

YEZE at SemEval-2026 Task 9: Detecting Multilingual, Multicultural and Multievent Online Polarization via Heterogeneous Ensembling

Fengze Guo Yue Chang

University of Tübingen

{fengze.guo, yue.chang}@student.uni-tuebingen.de

Abstract

This paper presents our system for *SemEval-2026 Task 9: Detecting Multilingual, Multicultural and Multievent Online Polarization*, which identifies polarized social media content in 22 languages through three subtasks: binary detection, target classification, and manifestation identification. We propose a heterogeneous ensemble of multilingual pretrained models, combining XLM-RoBERTa-large and mDeBERTa-v3-base. We investigate techniques such as multi-task learning, translation-based data augmentation, and class weighting to improve classification performance under severe label imbalance. Our findings indicate that independent task modeling combined with class weighting is more effective.

1 Introduction

Social media platforms have expanded public communication worldwide, but they have also amplified online polarization across languages and cultures (Howard, 2019). Polarized content differs not only in targets and stance but also in its linguistic manifestations, making automatic identification important for large-scale analysis of social dynamics and for assisting downstream moderation and policy research.

This paper presents our system for *SemEval-2026 Task 9* (Naseem et al., 2026a), a shared task dedicated to identifying online polarization in multilingual social media.¹ The shared task comprises three subtasks: (1) binary classification to determine whether a post is polarized, (2) multi-label classification of the polarization target type, and (3) multi-label identification of manifestation categories. The official dataset (Naseem et al., 2026b) consists of multilingual social media posts in 22 languages: Amharic, Arabic, Bengali, Burmese, Chinese, English, German, Hausa, Hindi, Italian,

¹The code is available at <https://github.com/FezeGo/SemEval-2026-Task9-Polar>.

Khmer, Nepali, Odia, Persian, Polish, Punjabi, Russian, Spanish, Swahili, Telugu, Turkish, and Urdu. Subtasks 1 and 2 cover all 22 languages, while Subtask 3 excludes Italian, Polish, Russian, and Burmese. Detailed input/output schemas and language coverage are provided in Appendices A and B. All subtasks are evaluated using the macro-averaged F1 score (Macro-F1).

Our system covers all subtasks, modeling each as an independent problem. Our final system combines two complementary multilingual encoders (XLM-RoBERTa-large (Conneau et al., 2020) and mDeBERTa-v3-base (He et al., 2023)) in a weighted ensemble optimized on the development set. For the multi-label subtasks, we use binary relevance with weighted binary cross-entropy to mitigate severe label imbalance.

This task poses several challenges for multilingual modeling. We observe substantial cross-lingual variation in label distributions, with strong prior shift in Subtask 1 and pronounced sparsity in Subtasks 2 and 3, making minority-label learning and calibration central under Macro-F1. Our experiments also suggest that shared multi-task training can introduce negative transfer when sparse fine-grained labels compete with the dominant binary objective, making fine-grained prediction less stable than coarse binary detection. Together, these observations indicate that fine-grained label sparsity, cross-task inconsistency, and negative transfer remain the main bottlenecks, especially for manifestation prediction.

In the official evaluation, we ranked in the Top 10 for 11/22, 16/22, and 17/18 languages in Subtasks 1, 2, and 3, respectively. The detailed per-language rankings are presented in Appendix C.

2 Background

Online polarization is an intricate socio-technical phenomenon. Beyond algorithmic amplification

within “information bubbles” (Build Up, 2025), social network structures often foster echo chambers where bipolar discourse is costly and infrequent (Garimella et al., 2018). Crucially, research suggests that mere exposure to opposing viewpoints can paradoxically increase polarization (Bail et al., 2018), underscoring the deep-seated nature of digital rifts. These divisions have severe offline consequences, such as intensifying ethnic mobilization and marginalizing vulnerable voices in conflict zones (Ali et al., 2025). In NLP, early work largely focused on identifying predictive features for hate speech and offensive content (Waseem and Hovy, 2016; Davidson et al., 2017). Recent frameworks like HateCheck (Röttger et al., 2021) have introduced functional testing to evaluate model robustness against complex linguistic phenomena such as negation and counter-speech.

Recent benchmarks (Naseem et al., 2026b) suggest that while binary detection is relatively mature, fine-grained target and manifestation prediction remains substantially more challenging. In addition, LLM-based approaches have been shown to be less reliable than specialized supervised models for these fine-grained labels. Motivated by these limitations, we focus on supervised multilingual encoder-based models, which provide more stable and controllable optimization in low-resource and imbalanced settings.

Such models are typically based on Transformer architectures (Vaswani et al., 2017), with pre-trained variants such as BERT (Devlin et al., 2019) offering strong contextual representations for downstream classification tasks.

