

Robustness of Synthetic vs. Real-World Pretrained Tabular Foundation Models Under Adversarial Noise

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: How does the robustness of tabular foundation models pretrained on synthetic data compare to those pretrained on real-world datasets when evaluated on TabMNAR with adversarial noise, measured by. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: TabularBench: Benchmarking Adversarial Robustness for Tabular Deep Learning in Real-world Use-cases. Research question: How does the robustness of tabular foundation models pretrained on synthetic data compare to those pretrained on real-world datasets when evaluated on TabMNAR with adversarial noise, measured by accuracy degradation across Gaussian and salt-and-pepper noise types?.

2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

12 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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.

References

- <http://arxiv.org/abs/2408.07579v1>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2207.03208v2>