

Scaling LoRA Layers in Qwen2-72B for Cross-Lingual Lexical Alignment in Swahili

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: To what extent does scaling the number of LoRA layers (e.g., 2 vs. 4 vs. 6) in Qwen2-72B impact cross-lingual lexical alignment for Swahili in XL-SUM, as evaluated by BLEU and ROUGE scores compared. 9 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.2/10. 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: To what extent does scaling the number of LoRA layers (e.g., 2 vs. 4 vs. 6) in Qwen2-72B impact cross-lingual lexical alignment for Swahili in XL-SUM, as evaluated by BLEU and ROUGE scores compared to full fine-tuning?.

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 4.2/10.

3 Results

12 papers retrieved. 9 claims extracted; 2 independently verified. Quality review score: 4.2/10.

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) showed a modest average cosine similarity of approximately 0.3153.	×	0.07
Layer 1 showed an average cosine similarity of 0.9808.	×	0.08
Layer 2 exhibited the peak average cosine similarity, reaching 0.99998.	×	0.08
Layer 31 showed an average similarity of 0.9876 in the pilot scan.	×	0.04
The baseline output similarity observed on the full evaluation set was approximately 0.32.	×	0.09
The average cosine similarity at the final output layer (Layer 31) of the base model was approximately 0.3211 for the tr	✓	0.15
Lugha-Llama-8B-wura is an open-source LLM specifically adapted for several African languages, including Swahili, built u	×	0.12
The model is loaded in 4-bit precision using bitsandbytes with NF4 quantization and torch.bfloat16 as the compute data t	×	0.02
The pilot study revealed that Lugha-Llama-8B-wura inherently achieves very high lexical alignment in its early layers, p	✓	0.19

References

- <http://arxiv.org/abs/2602.05988v1>
- <http://arxiv.org/abs/2506.15415v1>
- <http://arxiv.org/abs/2003.04866v1>