

# Pretraining Tabular Foundation Models on Synthetic Data with Controlled Noise Levels Enhances Robustness to Distribution Shifts

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: To what extent does pretraining tabular foundation models on synthetic data with varied noise levels impact robustness against distribution shifts in real-world structured datasets. 11 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. 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: To what extent does pretraining tabular foundation models on synthetic data with varied noise levels impact robustness against distribution shifts in real-world structured datasets?.

## 2 Methodology

Systematic literature search across multiple databases yielded 10 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

10 papers retrieved. 11 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 structured	×	0.11
TFMs can produce high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.07
Training TFMs relies on generating large amounts of diverse synthetic datasets constructed from structural causal models	×	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	×	0.05
State-of-the-art TFMs still lag behind tree-based methods on some benchmarks.	×	0.06
The proposed method, RTFM, was applied to TabPFN V2.	×	0.10
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 RTFM algorithm freezes the model weights (gW) to maximize the optimality gap.	×	0.05
The RTFM methodology uses a black-box optimization algorithm to search the SCM parameter space for parameters with large	×	0.02
In the described implementation, the estimated optimality gap could be computed in a matter of seconds when parallelized	×	0.03

## References

- <http://arxiv.org/abs/2312.05435v1>
- <http://arxiv.org/abs/2206.03669v3>
- <http://arxiv.org/abs/2512.03307v1>