

Scaling Adversarial Training Batch Sizes for Defense Transferability in Code Language Models

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: Does scaling adversarial training batch sizes improve the transferability of defenses from image classifiers to code language models like CodeT5 under PGD attacks. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: How Does Critical Batch Size Scale in Pre-training?. Research question: Does scaling adversarial training batch sizes improve the transferability of defenses from image classifiers to code language models like CodeT5 under PGD attacks?.

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 6.8/10.

3 Results

11 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.8/10.

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.

References

- <http://arxiv.org/abs/2410.21676v4>
- <http://arxiv.org/abs/2307.02055v1>
- <http://arxiv.org/abs/2408.13274v1>