

CausalMixFT Robustness Across Tabular Domains Under Varying Synthetic Data Generator Diversity

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: What is the impact of varying the diversity of synthetic data generators on the robustness of CausalMixFT-pretrained models across different tabular domains, as measured by validation accuracy on out-of-distribution benchmarks?.

2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 7.7/10.

3 Results

4 papers retrieved. 12 claims extracted; 10 independently verified. Quality review score: 7.7/10.

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 TFMs rely on in-context learning (ICL).	✓	0.18
TFMs can provide high-quality predictions on new datasets in milliseconds when GPU-accelerated.	✓	0.17
Current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribution	✓	0.17
Fixed priors underrepresent certain regions of the parameter space, potentially degrading performance on real-world data	✓	0.19
State-of-the-art TFMs still lag behind tree-based methods on some benchmarks.	✓	0.25
Training TFMs relies on generating a large amount of diverse synthetic datasets.	×	0.14
The generation process relies on constructing structural causal models (SCMs) from which datasets can be sampled.	✓	0.22
The structure of these SCMs is implicitly parameterized, giving significant control over the data generation process.	✓	0.26
RTFM can significantly improve the ranking of TabPFN on several real-world tabular benchmarks with only 90k additional tokens	✓	0.30
The optimality gap can be estimated by sampling a fixed number of generators and datasets.	✓	0.23
The estimated optimality gap can be computed in a matter of seconds when parallelized, given sufficient CPU cores.	×	0.13
	✓	0.22

References

- <http://arxiv.org/abs/1710.05833v2>
- <http://arxiv.org/abs/2309.05856v1>
- <http://arxiv.org/abs/2512.03307v1>