

Scaling Synthetic Pretraining Data Dimensionality for Tabular Foundation Model Robustness

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: Does scaling the dimensionality of synthetic pretraining data for tabular foundation models improve their robustness to covariate shift, as measured by accuracy on TabMWP benchmarks. 15 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: Causal Data Augmentation for Robust Fine-Tuning of Tabular Foundation Models. Research question: Does scaling the dimensionality of synthetic pretraining data for tabular foundation models improve their robustness to covariate shift, as measured by accuracy on TabMWP benchmarks?.

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. 15 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
Experiments were conducted on the Mitra model across 33 classification datasets with 10 folds each from the TabArena ben	×	0.07
The study totaled 2,310 fine-tuning runs.	×	0.12
Model performance is reported as normalized ROC-AUC relative to the pre-trained model.	×	0.07
CausalMixFT achieves a median improvement of $+0.12 \pm 0.63$ over the pre-trained model.	×	0.05
The default fine-tuning baseline achieves a median improvement of $+0.10 \pm 0.98$ over the pre-trained model.	×	0.09
Purely synthetic augmentation methods (CTGAN, SCM, TabEBM, TableAugment, and MixedModel) show negative median improvement	×	0.09
CausalMixFT has a variability of ± 0.63 , while default fine-tuning has a variability of ± 0.98 .	×	0.09
In average rank analysis, CausalMixFT ranks first overall, followed by the default fine-tuning baseline.	×	0.07
The normalization strategy uses the base model’s (Mitra’s) zero-shot performance as the baseline.	×	0.03
The normalization formula is defined as: $\text{score_normalized} = \text{metric_sign} \times (\text{score_method} / \text{score_baseline} - 1) \times 100\%$.	×	0.00
The method generates synthetic data using SCMs fitted to the target dataset.	×	0.14
Structural relations between features are estimated using the PC and FCI algorithms.	×	0.04
The estimation process produces a probabilistic adjacency matrix encoding edge strengths between variables.	×	0.05
DAGs are sampled and fitted using DoWhy’s SCM framework with additive noise models.	×	0.03
Numerical features are modeled with regressors and categorical features with classifiers within the SCM.	×	0.04

References

- <http://arxiv.org/abs/2601.04110v2>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2207.03208v2>