Within this paradigm, a common strategy is to jointly model related objectives using multi-task learning (MTL) (Caruana, 1997). However, prior work has shown that MTL can suffer from negative transfer and task interference, particularly when label distributions are highly imbalanced or heterogeneous (Yu et al., 2020). In contrast, we adopt independent task modeling combined with ensemble learning, which proves to be more robust under severe label sparsity.

3 System Overview

Our system targets multilingual polarization modeling across 22 languages under severe label imbalance. We follow three core design principles: (i) per-subtask specialization to ensure high accuracy on each individual task, (ii) imbalance-aware op-

timization to tackle class imbalance in multi-label subtasks, and (iii) heterogeneous ensembling to enhance robustness across languages and events.

3.1 Model Architecture

We treat the three subtasks as independent classification problems, with separate models trained for each. Our architecture leverages two complementary multilingual encoders to maximize performance across diverse languages and subtasks:

XLm-RoBERTa-large. We select XLm-RoBERTa-large (XLm-R) as the primary backbone because it is a strong multilingual encoder with broad cross-lingual coverage and stable transfer performance in prior work (Conneau et al., 2020).

mDeBERTa-v3-base. We also employ mDeBERTa-v3-base (mDeBERTa), which uses disentangled attention to encode content and relative position separately (He et al., 2023).

Together, these two encoders provide complementary representation spaces and tokenization behaviors, which are then combined through ensembling to increase robustness across different languages and tasks.

3.2 Imbalance-aware Optimization

As shown in Figure 1, the polarized rate varies substantially across languages, indicating strong cross-lingual prior shift. This motivates the use of imbalance-aware optimization, especially for the multi-label subtasks where sparsity is even more severe.

For Subtasks 2 and 3, we adopt a Binary Relevance (BR) formulation (Read et al., 2009), where multi-label predictions are decomposed into independent binary classification tasks. This allows each label to be treated as a separate problem, with a sigmoid activation applied to each logit to yield independent probabilities:

$$\hat{y}_c = P(c = 1 | x) = \sigma(z_c) \in [0, 1].$$

To handle severe class imbalance, we utilize Weighted Binary Cross-Entropy (WBCE) loss, a common approach in cost-sensitive and imbalanced learning (Elkan, 2001; He and Garcia, 2009), where the weight for each class is inversely proportional to the frequency of its occurrences in the training data:

$$\mathcal{L}(y_c, \hat{y}_c) = -[w_c \cdot y_c \log(\hat{y}_c) + (1 - y_c) \log(1 - \hat{y}_c)],$$




Figure 1: Positive rates for Subtask 1 across languages in the merged train+dev set.

where the weight w_c is computed as:

$$w_c = \frac{N_{\text{neg},c}}{\max(N_{\text{pos},c}, 1)}.$$

Here, $N_{\text{neg},c}$ and $N_{\text{pos},c}$ represent the number of negative and positive instances for each label, respectively.

We also evaluated the effectiveness of WBCE relative to other loss functions; a detailed comparison is provided in Section 5.3.

3.3 Heterogeneous Ensemble Strategy

For our final submission, we ensemble the two backbones using weighted probability averaging. For each label c and input post x , the ensembled posterior is:

$$\bar{P}(c = 1 | x) = \alpha P_{\text{XLM-R}}(c = 1 | x) + (1 - \alpha) P_{\text{mDeBERTa}}(c = 1 | x),$$

where α is selected on the *dev* set to maximize Macro-F1 (we use $\alpha = 0.7$).

3.4 Pipeline Summary

Our training and inference pipeline is:

- (i) fine-tune XLM-R and mDeBERTa separately for each subtask,
- (ii) tune the ensemble weight α on the *dev* set,
- (iii) generate test predictions by weighted ensembling and a global threshold $\tau = 0.5$.

Experimental details are provided in Section 4.

3.5 System Variants

We evaluate:

- (a) single-backbone models (XLM-R or mDeBERTa),

- (b) an MTL variant with a shared XLM-R encoder and task-specific heads,
- (c) the heterogeneous ensemble with independent per-subtask training (submitted).

4 Experimental Setup

4.1 Data and Partitioning

We restrict training to the provided task data and do not use external lexicons. The official dataset (Naseem et al., 2026b) provides *train*, *dev*, and *test* partitions.

Development Phase. We used an internal 85/15 split of the official *train* set for model selection. For Subtask 1, standard stratification was applied; for Subtasks 2 and 3, we used iterative stratification (Sechidis et al., 2011) to preserve rare label co-occurrences under extreme sparsity.

Test Phase. Upon the release of *dev* labels, we utilized the official *dev* set as a hold-out validation set for final hyperparameter tuning. The final system was retrained on the union of the official *train* and *dev* sets to generate predictions for the hidden *test* set.

