Adversarial Noise Scaling Enhances Robustness in Tabular Foundation Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: To what degree does scaling the adversarial noise intensity during synthetic data generation improve the robustness of tabular foundation models against adversarial attacks, as evaluated by. 11 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Robust Tabular Foundation Models. Research question: To what degree does scaling the adversarial noise intensity during synthetic data generation improve the robustness of tabular foundation models against adversarial attacks, as evaluated by robustness metrics on the TabM-NAR 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 4.5/10.

3 Results

13 papers retrieved. 11 claims extracted; 1 independently verified. Quality review score: 4.5/10.

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
Tabular foundation models (TFMs) have emerged as a promising direction for classification and regression tasks with stru	✓	0.15
TFMs rely on in-context learning (ICL).	×	0.02
TFMs can provide high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.07
Current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribution	×	0.06
Fixed priors underrepresent certain regions of the parameter space, potentially degrading performance on real-world data	×	0.05
State-of-the-art TFMs still lag behind tree-based methods on some benchmarks.	×	0.07
Training TFMs relies on generating a large amount of diverse synthetic datasets.	×	0.08
The generation process relies on constructing structural causal models (SCMs) from which datasets can be sampled.	×	0.03
The structure of these SCMs is implicitly parameterized, giving significant control over the data generation process.	×	0.07
RTFM can significantly improve the ranking of TabPFN on several real-world tabular benchmarks with only 90k additional t	×	0.09
The optimality gap can be computed in a matter of seconds when parallelized, given sufficient CPU cores (ncores = nds $\cdot$	×	0.04

References

- http://arxiv.org/abs/2411.15497v3
- http://arxiv.org/abs/2512.03307v1
- http://arxiv.org/abs/2407.13111v1