Adversarial Noise Injection in Synthetic Pretraining for Tabular Foundation Model Robustness

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: How does adversarial noise injection during synthetic pretraining affect the robustness of tabular foundation models on the Tab-OOD benchmark compared to clean synthetic data. 11 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.8/10. 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 adversarial noise injection during synthetic pretraining affect the robustness of tabular foundation models on the Tab-OOD benchmark compared to clean synthetic data?.

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 2.8/10.

3 Results

12 papers retrieved. 11 claims extracted; 0 independently verified. Quality review score: 2.8/10.

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.

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.	×	0.07
Training TFMs relies on generating diverse synthetic datasets constructed from structural causal models (SCMs).	×	0.07
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 algorithm is a model-agnostic two-stage adversarial training algorithm for TFMs.	×	0.12
Applying RTFM to TabPFN V2 using only 90k additional training datasets significantly improves its ranking on several rea	×	0.11
The maximization stage of the proposed method uses a black-box optimization algorithm to search the SCM parameter space	×	0.05
In the described implementation, estimating the optimality gap with nds=20 and e=7 takes a matter of seconds when parall	×	0.03
Table (p10) lists synthetic dataset generation parameters including feature counts ranging from 3 to 128 and activation	×	0.03

References

- http://arxiv.org/abs/2207.05295v2
- http://arxiv.org/abs/2504.20900v1
- http://arxiv.org/abs/2512.03307v1