

Quantized vs. Full-Precision Inference Impact on XGLUE Per-Token Latency

Assignee Research

July 25, 2026

Abstract

Intermediate-task training—fine-tuning a pretrained model on an intermediate task before fine-tuning again on the target task—often improves model performance substantially on language understanding tasks in monolingual English settings. We investigate whether English intermediate-task training is still helpful on non-English target tasks. Using nine intermediate language-understanding tasks, we evaluate intermediate-task transfer in a zero-shot cross-lingual setting on the XTREME benchmark. We see large improvements from intermediate training on the BUCC and Tatoeba sentence retrieval tas

1 Introduction

This paper examines: English Intermediate-Task Training Improves Zero-Shot Cross-Lingual Transfer Too. Research question: What is the impact of quantized (e.g., 4-bit, 8-bit) versus full-precision inference on the per-token latency difference between English intermediate-task training and direct zero-shot cross-lingual transfer on the XGLUE benchmark?.

2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 9.0/10.

3 Results

10 papers retrieved. 9 claims extracted; 9 independently verified. Quality review score: 9.0/10.

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
Intermediate-task training on SQuAD, MNLI, and HellaSwag yields large target-task improvements of 8.2, 7.5, and 7.0 point	✓	0.27
Multi-task intermediate-task training on all 9 tasks performs best, improving by 8.7 points.	✓	0.25
Applying intermediate-task training to BUCC and Tatoeba, the two sentence retrieval target tasks that have no training d	✓	0.29
TyDiQA shows consistent improvements with many intermediate tasks, whereas XNLI does not see benefits from intermediate	✓	0.17
Evaluating our best performing models for each target task on the XTREME benchmark yields an average improvement of 5.4	✓	0.33
Training on English intermediate tasks outperforms the more complex alternatives of (i) continuing multilingual MLM duri	✓	0.32
We use the pretrained XLM-R Large model (Conneau et al., 2020) as a starting point for all our experiments, as it curren	✓	0.32
We follow a three-phase approach to training: (i) we use a publicly available MLM; (ii) we perform intermediate-task tra	✓	0.42
We study the effect of intermediate-task training (STILTs; Phang et al., 2018) with nine different English intermediate	✓	0.27

References

- <http://arxiv.org/abs/2203.15952v4>
- <http://arxiv.org/abs/2004.01401v3>

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