Scaling Synthetic Data Diversity Enhances Robustness in Tabular Foundation Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: Does scaling synthetic data diversity improve the robustness of tabular foundation models against distribution shifts more effectively than increasing data volume alone. 7 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Causal Data Augmentation for Robust Fine-Tuning of Tabular Foundation Models. Research question: Does scaling synthetic data diversity improve the robustness of tabular foundation models against distribution shifts more effectively than increasing data volume alone?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.2/10.

3 Results

15 papers retrieved. 7 claims extracted; 1 independently verified. Quality review score: 5.2/10.

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
CausalMixFT achieves the highest median improvement of (+0.12$\pm$0.63) over the pretrained model on 33 classification data	×	0.09
Default fine-tuning has a variability of $\pm$0.98, while CausalMixFT has a variability of $\pm$0.63, indicating greater instabi	×	0.09
CausalMixFT ranks first overall in average ranks across datasets, followed by the default fine-tuning baseline, with pur	×	0.08
The normalization strategy used to compare performance across different data generators is based on the zero-shot perfor	×	0.05
CausalMixFT extends the fine-tuning framework by mixing real and causally grounded synthetic samples, using SCMs fitted	✓	0.19
SCM-Based Synthetic Augmentation (CausalMixFT) uses the PC and FCI algorithms to estimate structural relations between f	×	0.04
DAGs are sampled and fitted using DoWhy's SCM framework with additive noise models, with numerical features modeled with	×	0.04

References

- http://arxiv.org/abs/2601.04110v2
- http://arxiv.org/abs/2512.03307v1
- http://arxiv.org/abs/2312.07577v3