4.2 Evaluation Metrics

All subtasks are evaluated using Macro-F1:

$$\text{Macro-F1} = \frac{1}{|L|} \sum_{l \in L} F1_l.$$

For Subtasks 2 and 3, $F1_l$ is computed independently for each label and then averaged (label-wise macro), which emphasizes rare categories.

Lang.	POLARDETECT (Subtask 1)				POLARTYPE (Subtask 2)				POLARMANIFEST (Subtask 3)			
	XLM-R	mDeB	Ensem	MTL	XLM-R	mDeB	Ensem	MTL	XLM-R	mDeB	Ensem	MTL
Indo-Aryan Family												
hin	0.801	0.799	0.806	0.785	0.764	0.747	0.776	0.726	0.737	0.728	0.741	0.732
urd	0.785	0.748	0.789	0.794	0.779	0.770	0.781	0.785	0.811	0.804	0.815	0.815
ben	0.840	0.836	0.840	0.838	0.351	0.375	0.352	0.350	0.195	0.208	0.202	0.186
nep	0.905	0.888	0.908	0.909	0.731	0.747	0.750	0.721	0.649	0.577	0.653	0.629
pan	0.777	0.773	0.786	0.780	0.414	0.401	0.421	0.382	0.494	0.498	0.504	0.491
ori	0.818	0.743	0.812	0.792	0.568	0.443	0.562	0.556	0.301	0.235	0.277	0.243
European												
eng	0.761	0.787	0.781	0.792	0.489	0.484	0.503	0.481	0.501	0.483	0.500	0.499
deu	0.722	0.685	0.717	0.711	0.567	0.559	0.573	0.561	0.487	0.485	0.493	0.494
spa	0.774	0.770	0.783	0.784	0.626	0.621	0.637	0.631	0.506	0.468	0.505	0.499
ita	0.525	0.662	0.591	0.548	0.323	0.293	0.319	0.310	–	–	–	–
rus	0.764	0.772	0.806	0.794	0.568	0.516	0.558	0.567	–	–	–	–
pol	0.798	0.786	0.810	0.819	0.557	0.508	0.578	0.567	–	–	–	–
Afro-Asiatic & Iranian												
arb	0.814	0.818	0.819	0.830	0.614	0.604	0.626	0.600	0.608	0.591	0.610	0.596
fas	0.790	0.794	0.814	0.780	0.569	0.576	0.580	0.552	0.414	0.365	0.400	0.400
amh	0.782	0.757	0.771	0.784	0.608	0.629	0.649	0.631	0.544	0.510	0.546	0.522
hau	0.774	0.686	0.740	0.774	0.375	0.318	0.381	0.343	0.122	0.115	0.122	0.187
Other Families												
zho	0.896	0.867	0.897	0.892	0.771	0.723	0.782	0.741	0.555	0.455	0.568	0.540
mya	0.862	0.874	0.879	0.871	0.672	0.672	0.692	0.619	–	–	–	–
khm	0.742	0.673	0.727	0.718	0.557	0.675	0.667	0.644	0.361	0.283	0.353	0.341
tel	0.880	0.860	0.884	0.860	0.419	0.414	0.429	0.423	0.389	0.369	0.378	0.350
tur	0.761	0.765	0.777	0.791	0.556	0.549	0.563	0.558	0.504	0.489	0.500	0.506
swa	0.763	0.778	0.778	0.784	0.460	0.484	0.483	0.433	0.547	0.536	0.547	0.540
Average	0.788	0.778	0.796	0.792	0.565	0.550	0.575	0.554	0.485	0.456	0.484	0.476

Table 1: Official Macro-F1 results by language for Subtasks 1–3. **Ensem** is our final submission; *mDeB* denotes mDeBERTa-v3-base. **blue** indicates that the ensemble is tied for or achieves the best score, while **orange** indicates that the best non-ensemble model achieves a strictly higher score than the ensemble.

4.3 Implementation

Preprocessing. We applied minimal preprocessing: rows with missing or empty text were removed, while emojis were retained to preserve affective cues. No task-specific normalization beyond shared-task standardization (e.g., URLs as [URL]) was performed. Tokenization used the default tokenizer with truncation and dynamic padding, and a maximum sequence length of 256.

Training setup. Experiments were conducted on NVIDIA A100 GPUs using PyTorch and Hugging Face Transformers (bf16/tf32). Models were finetuned with AdamW and a linear scheduler (warmup ratio 0.1) with weight decay set to 0. On the internal validation set, we use early stopping with patience 2. The learning rate was 1×10^{-5} for XLM-R (epochs 4, batch size 32) and 2×10^{-5} for mDeBERTa (epochs 5, batch size 64). Additional implementation details are given in Appendix D.

