

Scaling AdaptToken from 3B to 7B Parameters: Accuracy-Latency Trade-offs on AdvGLUE with Adversarial Training

Assignee Research

June 6, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: What is the impact of scaling AdaptToken-3B's model size (e.g., 3B to 7B parameters) on AdvGLUE benchmark performance when combining mean shift-based feature space analysis with adversarial training. 7 claims were extracted from source literature; 7 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 8.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: On the Impact of White-box Deployment Strategies for Edge AI on Latency and Model Performance. Research question: What is the impact of scaling AdaptToken-3B's model size (e.g., 3B to 7B parameters) on AdvGLUE benchmark performance when combining mean shift-based feature space analysis with adversarial training, measured by accuracy and inference latency?.

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 8.7/10.

3 Results

10 papers retrieved. 7 claims extracted; 7 independently verified. Quality review score: 8.7/10.

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
The study empirically assesses the accuracy vs latency trade-off of white-box and black-box operators and their combinat	✓	0.31
The study includes inference experiments with 3 white-box operators (QAT, Pruning, Knowledge Distillation), 2 black-box	✓	0.35
The experiments are conducted across 3 tiers (Mobile, Edge, Cloud) on 4 commonly-used Computer Vision and Natural Langua	✓	0.29
The combination of Distillation and SPTQ operators (DSPTQ) should be preferred over non-hybrid operators when lower late	✓	0.40
Among the non-hybrid operators, the Distilled operator is a better alternative in both mobile and edge tiers for lower l	✓	0.45
Operators involving distillation show lower latency in resource-constrained tiers (Mobile, Edge) compared to the operato	✓	0.39
For textual subject models, which have low input data size requirements, the Cloud tier is a better alternative for the	✓	0.38

References

- <http://arxiv.org/abs/2511.12188v1>
- <http://arxiv.org/abs/2411.00907v3>
- <http://arxiv.org/abs/2005.10190v4>