

Performance comparison of multimodal and textual intermediate tasks in zero-shot cross-lingual transfer on XTREME-R

Assignee Research

July 23, 2026

Abstract

In zero-shot cross-lingual transfer, a supervised NLP task trained on a corpus in one language is directly applicable to another language without any additional training. A source of cross-lingual transfer can be as straightforward as lexical overlap between languages (e.g., use of the same scripts, shared subwords) that naturally forces text embeddings to occupy a similar representation space. Recently introduced cross-lingual language model (XLM) pretraining brings out neural parameter sharing in Transformer-style networks as the most important factor for the transfer. In this paper, we aim

1 Introduction

This paper examines: Analyzing Zero-shot Cross-lingual Transfer in Supervised NLP Tasks. Research question: How does the performance of multimodal intermediate-task training (e.g., image-text alignment) compare to purely textual intermediate tasks in zero-shot cross-lingual transfer accuracy on the XTREME-R 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 7.7/10.

3 Results

11 papers retrieved. 15 claims extracted; 12 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
XLM by Conneau & Lample [2] has reported the state-of-the-art on downstream benchmarks in cross-lingual language underst	✓	0.21
Recent studies show that parameter sharing induced by the Transformer architecture is the most attributable factor for c	✓	0.22
XLM-RoBERTa (XLM-R) [1] produces token embeddings (up to 512 token vectors of 768 dimensions each) for a given input.	✓	0.18
Sentence embeddings for tasks like semantic textual similarity (STS) are obtained by averaging the token embedding outpu	✓	0.22
A siamese network architecture by Sentence-BERT [7] is adopted for sentence-pair modeling to avoid combinatorial explosi	✓	0.22
Cross-lingual mapping for sentence-level transformations is learned directly from sentence-pair examples.	✓	0.20
A projection matrix is computed using linear algebraic methods to achieve fine-grained alignment of sentence embeddings	✓	0.19
A single-layer neural net is used to iteratively learn the projection matrix by gradient descent.	✓	0.24
The STSb sentence pairs are labeled with a similarity score ranging from 0 to 5.	✓	0.21
The paper extends semantic textual similarity, SQuAD [3] & KorQuAD [4] question answering, and sentiment classifications	×	0.14
The paper computes a projection that directly maps sentence embeddings of one language to those of another.	×	0.15
The paper analyzes the effect of fine-grained alignment of sentences across different languages on the quality of zero-s	✓	0.22
The paper provides rigorous results on cross-lingual transfer in three important supervised NLP tasks: STS, MRC, and sen	✓	0.18
The paper proposes to directly compute a cross-lingual mapping that aligns sentence embeddings of different languages.	✓	0.22
The paper shows benefits of the fine-grained cross-lingual sentence alignment for sentence-pair regression tasks.	×	0.14

References

- <http://arxiv.org/abs/2005.13013v2>
- <http://arxiv.org/abs/2101.10649v1>
- <http://arxiv.org/abs/2104.08645v2>