

Targeted Lexical Injection for Zero-Shot Named Entity Recognition in Swahili and Zulu

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: What is the impact of Targeted Lexical Injection on zero-shot named entity recognition F1 scores for Swahili and Zulu within the XTREME-R suite relative to full parameter fine-tuning. 17 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Targeted Lexical Injection: Unlocking Latent Cross-Lingual Alignment in Lugha-Llama via Early-Layer LoRA Fine-Tuning. Research question: What is the impact of Targeted Lexical Injection on zero-shot named entity recognition F1 scores for Swahili and Zulu within the XTREME-R suite relative to full parameter fine-tuning?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.8/10.

3 Results

11 papers retrieved. 17 claims extracted; 1 independently verified. Quality review score: 3.8/10.

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 |
|---|----------|------------|
| Layer 0 (input embeddings) of the Lugha-Llama-8B-wura model showed an average cosine similarity of approximately 0.3153 | × | 0.13 |
| Layer 1 of the Lugha-Llama-8B-wura model showed an average cosine similarity of 0.9808 in the pilot study. | × | 0.14 |
| Layer 2 exhibited the peak average cosine similarity of 0.99998 in the pilot study. | × | 0.10 |
| Layer 31 showed an average cosine similarity of 0.9876 in the pilot scan. | × | 0.07 |
| The baseline output similarity on the full evaluation set prior to TLI fine-tuning was approximately 0.32. | ✓ | 0.16 |
| The average cosine similarity at Layer 31 for the trained set prior to TLI fine-tuning was approximately 0.3211. | × | 0.14 |
| The average cosine similarity at Layer 31 for the control set prior to TLI fine-tuning was approximately 0.3143. | × | 0.14 |
| The control set used for evaluation consisted of 63 unseen word pairs. | × | 0.12 |
| A paired t-test was conducted to determine the statistical significance of changes in mean cosine similarity before and | × | 0.07 |
| The base model used is Lugha-Llama-8B-wura. | × | 0.08 |
| Lugha-Llama-8B-wura is built upon the Llama-3 architecture. | × | 0.08 |
| The model was loaded in 4-bit precision using bitsandbytes with NF4 quantization. | × | 0.02 |
| The compute data type used was torch.bfloat16. | × | 0.02 |
| The pilot study extracted embeddings from Layers 0 through 31 of the Lugha-Llama model. | × | 0.08 |
| Layer 0 represents the initial input embeddings in the Lugha-Llama model. | × | 0.06 |
| Embeddings for the final evaluation were extracted from Layer 31 (the last hidden state output). | × | 0.05 |
| Embeddings were mean-pooled over attention-masked tokens and L2-normalized. | × | 0.01 |

References

- <http://arxiv.org/abs/1908.10261v1>
- <http://arxiv.org/abs/2111.07734v1>
- <http://arxiv.org/abs/2506.15415v1>