

Scaling Synthetic Adversarial Pretraining Data for Tabular Foundation Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: What is the impact of scaling the synthetic adversarial pretraining dataset size on the downstream task accuracy of tabular foundation models, as measured by AUC scores on the TabRobust benchmark. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: TabDPT: Scaling Tabular Foundation Models on Real Data. Research question: What is the impact of scaling the synthetic adversarial pretraining dataset size on the downstream task accuracy of tabular foundation models, as measured by AUC scores on the TabRobust benchmark?.

2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.8/10.

3 Results

13 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.8/10.

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/2402.04177v3>
- <http://arxiv.org/abs/2410.18164v3>
- <http://arxiv.org/abs/2512.03307v1>