

Calibration Scaling of Tabular Foundation Models on Synthetic Adversarial Data

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does the calibration performance of tabular foundation models pretrained on synthetic adversarial datasets scale with increasing model size on the TabTime benchmark. 12 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.1/10. 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 calibration performance of tabular foundation models pretrained on synthetic adversarial datasets scale with increasing model size on the TabTime benchmark?.

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 4.1/10.

3 Results

14 papers retrieved. 12 claims extracted; 0 independently verified. Quality review score: 4.1/10.

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 TFMs rely on in-context learning (ICL).	×	0.14
TFMs can provide high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.03
Current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribution	×	0.07
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
RTFM is a two-stage adversarial training algorithm for TFMs.	×	0.05
RTFM can significantly improve the ranking of TabPFN on several real-world tabular benchmarks with only 90k additional tokens	×	0.12
The data generation process for TFMs relies on constructing structural causal models (SCMs).	×	0.12
The structure of SCMs is implicitly parameterized, giving significant control over the data generation process.	×	0.07
The optimality gap concept is used to target regions where the TFM underperforms relative to the best achievable performance	×	0.06
The estimated optimality gap can be computed in a matter of seconds when parallelized, given sufficient CPU cores.	×	0.10
	×	0.05

References

- <http://arxiv.org/abs/2606.00815v1>
- <http://arxiv.org/abs/2305.06090v1>
- <http://arxiv.org/abs/2512.03307v1>