Adversarial Pretraining of Tabular Foundation Models on Synthetic Data and Robustness Gains

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How does pretraining tabular foundation models on synthetic adversarial data affect their robustness scores on the TabTime benchmark compared to real-world pretraining. 17 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: How does pretraining tabular foundation models on synthetic adversarial data affect their robustness scores on the TabTime benchmark compared to real-world pretraining?.

2 Methodology

Systematic literature search across multiple databases yielded 15 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

15 papers retrieved. 17 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	×	0.13
TFMs rely on in-context learning (ICL).	×	0.03
TFMs can provide high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.08
TFMs are pretrained using synthetic data generated from structural causal models (SCMs).	×	0.13
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	×	0.07
State-of-the-art TFMs still lag behind tree-based methods on some benchmarks.	×	0.06
The work introduces an optimality gap concept to target regions where the TFM underperforms relative to the best achieva	×	0.09
The proposed RTFM algorithm can significantly improve the ranking of TabPFN on several real-world tabular benchmarks wit	×	0.10
The data generation process for TFMs relies on constructing structural causal models (SCMs).	×	0.11
The structure of SCMs is implicitly parameterized, giving significant control over the data generation process.	×	0.04
The work leverages the significant control provided by the data generation process to frame TFM training from an adversa	×	0.09
The work proposes an efficient, model-agnostic two-stage adversarial training algorithm for TFMs, called ROBUST TABULAR	✓	0.20
The work applies RTFM to TabPFN V2, showing significant improvement in ranking on several real-world tabular benchmarks.	×	0.09
The maximization stage of RTFM involves freezing the model to maximize the optimality gap in its current state.	×	0.05
The work uses a black-box optimization algorithm to efficiently search the parameter space for parameters with large opt	×	0.05
For nds = 20 and e = 7, the estimated optimality gap can be computed in a matter of seconds when parallelized, given suf	×	0.04

References

- http://arxiv.org/abs/2601.04110v2
- http://arxiv.org/abs/2512.03307v1
- http://arxiv.org/abs/2207.03208v2