AI4PC-Howard University at SemEval-2026 Task 9: Multilingual Polarization Detection via Large Language Model Inference

Multilingual Polarization Detection with LLMs

Evaluating Weak Supervision and Direct LLM Inference Across Languages

Surangana Aryal, Howard University, surangana.aryal@bison.howard.edu

Saurav Keshari Aryal, Howard University, saurav.aryal@howard.edu

Soluchi Fidellbeab, Howard University, soluchi.fidellbeab@bison.howard.edu

This paper describes the PolarNLP system submitted to SemEval-2026 Task 9, Subtask 1, which focuses on detecting political polarization in multilingual text. The task spans 22 typologically diverse languages and poses challenges related to domain shift, class imbalance, and cross-lingual generalization. We explored two modeling strategies: (i) a weakly supervised teacher–student approach that uses a large language model (LLM) to generate pseudolabels for training a multilingual classifier, and (ii) direct LLM-based inference augmented with language-agnostic stylistic features. While the teacher–student approach achieved reasonable in-distribution performance, it failed to generalize to the heldout test set, collapsing toward the majority class. Consequently, our final submission relies on direct LLM inference. We present a detailed analysis of both approaches, highlighting the limitations of weak supervision for polarization detection and the relative robustness of direct LLM reasoning in multilingual settings.

CCS CONCEPTS • Computing methodologies → Artificial intelligence → Natural language processing • Computing methodologies → Machine learning → Learning paradigms → Supervised learning • Applied computing → Law, social and behavioral sciences

Additional Keywords and Phrases: Political polarization detection, Large language models (LLMs), Teacher–student learning, Weak supervision

REFERENCES

[1] SemEval-2026 Task 9: Multilingual Political Polarization Detection. 2026. Codabench Competition. Retrieved from https://www.codabench.org/competitions/10522/

[2] AI at Meta. 2024. Llama 3.1: Open Foundation and Instruction Models. Retrieved from https://ai.meta.com/llama/

[3] Bartowski. 2024. Meta-Llama-3.1-8B-Instruct-GGUF. Hugging Face Repository. Retrieved from https://huggingface.co/bartowski/Meta-Llama-3.1-8B-Instruct-GGUF

[4] Georgi Gerganov. 2023. llama.cpp: Inference of Meta's LLaMA model in pure C/C++. GitHub Repository. Retrieved from https://github.com/ggerganov/llama.cpp

[5] Hugging Face. 2024. Hugging Face Hub. Retrieved from https://huggingface.co/

[6] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, et al. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In Proceedings of ACL 2020. (XLM-RoBERTa)

[7] Google. 2024. Google Colaboratory. Retrieved from https://colab.research.google.com/