

Synthetic Parallel Corpora and Zero-Shot Cross-Lingual Robustness in Multimodal Image-Text Retrieval

Assignee Research

July 21, 2026

Abstract

Zero-shot cross-lingual transfer is a central task in multilingual NLP, allowing models trained in languages with more sufficient training resources to generalize to other low-resource languages. Earlier efforts on this task use parallel corpora, bilingual dictionaries, or other annotated alignment data to improve cross-lingual transferability, which are typically expensive to obtain. In this paper, we propose a simple yet effective method, SALT, to improve the zero-shot cross-lingual transfer of the multilingual pretrained language models without the help of such external data. By incorporati

1 Introduction

This paper examines: Self-Augmentation Improves Zero-Shot Cross-Lingual Transfer. Research question: What is the impact of data augmentation with synthetic parallel corpora on the robustness of zero-shot cross-lingual transfer for multimodal models (e.g., CLIP, BLIP) on image-text retrieval benchmarks like Flickr30k?.

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 7.5/10.

3 Results

14 papers retrieved. 10 claims extracted; 7 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
SALT achieves 1.4% improvement on 15 languages of XNLI and 4.1% on 6 languages of PAWS-X in terms of average accuracy	✓	0.20
SALT adopts a modified masked language modeling (MLM) method to distill cross-lingual token translation pairs.	✓	0.25
SALT makes two modifications to the original MLM learning objective: predicting tokens only in a specific target language	✓	0.23
Generated token-level cross-lingual substitutions with high enough predicted probabilities are used for code-switching.	✓	0.26
SALT predicts synonyms in the source language with a different threshold for synonym prediction.	×	0.12
SALT proposes an online self-augmentation technique based on cross-lingual embedding mixup.	✓	0.16
The mixed-up token embedding is generated as $h_i = r \cdot h_{si} + (1 - r) \cdot h_{ti}$, where $r = \{r_j\}$, $r_j \in [0, 1]$ is a random vector	✓	0.27
The interpolation coefficient vector r is dynamically sampled in each step of training for each instance.	✓	0.16
SALT achieves an average accuracy of 69.8 on XNLI with 15 languages, which is higher than the mBERT baseline of 67.8.	×	0.12
SALT achieves an average accuracy of 66.8 on XNLI without English, which is higher than the mBERT baseline of 65.4.	×	0.11

References

- <http://arxiv.org/abs/2410.14154v1>
- <http://arxiv.org/abs/2309.10891v1>
- <http://arxiv.org/abs/2104.08645v2>