

Comparative Analysis of Prefix-Tuning and Adapter-Based Fine-Tuning for Zero-Shot Cross-Lingual Generation on Low-Resource

Assignee Research

June 14, 2026

Abstract

With the release of new large language models (LLMs) like Llama and Mistral, zero-shot cross-lingual transfer has become increasingly feasible due to their multilingual pretraining and strong generalization capabilities. However, adapting these decoder-only LLMs to new tasks across languages remains challenging. While parameter-efficient fine-tuning (PeFT) techniques like Low-Rank Adaptation (LoRA) are widely used, prefix-based techniques such as soft prompt tuning, prefix tuning, and Llama Adapter are less explored, especially for zero-shot transfer in decoder-only models. We present a compre

1 Introduction

This paper examines: Zero-Shot Cross-Lingual Transfer using Prefix-Based Adaptation. Research question: How does prefix-tuning compare to adapter-based fine-tuning in zero-shot cross-lingual generation performance on low-resource African languages within the XTREME benchmark?.

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 7.5/10.

3 Results

11 papers retrieved. 26 claims extracted; 19 independently verified. Quality review score: 7.5/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
Zhao and Schtze (2021) systematically compared discrete prompting, soft prompting, and fine-tuning on the few-shot mult	✓	0.29
Tu et al. (2022) compared prompt tuning with fine-tuning across diverse NLU tasks on XLM-R and mBERT.	✓	0.26
Tu et al. (2022) evaluated prefix tuning on the encoder-only XLM-R model and showed its effectiveness over full fine-tun	✓	0.35
Tu et al. (2022) investigated the decoder-based multilingual model XGLM, limiting their analysis to a single small model	✓	0.26
Tu et al. (2022) showed that prompt tuning can sometimes surpass fine-tuning, particularly for low-resource languages.	✓	0.32
Experiments were conducted on Llama 3.1 (8B) and Mistral v0.3 (7B).	✓	0.20
To study model scaling, Llama 3.2 (1B) and Mistral Small (24B) were evaluated.	✓	0.17
Mistral v0.3 (7B) has an extended vocabulary compared to Mistral v0.1.	✓	0.21
Mistral Small (24B) is categorized as a 'small' LLM (under 70B).	×	0.15
The experiments were limited to base model variants only.	×	0.11
XQUAD (Artetxe et al., 2019) was used as a benchmark for cross-lingual question answering.	✓	0.19
XNLI (Conneau et al., 2018) was used as a benchmark for cross-lingual natural language inference.	✓	0.20
Belebele (Bandarkar et al., 2024) was used as a benchmark for cross-lingual machine reading comprehension.	✓	0.21
MGSM (Shi et al., 2023) was used to assess reasoning capabilities in multilingual settings.	✓	0.20
Prefix-based adaptation methods and LoRA with rank 4 were fine-tuned using the English SQuAD training set containing 87.	✓	0.22
A subset of the English NLI training data containing 100K samples was used for XNLI evaluations.	✓	0.22
For Belebele, a training set containing 67.5K English samples was used.	✓	0.22
The GSM8K English training dataset with 7.47K samples was used for MGSM evaluation.	✓	0.17
Learning rates of 3e-3, 1e-3, and 3e-4 were experimented with.	✓	0.17
On the Kyrgyz language benchmark, the Base model scored 37.2, LoRA4 scored 52.9, Soft Prompt scored 50.2, Llama Adapter	×	0.09

References

- <http://arxiv.org/abs/2312.16159v1>
- <http://arxiv.org/abs/2310.09917v3>
- <http://arxiv.org/abs/2510.24619v1>