Thresholds and ensembling. For multi-label settings, weighted binary cross-entropy was used with a positive weight computed from label frequencies. Per-label threshold tuning was not adopted due

to performance degradation, so a global threshold $\tau = 0.5$ was used. The ensemble coefficient α was set to 0.7 after grid search.

5 Results and Discussion

5.1 Official Evaluation Results

Table 1 reports official test Macro-F1 by language for all subtasks. We use the ensemble as the final submission across all subtasks to maintain a single consistent pipeline. In Subtask 2, it reaches Top 5 in Amharic (4th), Urdu (5th), Odia (5th), and Polish (5th). In Subtask 3, it ranks in the Top 5 for Amharic (4th), Arabic (3rd), English (5th), Khmer (5th), Spanish (4th), and Urdu (3rd).

5.2 Data Imbalance and Cross-lingual Conditions

Figure 2 summarizes train+dev label distributions across subtasks and languages. Figure 2(a) shows strong cross-lingual prior shift in Subtask 1 (10.75/49.45/90.79% min/med/max polarized rate) and overall skew across subtasks. As shown in Figure 2(b), *Vilification* and *Stereotype* are more frequent than the remaining manifestation labels.

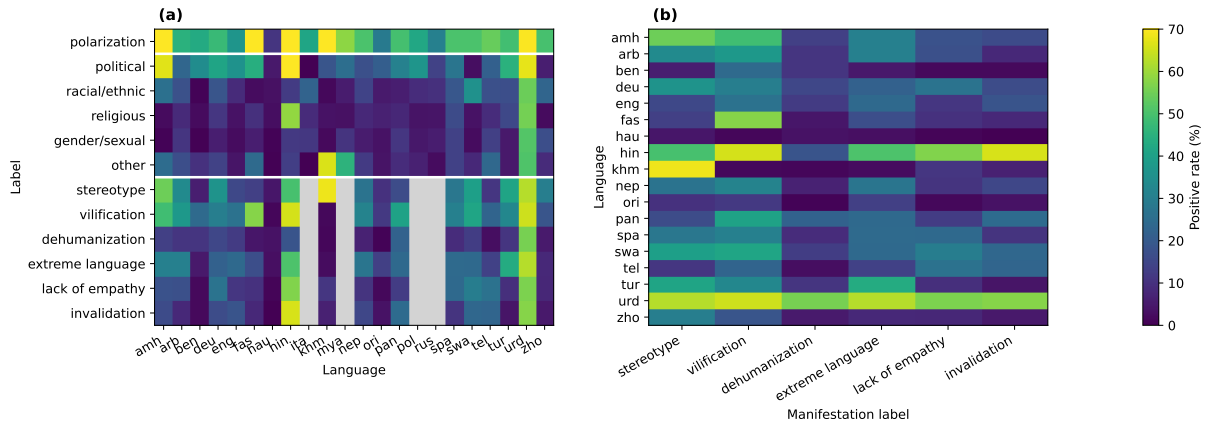


Figure 2: Dataset imbalance analysis on the merged train+dev set. (a) Task-level skew (Language \times Task). (b) Subtask 3 label imbalance (Language \times Label).

Overall, difficulty stems from sparsity interacting with cross-lingual heterogeneity (Hu et al., 2020). Despite partial intra-family consistency (Pires et al., 2019; Wu and Dredze, 2019), extreme sparsity (Lauscher et al., 2020) and script/tokenization effects (Rust et al., 2021) still limit fine-grained transfer.

5.3 Optimization Analysis

Loss functions. Development tests (Table 2) show Focal Loss (Lin et al., 2018) benefits well-represented languages but is unstable under sparsity. Conversely, WBCE significantly boosts low-resource settings (e.g., yielding a +0.212 increase in Macro-F1 for Telugu) and improves the 22-language average by 3.5 points (0.492 \rightarrow 0.527).

Lang.	Base	Focal	WBCE	Δ
zho	0.6905	0.6893	0.7218	+0.0313
hin	0.7724	0.8127	0.7996	+0.0272
tel	0.2253	0.2986	0.4372	+0.2119
amh	0.3760	0.4568	0.4589	+0.0829
hau	0.1115	0.2719	0.2513	+0.1398
Avg.	0.4351	0.5053	0.5338	+0.0987

Table 2: Optimization objective ablation for Subtask 2 (Macro-F1). **Base** = BCE loss; **Focal** = Focal Loss; Δ = the improvement of WBCE over Base.

