

Scaling Laws and Trade-offs in Pretraining Tabular Foundation Models on Synthetic Data

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: How does the computational efficiency (measured in FLOPs or training time) of pretraining tabular foundation models on synthetic data scale with dataset size, and what trade-offs exist between. 10 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.2/10. 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 the computational efficiency (measured in FLOPs or training time) of pretraining tabular foundation models on synthetic data scale with dataset size, and what trade-offs exist between efficiency and robustness on benchmarks like TabMNAR and Monash?.

2 Methodology

Systematic literature search across multiple databases yielded 8 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.2/10.

3 Results

8 papers retrieved. 10 claims extracted; 0 independently verified. Quality review score: 3.2/10.

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.	×	0.13
TFMs rely on in-context learning (ICL).	×	0.04
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).	×	0.15
Current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribution	×	0.05
Fixed priors underrepresent certain regions of the parameter space, potentially degrading performance on real-world data	×	0.06
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 achievable model.	×	0.10
The proposed RTFM algorithm can significantly improve the ranking of TabPFN on several real-world tabular benchmarks with structured data.	×	0.12
The optimality gap can be computed in a matter of seconds when parallelized, given sufficient CPU cores ($n_{\text{cores}} = n_{\text{tasks}} \cdot d$)	×	0.04

References

- <http://arxiv.org/abs/2504.20900v1>
- <http://arxiv.org/abs/2512.03307v1>

- <http://arxiv.org/abs/2207.03208v2>