

Fine-Tuning Lughha-Llama with TLI on Multilingual Code Corpora and MBPP Performance

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: To what extent does fine-tuning Lughha-Llama with TLI on multilingual code corpora improve its performance on the MBPP benchmark compared to code-specific models like CodeLlama or CodeT5+. 17 claims were extracted from source literature; 2 were 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 Lughha-Llama via Early-Layer LoRA Fine-Tuning. Research question: To what extent does fine-tuning Lughha-Llama with TLI on multilingual code corpora improve its performance on the MBPP benchmark compared to code-specific models like CodeLlama or CodeT5+?.

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; 2 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	✓	0.20
Layer 1 of the Lugha-Llama-8B-wura model showed an average cosine similarity of 0.9808 for Swahili-English word pairs in	✓	0.22
Layer 2 exhibited the peak average cosine similarity of 0.99998 in the pilot study of the Lugha-Llama-8B-wura model.	×	0.14
Layer 31 showed an average cosine similarity of 0.9876 in the pilot scan of the Lugha-Llama-8B-wura model.	×	0.11
The baseline output similarity on the full evaluation set for the base model was approximately 0.32.	×	0.10
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.	×	0.11
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 in the study is Lugha-Llama-8B-wura.	×	0.08
Lugha-Llama is built upon the Llama-3 architecture.	×	0.05
The model was loaded in 4-bit precision using bitsandbytes with NF4 quantization.	×	0.02
The compute data type used for the model was torch.bfloat16.	×	0.02
The pilot study extracted embeddings from every transformer layer, ranging from Layer 0 to Layer 31.	×	0.05
Layer 0 represents the initial input embeddings in the Lugha-Llama model architecture.	×	0.05
For the final evaluation, word embeddings were extracted from Layer 31 (the final output layer).	×	0.08
Embeddings used for evaluation were mean-pooled over attention-masked tokens and L2-normalized.	×	0.01

References

- <http://arxiv.org/abs/2508.11281v3>
- <http://arxiv.org/abs/2509.22472v1>
- <http://arxiv.org/abs/2506.15415v1>