MTL vs. Independent Modeling. The multi-task learning (MTL) variant suggests negative transfer for the sparsest labels. While shared representations can slightly benefit Subtask 1, Subtasks 2 and 3 suffer when low-frequency labels compete with the dominant binary objective during shared

updates (Yu et al., 2020). We use independent per-subtask modeling to preserve fine-grained signals.

5.4 Error Analysis and Post-hoc Evaluation

Post-hoc evaluation confirms strong cross-lingual variability; additional figures are shown in Appendix E. Weak transfer from Subtask 1 to Subtasks 2 and 3 (e.g., Bengali and Telugu) highlights fine-grained label sparsity. Category-wise, *Political* (Subtask 2), *Vilification* and *Stereotype* (Subtask 3) are the most reliable, while *Gender/Sexual*, *Religious*, and *Other* (Subtask 2), as well as *Dehumanization*, *Lack of Empathy*, and *Invalidation* (Subtask 3), remain challenging.

Miscalibration and label collapse. Low-resource settings exhibit miscalibration (high recall but low precision) and label collapse, where sparse categories fall to near-zero F1, suggesting expanded but noisy decision regions (Desai and Durrett, 2020).

Cross-task inconsistency. Independent modeling causes hierarchical violations: some instances labeled non-polarized in Subtask 1 receive positive Subtask 2 or 3 labels. This suggests a trade-off between robustness and coherence, motivating future post-hoc gating or hierarchical calibration.

Translation augmentation. Machine translation augmentation (using 4,000 Gemini-generated samples for Hausa) failed to yield significant performance gains. This result suggests that synthetic “translationese” and the resulting pragmatic shifts (Ji, 2023) may obscure region-specific idioms and subtle linguistic cues.

6 Conclusion

We presented a multilingual system for SemEval-2026 Task 9 that addresses polarization detection and characterization across 22 languages. Our final system combines XLM-R and mDeBERTa in a heterogeneous ensemble and uses imbalance-aware optimization to improve robustness. Across official evaluations, the system remains competitive and shows that independent per-subtask modeling is a strong practical alternative to shared MTL.

Our analyses show that the main challenges are not only cross-lingual variation, but also fine-grained label sparsity, calibration instability, and cross-task inconsistency. Post-hoc evaluation on the released test gold labels further confirms that high-coverage labels are substantially easier to learn, while sparse manifestation labels remain the most fragile. In future work, we plan to explore lightweight hierarchical calibration (e.g., gating-based consistency constraints) and culturally grounded data synthesis to better capture region-specific socio-political idioms for fine-grained polarization analysis.

Acknowledgments

We would like to express our sincere gratitude to Dr. Çağrı Çöltekin for his invaluable guidance, insightful comments, and continuous support throughout this study.

References

- Adem Chanie Ali, Seid Muhie Yimam, Abinew Ali Ayele, Chris Biemann, and Martin Semmann. 2025. [Silenced voices: social media polarization and women’s marginalization in peacebuilding during the Northern Ethiopia War](#). *i-com*, 24(2):407–432.
- Christopher A. Bail, Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. 2018. [Exposure to opposing views on social media can increase political polarization](#). *Proceedings of the National Academy of Sciences*, 115(37):9216–9221.
- Build Up. 2025. [Polarization footprint europe report](#). Technical report, Build Up.
- Rich Caruana. 1997. [Multitask learning](#). *Machine Learning*, 28(1):41–75.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. [Automated hate speech detection and the problem of offensive language](#). In *Proceedings of the Eleventh International AAAI Conference on Web and Social Media (ICWSM)*, pages 512–515.
- Shrey Desai and Greg Durrett. 2020. [Calibration of pre-trained transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 295–302, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Charles Elkan. 2001. [The foundations of cost-sensitive learning](#). In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2, ICAI’01*, page 973–978, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. [Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship](#). *Preprint*, arXiv:1801.01665.
- Haibo He and Edwardo A. Garcia. 2009. [Learning from imbalanced data](#). *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing](#). *Preprint*, arXiv:2111.09543.
- Jeffrey W. Howard. 2019. [Free Speech and Hate Speech](#). *Annual Review of Political Science*, 22:93–109.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization](#). *Preprint*, arXiv:2003.11080.
- Meng Ji. 2023. [Cultural and linguistic bias of neural machine translation technology](#). In *Translation Technology in Accessible Health Communication*, pages 100–128. Cambridge University Press.

- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online. Association for Computational Linguistics.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2018. [Focal loss for dense object detection](#). *Preprint*, arXiv:1708.02002.
- Usman Naseem, Robert Geislinger, Juan Ren, Sarah Kohail, Rudy Garrido Veliz, P Sam Sahil, Yiran Zhang, Marco Antonio Stranisci, Idris Abdulmumin, Özge Alacam, Cengiz Acartürk, Aisha Jabr, Saba Anwar, Abinew Ali Ayele, Elena Tutubalina, Aung Kyaw Htet, Xintong Wang, Surendrabikram Thapa, Tanmoy Chakraborty, Dheeraj Kodati, Sahar Moradizeyveh, Firoj Alam, Ye Kyaw Thu, Shantipriya Parida, Ihsan Ayyub Qazi, Nelson Odhiambo Onyango, Clemencia Siro, Ibrahim Said Ahmad, Lilian Wanzare, Adem Chanie Ali, Martin Semmann, Chris Biemann, Shamsuddeen Hassan Muhammad, and Seid Muhie Yimam. 2026a. [SemEval-2026 Task 9: Detecting Multilingual, Multicultural and Multient Online Polarization](#). In *Proceedings of the 20th International Workshop on Semantic Evaluation (SemEval-2026)*, San Diego, CA, USA. Association for Computational Linguistics.
- Usman Naseem, Robert Geislinger, Juan Ren, Sarah Kohail, Rudy Garrido Veliz, P Sam Sahil, Yiran Zhang, Marco Antonio Stranisci, Idris Abdulmumin, Özge Alacam, Cengiz Acartürk, Aisha Jabr, Saba Anwar, Abinew Ali Ayele, Simona Frenda, Alessandra Teresa Cignarella, Elena Tutubalina, Oleg Rogov, Aung Kyaw Htet, Xintong Wang, Surendrabikram Thapa, Kritesh Rauniyar, Tanmoy Chakraborty, Arfeen Zeeshan, Dheeraj Kodati, Satya Keerthi, Sahar Moradizeyveh, Firoj Alam, Arid Hasan, Syed Ish-tiaque Ahmed, Ye Kyaw Thu, Shantipriya Parida, Ihsan Ayyub Qazi, Lilian Wanzare, Nelson Odhiambo Onyango, Clemencia Siro, Jane Wanjiru Kimani, Ibrahim Said Ahmad, Adem Chanie Ali, Martin Semmann, Chris Biemann, Shamsuddeen Hassan Muhammad, and Seid Muhie Yimam. 2026b. [Polar: A benchmark for multilingual, multicultural, and multi-event online polarization](#). *Preprint*, arXiv:2505.20624.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. 2009. [Classifier chains for multi-label classification](#). In *Machine Learning and Knowledge Discovery in Databases*, pages 254–269, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Paul Röttger, Bertie Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2021. [HateCheck: Functional tests for hate speech detection models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 41–58. Association for Computational Linguistics.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. [How good is your tokenizer? On the monolingual performance of multilingual language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3118–3135, Online. Association for Computational Linguistics.
- Konstantinos Sechidis, Grigorios Tsoumakas, and Ioannis P. Vlahavas. 2011. [On the stratification of multi-label data](#). In *ECML/PKDD*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Zeerak Waseem and Dirk Hovy. 2016. [Hateful symbols or hateful people? predictive features for hate speech detection on Twitter](#). In *Proceedings of the NAACL Student Research Workshop*, pages 88–93, San Diego, California. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2019. [Beto, Bentz, Becas: The surprising cross-lingual effectiveness of BERT](#). 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 833–844, Hong Kong, China. Association for Computational Linguistics.
- Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. [Gradient surgery for multi-task learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 5824–5836. Curran Associates, Inc.

A Input/Output Schemas

A subset of the training data for each subtask has been selected as illustrative examples. The corresponding input and output schemas for the three subtasks are presented below.

A.1 Subtask 1

This subtask comprises a binary classification task to determine whether a given text exhibits polarized content. A text is labeled as polarized (*Polarization* = 1) only if it explicitly expresses an opinion indicative of attitude polarization, accounting for the overall context and semantic meaning. Texts lacking such characteristics are labeled as non-polarized (*Polarization* = 0).

ID	Text	Polarization
eng_9739...	Is detecting imperialism in the dnd chat.	–
eng_891d...	Start by not listening to msnbc.	–

Table 3: Subtask 1 input examples.

ID	Polarization
eng_9739...	0
eng_891d...	1

Table 4: Subtask 1 output examples.

A.2 Subtask 2

This multi-label classification task identifies the specific targets of polarization in each text. Each text is evaluated across five predefined categories: *political*, *racial/ethnic*, *religious*, *gender/sexual*, and *other*. A label of 1 denotes the presence of a given polarization type, while 0 indicates its absence. Multiple labels may be assigned to a single text if multiple types are present.

ID	Text	Political	Racial/Ethnic	Religious	Gender/Sexual	Other
eng_9739...	Is detecting imperialism in the dnd chat.	–	–	–	–	–
eng_891d...	Start by not listening to msnbc.	–	–	–	–	–

Table 5: Subtask 2 input examples.

