Scaling Noise Magnitude Enhances Adversarial Robustness in Tabular Foundation Models

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: Does scaling the noise magnitude in synthetic data generation improve the adversarial robustness of tabular foundation models, as measured by accuracy on TabMNAR and other adversarial benchmarks like. 12 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.5/10. 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: Does scaling the noise magnitude in synthetic data generation improve the adversarial robustness of tabular foundation models, as measured by accuracy on TabMNAR and other adversarial benchmarks like TabRobust?.

2 Methodology

Systematic literature search across multiple databases yielded 8 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.5/10.

3 Results

8 papers retrieved. 12 claims extracted; 0 independently verified. Quality review score: 3.5/10.

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) rely on in-context learning (ICL) for classification and regression tasks with structur	×	0.12
TFMs can produce high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.08
Training TFMs relies on generating diverse synthetic datasets constructed from structural causal models (SCMs).	×	0.08
All current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribut	×	0.06
Fixed priors in TFM training underrepresent certain regions of the parameter space, potentially degrading performance on	×	0.06
State-of-the-art TFMs lag behind tree-based methods on some benchmarks.	×	0.06
The proposed RTFM method was applied to TabPFN V2.	×	0.11
Training TabPFN V2 with RTFM using only 90k additional training datasets significantly improved its ranking on several r	×	0.11
The maximization stage of the proposed algorithm uses a black-box optimization algorithm to search the SCM parameter spa	×	0.02
In the described implementation, the estimated optimality gap could be computed in a matter of seconds when parallelized	×	0.04
The implementation used 20 sampled generators (nds=20) and 7 baseline estimators (e=7) to compute the optimality gap.	×	0.05
The number of CPU cores used for parallelization was set to the product of the number of sampled generators and baseline	×	0.02

References

- http://arxiv.org/abs/2507.07829v1
- http://arxiv.org/abs/2512.03307v1

- http://arxiv.org/abs/2307.00161v1