

Scaling Synthetic Pretraining Data for Few-Shot Tabular Foundation Models

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: Does increasing the scale of synthetic pretraining data improve few-shot adaptation rates of tabular foundation models on heterogeneous TabBench tasks. 12 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. 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: Does increasing the scale of synthetic pretraining data improve few-shot adaptation rates of tabular foundation models on heterogeneous TabBench tasks?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

14 papers retrieved. 12 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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) rely on in-context learning (ICL) for classification and regression tasks with structure	×	0.12
TFMs can produce high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.07
Training TFMs relies on generating diverse synthetic datasets constructed from structural causal models (SCMs).	×	0.07
All current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribution	×	0.06
Fixed priors in TFM training underrepresent certain regions of the parameter space, potentially degrading performance on	×	0.06
State-of-the-art TFMs lag behind tree-based methods on some benchmarks.	×	0.06
The proposed method, RTFM, was applied to TabPFN V2.	×	0.11
Training TabPFN V2 with RTFM using only 90k additional training datasets significantly improved its ranking on several r	×	0.11
The maximization stage of the proposed algorithm uses a black-box optimization algorithm to search the SCM parameter space	×	0.04
In the described implementation, estimating the optimality gap with $n_{ds}=20$ and $e=7$ takes a matter of seconds when parallel	×	0.03
The benchmark table includes synthetic dataset configurations with feature counts ranging from 3 to 128.	×	0.03
The benchmark table includes synthetic dataset configurations using activation functions such as tanh, identity, elu, an	×	0.03

References

- <http://arxiv.org/abs/2208.01009v2>
- <http://arxiv.org/abs/2311.10051v1>

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