

Impact of Model Architecture on Inference Latency and Accuracy in Zero-Shot Cross-Lingual Transfer

Assignee Research

July 22, 2026

Abstract

Multilingual BERT (mBERT), a language model pre-trained on large multilingual corpora, has impressive zero-shot cross-lingual transfer capabilities and performs surprisingly well on zero-shot POS tagging and Named Entity Recognition (NER), as well as on cross-lingual model transfer. At present, the mainstream methods to solve the cross-lingual downstream tasks are always using the last transformer layer’s output of mBERT as the representation of linguistic information. In this work, we explore the complementary property of lower layers to the last transformer layer of mBERT. A feature aggregat

1 Introduction

This paper examines: Feature Aggregation in Zero-Shot Cross-Lingual Transfer Using Multilingual BERT. Research question: What is the impact of model architecture (e.g., Transformer vs. Mixture-of-Experts) on inference latency and accuracy when fine-tuning on English versus multilingual intermediate tasks for zero-shot cross-lingual transfer on XTREME-R?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 7.7/10.

3 Results

14 papers retrieved. 10 claims extracted; 10 independently verified. Quality review score: 7.7/10.

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
Lower layers of mBERT provide more cross-lingual information while upper layers provide more language structure informat	✓	0.28
The output of layers before the last layer can provide supplementary information to the last layer of mBERT for differen	✓	0.30
The feature aggregation module based on an attention mechanism improves the performance of mBERT on all four cross-lingu	✓	0.31
The best results of aggregation models in each task outperform the baseline by 1 to 3 absolute percentage points.	✓	0.19
The best performances of the four tasks are obtained with different fusion layers.	✓	0.21
Languages that belong to the same language family as English are denoted as 'enf' and those that belong to different lan	✓	0.21
The DLFA module integrates representations from the last and one of the lower layers of mBERT, and the fusion embeddings	✓	0.20
The AIF module extracts global and local information via two branches and element-wisely multiplies the result with the	✓	0.31
The AIF module uses two convolution layers to obtain information dynamically according to the requirements of different	✓	0.23
The representation generated from transformers in mBERT is a tensor with the dimension as $B \times T \times E$, where B is the batc	✓	0.29

References

- <http://arxiv.org/abs/2606.01947v1>
- <http://arxiv.org/abs/2205.08497v1>
- <http://arxiv.org/abs/2005.00633v1>