Synthetic Data Pretrained Tabular Foundation Models vs Fine-Tuned Models on TabBench

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 9 peer-reviewed papers addressing the following research question: How do tabular foundation models pretrained on synthetic data compare in zero-shot performance on TabBench against fine-tuned models in terms of accuracy and inference latency. 11 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.5/10. 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 do tabular foundation models pretrained on synthetic data compare in zero-shot performance on TabBench against fine-tuned models in terms of accuracy and inference latency?.

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. 11 claims extracted; 0 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) rely on in-context learning (ICL) for classification and regression tasks with structur	×	0.10
GPU-accelerated tabular foundation models can generate high-quality predictions on new datasets in milliseconds.	×	0.12
Training TFMs relies on generating diverse synthetic datasets constructed from structural causal models (SCMs).	×	0.13
All current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribut	×	0.06
Fixed priors in TFM training underrepresent certain regions of the parameter space, potentially degrading performance on	×	0.05
State-of-the-art TFMs lag behind tree-based methods on some benchmarks.	×	0.05
The proposed RTFM method was applied to TabPFN V2.	×	0.11
Training TabPFN V2 with only 90k additional datasets using RTFM significantly improved its ranking on several real-world	×	0.11
The maximization stage of the proposed algorithm freezes the model weights (gW) to maximize the optimality gap.	×	0.04
The proposed method uses a black-box optimization algorithm to search the SCM parameter space for parameters with large	×	0.01
In the described implementation, the estimated optimality gap could be computed in seconds when parallelized across 140	×	0.03

References

- http://arxiv.org/abs/2504.20900v1
- http://arxiv.org/abs/2512.03307v1
- http://arxiv.org/abs/2402.01204v4