ID	Political	Racial/Ethnic	Religious	Gender/Sexual	Other
eng_9739...	0	0	0	0	0
eng_891d...	1	0	0	0	1

Table 6: Subtask 2 output examples.

A.3 Subtask 3

This subtask identifies the specific manifestations of polarization in each text. Each text is evaluated across six predefined categories: *Stereotype*, *Vilification*, *Dehumanization*, *Extreme_Language*, *Lack_of_Empathy*, and *Invalidation*. A label of 1 denotes presence, while 0 indicates absence. Multiple labels may be assigned to a single text if multiple manifestations are present.

For brevity, texts and category names are abbreviated in the input table.

ID	Text	Stereo.	Vilifi.	Dehuman.	Extreme_Language	Lack_of_Empathy	Invalidation
eng_9739...	Is...	-	-	-	-	-	-
eng_891d...	Start...	-	-	-	-	-	-

Table 7: Subtask 3 input examples. Abbreviations: Stereo.=Stereotype; Vilifi.=Vilification; Dehuman.=Dehumanization.

ID	Stereotype	Vilification	Dehumanization	Extreme_Language	Lack_of_Empathy	Invalidation
eng_9739...	0	0	0	0	0	0
eng_891d...	1	0	0	0	1	0

Table 8: Subtask 3 output examples.

B Data Statistics

Table 9 provides the detailed instance counts for each language across the training, development, and test sets.

Language	Code	Train	Dev	Test	Total
Amharic	amh	3332	166	1501	4999
Arabic	arb	3380	169	1521	5070
Bengali	ben	3333	166	1501	5000
Burmese	mya	2889	144	1301	4334
Chinese	zho	4280	214	1927	6421
English	eng	3222	160	1452	4834
German	deu	3180	159	1432	4771
Hausa	hau	3651	182	1644	5477
Hindi	hin	2744	137	1236	4117
Italian	ita	3334	166	1538	5038
Khmer	khm	6640	332	2988	9960
Nepali	nep	2005	100	903	3008
Odia	ori	2368	118	1066	3552
Persian	fas	3295	164	1484	4943
Polish	pol	2391	119	1077	3587
Punjabi	pan	1700	100	809	2609
Russian	rus	3348	167	1508	5023
Spanish	spa	3305	165	1488	4958
Swahili	swa	6991	349	3147	10487
Telugu	tel	2366	118	1066	3550
Turkish	tur	2364	115	1093	3572
Urdu	urd	3563	177	1606	5346
Total	22	67736	3744	33782	105262

Table 9: Detailed dataset statistics for all 22 languages.

These languages cover a broad range of linguistic and cultural contexts, providing a diverse basis for detecting and analyzing online polarization in multilingual settings.

C Detailed Official Rankings

Table 10 provides the official system rankings across all three subtasks for the 22 languages included in the dataset. Subtask 3 was evaluated on a subset of 18 languages; instances where a language was not included in that subtask are marked with a dash (—).

Language	Subtask 1	Subtask 2	Subtask 3
Amharic (amh)	13	4	4
Arabic (arb)	20	8	3
Bengali (ben)	10	6	11
Burmese (mya)	7	9	—
Chinese (zho)	12	11	10
English (eng)	34	8	5
German (deu)	17	6	7
Hausa (hau)	24	7	10
Hindi (hin)	14	10	6
Italian (ita)	20	8	—
Khmer (khm)	14	9	5
Nepali (nep)	9	18	6
Odia (ori)	4	5	6
Persian (fas)	8	11	7
Polish (pol)	9	5	—
Punjabi (pan)	8	14	6
Russian (rus)	7	11	—
Spanish (spa)	10	10	4
Swahili (swa)	17	9	6
Telugu (tel)	10	9	7
Turkish (tur)	17	17	6
Urdu (urd)	8	5	3

Table 10: Official system rankings per language for Subtask 1, Subtask 2, and Subtask 3.

D Implementation Details

D.1 Preprocessing

We apply minimal preprocessing across all sub-tasks. Rows with missing text are removed, and examples whose cleaned text becomes empty are discarded. No task-specific normalization is applied beyond the standardization already present in the shared-task data: URLs are already normalized as [URL], and the released files we use do not contain user placeholders such as @USER. We retain emojis because they may carry affective or evaluative cues relevant to polarization.

For tokenization, we use the default tokenizer associated with each backbone. Texts are encoded with `truncation=True`, `padding=False`, and `max_length=256`; dynamic padding is applied at batch time. In practice, most posts are short, so truncation is triggered only for a small fraction of examples.

