

To what extent does the pretraining data split ratio between synthetic and real tabular data affect the downstream

Assignee Research

June 10, 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: To what extent does the pretraining data split ratio between synthetic and real tabular data affect the downstream robustness of TFMs on the TabMNAR benchmark under 50% random missing data perturbations?.

2 Methodology

Systematic literature search across multiple databases yielded 9 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.5/10.

3 Results

9 papers retrieved. 14 claims extracted; 1 independently verified. Quality review score: 3.5/10.

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.	✓	0.17
TFMs rely on in-context learning (ICL).	×	0.02
TFMs can provide high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.07
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.	×	0.04
State-of-the-art TFMs still lag behind tree-based methods on some benchmarks.	×	0.07
Training TFMs relies on generating a large amount of diverse synthetic datasets.	×	0.09
The generation process relies on constructing structural causal models (SCMs) from which datasets can be sampled.	×	0.03
The structure of these SCMs is implicitly parameterized, giving significant control over the data generation process.	×	0.06
RTFM is a two-stage adversarial training algorithm for TFMs.	×	0.12
RTFM can significantly improve the ranking of TabPFN on several real-world tabular benchmarks with only 90k additional tokens.	×	0.09
The optimality gap concept is used to target regions where the TFM underperforms relative to the best achievable performance.	×	0.10
The black-box optimization algorithm is used to efficiently search the space for parameters with large optimality gaps.	×	0.05
For $n_{\text{ds}} = 20$ and $e = 7$, the estimated optimality gap $b\delta\theta_i$ could be computed in a matter of seconds when parallelized, given	×	0.04

References

- <http://arxiv.org/abs/2510.21391v1>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2411.15497v3>