

Comparative Performance of Tabular Foundation Models Pretrained on Synthetic Versus Real-World Imbalanced Data on OpenML-CC18

Assignee Research

June 11, 2026

Abstract

The development of tabular foundation models (TFMs) has accelerated in recent years, showing strong potential to outperform traditional ML methods for structured data. A key finding is that TFMs can be pretrained entirely on synthetic datasets, opening opportunities to design data generators that encourage desirable model properties. Prior work has mainly focused on crafting high-quality priors over generators to improve overall pretraining performance. Our insight is that parameterizing the generator distribution enables an adversarial robustness perspective: during training, we can adapt the

1 Introduction

This paper examines: Robust Tabular Foundation Models. Research question: How does the performance of tabular foundation models pretrained on synthetic data compare to those pretrained on real-world imbalanced datasets when evaluated on OpenML-CC18 benchmark tasks using metrics like accuracy, F1-score, and AUC-ROC?.

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 7.3/10.

3 Results

12 papers retrieved. 12 claims extracted; 9 independently verified. Quality review score: 7.3/10.

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 structure	✓	0.20
TFMs can produce high-quality predictions on new datasets in milliseconds when GPU-accelerated.	✓	0.18
Training TFMs relies on generating large amounts of diverse synthetic datasets constructed from structural causal models	✓	0.20
All current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribution	✓	0.20
Fixed priors in TFM training underrepresent certain regions of the parameter space, potentially degrading performance on	✓	0.22
State-of-the-art TFMs lag behind tree-based methods on some benchmarks.	×	0.14
The proposed method formalizes adversarial training over the SCM parameter space to adapt to challenging regions such as	✓	0.21
The proposed algorithm is named ROBUST TABULAR FOUNDATION MODELS (RTFM).	×	0.12
Applying RTFM to TabPFN V2 with only 90k additional training datasets significantly improves its ranking on several real	✓	0.19
The maximization stage of the proposed algorithm uses a black-box optimization algorithm to search the parameter space	✓	0.26
In the described implementation, the estimated optimality gap could be computed in a matter of seconds when parallelized	✓	0.17
Table (p10) lists synthetic dataset generation parameters including feature counts ranging from 3 to 128 and activation	×	0.05

References

- <http://arxiv.org/abs/2409.05215v1>
- <http://arxiv.org/abs/2512.03307v1>

- <http://arxiv.org/abs/2502.17119v2>