

Targeted Lexical Injection via LoRA Rank Modulation in Low-Resource Bantu Languages

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How does varying the LoRA rank during early-layer Targeted Lexical Injection impact zero-shot accuracy on the XTREME-R benchmark for low-resource Bantu languages compared to full fine-tuning. 18 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.4/10. 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: How does varying the LoRA rank during early-layer Targeted Lexical Injection impact zero-shot accuracy on the XTREME-R benchmark for low-resource Bantu languages compared to full fine-tuning?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.4/10.

3 Results

15 papers retrieved. 18 claims extracted; 2 independently verified. Quality review score: 4.4/10.

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	×	0.14
Layer 1 of the Lugha-Llama-8B-wura model showed an average cosine similarity of 0.9808 in the pilot study.	✓	0.15
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 was approximately 0.32.	×	0.12
The average cosine similarity at Layer 31 for the trained set prior to TLI fine-tuning was approximately 0.3211.	×	0.13
The average cosine similarity at Layer 31 for the control set prior to TLI fine-tuning was approximately 0.3143.	×	0.13
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	×	0.04
The base model used is Lugha-Llama-8B-wura.	×	0.10
Lugha-Llama-8B-wura is built upon the Llama-3 architecture.	×	0.10
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.	×	0.05
Layer 0 represents the initial input embeddings.	×	0.02
For the evaluation phase, word embeddings were extracted from the final output layer (Layer 31).	×	0.09
Embeddings were mean-pooled over attention-masked tokens and L2-normalized.	×	0.01
Cosine similarity between L2-normalized Swahili and English word embeddings was used as the primary metric for lexical a	✓	0.18

References

- <http://arxiv.org/abs/2409.10965v1>
- <http://arxiv.org/abs/2506.15415v1>
- <http://arxiv.org/abs/2604.06202v1>