D.2 Hyperparameter Search and Training

All models are implemented in PyTorch using the Hugging Face Transformers library. We tune the learning rate over $\{1e-5, 2e-5\}$ and use AdamW with a linear scheduler and warmup ratio 0.1. We set weight decay to 0, early stopping uses patience 2, and gradient clipping is not applied.

XLM-R is trained for 4 epochs, and mDeBERTa for 5 epochs. We use per-device batch sizes of 32 and 64, respectively, with gradient accumulation only when needed due to memory constraints. We also test three random seeds (42, 2025, 3072), but do not use seed ensembling because trends are consistent across runs.

For multi-label settings, weighted binary cross-entropy is implemented through `BCEWithLogitsLoss` with a label-dependent `pos_weight` computed from class frequencies in the corresponding training split. For the final run, `pos_weight` is recomputed on the merged train+dev data. Per-label threshold tuning was explored on development data but was not adopted because it reduced overall Macro-F1 under extreme label sparsity.

D.3 Final Configuration Summary

After development-stage tuning, we use the following final configuration for the submitted system:

- **XLM-RoBERTa-large:** batch size 32, 4 epochs, learning rate 1×10^{-5} .

- **mDeBERTa-v3-base:** batch size 64, 5 epochs, learning rate 2×10^{-5} .
- **Thresholding:** global threshold $\tau = 0.5$ for Subtask 1 and all labels in Subtasks 2 and 3.
- **Ensembling:** weighted probability averaging with $\alpha = 0.7$.

E Additional Post-hoc Visualizations

These figures provide supplementary post-hoc analyses on the released test gold labels and complement the discussion in Section 5. They visualize per-language performance, label-wise F1, and precision–recall gaps that are summarized only briefly in the main text.

E.1 Per-language Performance Across Subtasks

Figure 3 highlights substantial cross-lingual variability and the mismatch between strong Subtask 1 performance and weaker fine-grained performance in some languages. In particular, several languages remain competitive in binary polarization detection but drop noticeably in Subtasks 2 and 3, which is consistent with the role of label sparsity in the multi-label settings.





Figure 3: Post-hoc Macro-F1 by language and subtask on the released test gold labels.

E.2 Additional Analysis for Subtask 2

Figures 4a and 4b provide a label-level view of Subtask 2. Figure 4a shows that some target types are substantially easier to predict than others, consistent with the coverage–learnability pattern discussed in the main text. Figure 4b further shows that several labels exhibit recall-dominant behavior, which is consistent with over-prediction under sparse supervision.



(a) Average label-wise F1 for Subtask 2, showing that some target categories are substantially more stable than others.

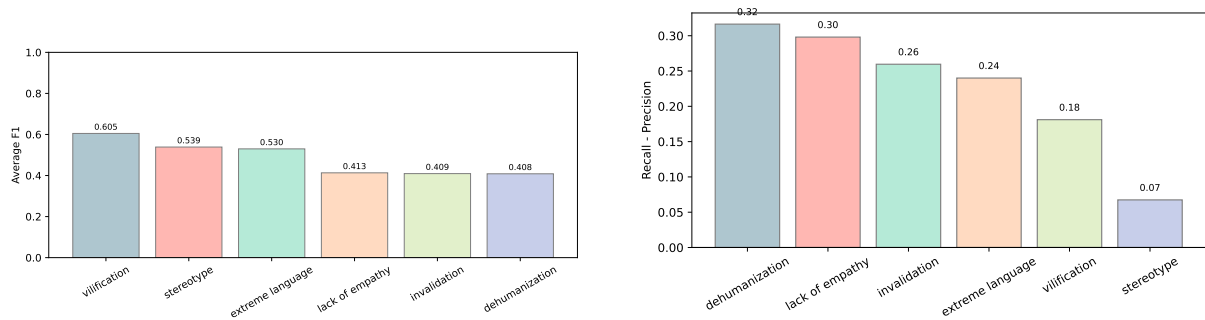
(b) Average precision–recall gap for Subtask 2. Positive values indicate labels for which recall tends to exceed precision, suggesting over-prediction under sparsity.

Figure 4: Supplementary post-hoc visualizations for Subtask 2.

E.3 Additional Analysis for Subtask 3

Figures 5a and 5b provide a corresponding label-level view of Subtask 3. Figure 5a illustrates the relative difficulty of different manifestation categories, while Figure 5b shows that recall often exceeds precision

for the harder labels, supporting the main-text observation that sparse manifestation labels are especially prone to miscalibration.



(a) Average label-wise F1 for Subtask 3, illustrating the relative difficulty of different manifestation categories.

(b) Average precision-recall gap for Subtask 3. Positive values indicate recall-dominant behavior, consistent with miscalibration on sparse labels.

Figure 5: Supplementary post-hoc visualizations for Subtask 3.