1.14	0.44	0.29

Table 16: Cross-lingual transfer cases with corresponding dataset properties for a representative sample.