

Early-Layer LoRA Fine-Tuning of Lugha-Llama on Swahili XTREME-R Benchmarks

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: How does the performance of Lugha-Llama with Early-Layer LoRA fine-tuning compare to other parameter-efficient fine-tuning methods (e.g., Adapter, Prefix-Tuning) on the Swahili XTREME-R task in terms of accuracy and training efficiency? 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 1.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: LoRA-FAIR: Federated LoRA Fine-Tuning with Aggregation and Initialization Refinement. Research question: How does the performance of Lugha-Llama with Early-Layer LoRA fine-tuning compare to other parameter-efficient fine-tuning methods (e.g., Adapter, Prefix-Tuning) on the Swahili XTREME-R task in terms of accuracy and training efficiency?.

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 1.5/10.

3 Results

12 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 1.5/10.

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.

References

- <http://arxiv.org/abs/2606.01947v1>
- <http://arxiv.org/abs/2411.14961v3>
- <http://arxiv.org/abs/2506.15415v1>