

Calibration Performance of Synthetic-Pretrained Multimodal Tabular Foundation Models on Perturbed OOD TabBench Subsets

Assignee Research

June 12, 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: Do multimodal tabular foundation models pretrained on synthetic data retain calibration performance better than real-data-only models when evaluated on perturbed OOD subsets of TabBench?.

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 8.5/10.

3 Results

12 papers retrieved. 13 claims extracted; 11 independently verified. Quality review score: 8.5/10.

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 structured data.	✓	0.18
TFMs rely on in-context learning (ICL).	✓	0.16
TFMs can provide high-quality predictions on new datasets in milliseconds when GPU-accelerated.	✓	0.17
TFMs are pretrained using synthetic data generated from structural causal models (SCMs).	✓	0.20
Current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribution.	✓	0.20
Fixed priors underrepresent certain regions of the parameter space, potentially degrading performance on real-world data.	✓	0.25
State-of-the-art TFMs still lag behind tree-based methods on some benchmarks.	×	0.13
The work introduces an optimality gap concept to target regions where the TFM underperforms relative to the best achievable model.	×	0.12
The work proposes an efficient, model-agnostic two-stage adversarial training algorithm for TFMs, called ROBUST TABULAR.	✓	0.20
Applying RTFM to TabPFN V2 with 90k additional training datasets significantly improves the ranking of TabPFN on several benchmarks.	✓	0.21
The maximization stage involves freezing gW to maximize the optimality gap in its current state.	✓	0.15
A black-box optimization algorithm is used to efficiently search the parameter space for parameters with large optimality gaps.	✓	0.29
For $n_{ds} = 20$ and $e = 7$, the estimated optimality gap δ_{θ_i} can be computed in a matter of seconds when parallelized, given $n_{ds} = 20$ and $e = 7$.	✓	0.30

References

- <http://arxiv.org/abs/2312.07577v3>
- <http://arxiv.org/abs/2512.03307v1>

- <http://arxiv.org/abs/2210.12369v1>