

# Foundation Model Robustness Under Synthetic vs. Real Data Training Regimes

Assignee Research

June 8, 2026

## Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: To what extent does training on generative AI-synthesized datasets degrade the robustness of foundation models against distributional shifts compared to real-data baselines. 13 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.7/10. 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 training on generative AI-synthesized datasets degrade the robustness of foundation models against distributional shifts compared to real-data baselines?.

## 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 2.7/10.

## 3 Results

14 papers retrieved. 13 claims extracted; 0 independently verified. Quality review score: 2.7/10.

## 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	×	0.13
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
State-of-the-art TFMs lag behind tree-based methods on some benchmarks.	×	0.05
The proposed method, RTFM, was applied to TabPFN V2.	×	0.10
Training RTFM on TabPFN V2 required only 90k additional training datasets.	×	0.09
Applying RTFM to TabPFN V2 significantly improved its ranking on several real-world tabular benchmarks.	×	0.06
The maximization stage of the proposed algorithm uses a black-box optimization algorithm to search the parameter space.	×	0.06
In the described implementation, the estimated optimality gap could be computed in a matter of seconds when parallelized	×	0.04
The implementation used ncores equal to the product of nds and e (ncores = nds $\cdot$ e) for parallelization.	×	0.01
The benchmark table includes synthetic dataset configurations with activation functions such as tanh, identity, elu, and	×	0.02
The benchmark table includes synthetic dataset configurations with noise distributions including uniform, exponential, a	×	0.02

## References

- <http://arxiv.org/abs/2410.02152v1>

- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2303.04707v2>