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

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does the robustness of tabular foundation models pretrained on synthetic data with scaled adversarial noise compare to those pretrained on real-world tabular datasets, as measured by accuracy. 12 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.0/10. 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: How does the robustness of tabular foundation models pretrained on synthetic data with scaled adversarial noise compare to those pretrained on real-world tabular datasets, as measured by accuracy degradation on adversarial samples from the TabMNAR benchmark?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.0/10.

3 Results

14 papers retrieved. 12 claims extracted; 0 independently verified. Quality review score: 3.0/10.

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	×	0.10
TFMs can produce high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.07
Training TFMs relies on generating diverse synthetic datasets constructed from structural causal models (SCMs).	×	0.12
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	×	0.05
State-of-the-art TFMs lag behind tree-based methods on some benchmarks.	×	0.06
The proposed RTFM algorithm applied to TabPFN V2 used only 90k additional training datasets.	×	0.11
Applying RTFM to TabPFN V2 significantly improved its ranking on several real-world tabular benchmarks.	×	0.11
In the maximization stage of the proposed method, the model gW is frozen to maximize the optimality gap.	×	0.04
The study used a black-box optimization algorithm to search the SCM parameter space for parameters with large optimality	×	0.02
The estimated optimality gap computation was parallelized using a number of CPU cores equal to the product of the number	×	0.06
With nds=20 and e=7, the estimated optimality gap could be computed in a matter of seconds when parallelized.	×	0.04

References

- http://arxiv.org/abs/2601.04110v2

- http://arxiv.org/abs/2208.00428v1
- http://arxiv.org/abs/2512.03307v1