

Sequential Fine-Tuning Order and F1-Score Degradation in StressGAN Adversarial Examples for Low-Resource Languages

Assignee Research

June 28, 2026

Abstract

Euphemisms are culturally variable and often ambiguous, posing challenges for language models, especially in low-resource settings. This paper investigates how cross-lingual transfer via sequential fine-tuning affects euphemism detection across five languages: English, Spanish, Chinese, Turkish, and Yoruba. We compare sequential fine-tuning with monolingual and simultaneous fine-tuning using XLM-R and mBERT, analyzing how performance is shaped by language pairings, typological features, and pretraining coverage. Results show that sequential fine-tuning with a high-resource L1 improves L2 perfo

1 Introduction

This paper examines: When Does Language Transfer Help? Sequential Fine-Tuning for Cross-Lingual Euphemism Detection. Research question: To what extent does the order of language exposure in sequential fine-tuning impact F1-score degradation on StressGAN-generated adversarial examples for low-resource languages like Yoruba compared to high-resource languages?.

2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 7.7/10.

3 Results

12 papers retrieved. 5 claims extracted; 4 independently verified. Quality review score: 7.7/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
The model was tested on English (EN), Mandarin Chinese (ZH), Spanish (ES), Turkish (TR), and Yorub (YO).	✓	0.18
The number of examples for the 2025 PETs Datasets are: ZH (3211), EN (3098), ES (2952), TR (2436), YO (2598).	×	0.15
The performance of XLM-R and mBERT on different languages are: EN (XLM-R: 0.821, mBERT: 0.791), ES (XLM-R: 0.768, mBERT:	✓	0.18
The performance of XLM-R and mBERT on different language pairs are: EN & ES (XLM-R: 0.821, mBERT: 0.781), EN & ZH (XLM-R	✓	0.18
The performance of XLM-R and mBERT on sequential fine-tuning for different language pairs are: TR \rightarrow EN (XLM-R: 0.835), ES	✓	0.22

References

- <http://arxiv.org/abs/2508.11831v1>
- <http://arxiv.org/abs/2501.05260v1>
- <http://arxiv.org/abs/1909.07342v1>