

Scaling SCM-Generated Synthetic Data for Causal Transfer in Tabular Foundation Models

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: Does scaling the size of the SCM-generated synthetic data in CausalMixFT lead to consistent performance improvements in TFMs across different tabular domains, as evaluated by metrics like mean. 14 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. 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 size of the SCM-generated synthetic data in CausalMixFT lead to consistent performance improvements in TFMs across different tabular domains, as evaluated by metrics like mean absolute error on benchmarks like TabFGSM?.

2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

13 papers retrieved. 14 claims extracted; 1 independently verified. Quality review score: 4.5/10.

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 pre-trained model on 33 classification data	×	0.09
Default fine-tuning has a variability of ± 0.98 , while CausalMixFT has a variability of ± 0.63 .	×	0.09
CausalMixFT ranks first overall in average ranks across datasets, followed by the default fine-tuning baseline.	×	0.08
Purely synthetic generators, including CTGAN, SCM, TabEBM, TableAugment, and Mixed-Model, show negative median improvement	×	0.08
The normalization strategy used to compare performance across different data generators is based on Gorishniy et al. [12]	×	0.04
The base model’s (Mitra’s) zero-shot performance is used as the performance baseline for normalization.	×	0.05
The normalized performance is computed as $\text{score}_{\text{normalized}} = \text{metricsign} \times (\text{score}_{\text{method}} / \text{score}_{\text{baseline}} - 1) \times 100\%$.	×	0.03
CausalMixFT extends the fine-tuning framework of Bhlér et al. [5] by mixing real and causally grounded synthetic sample	×	0.11
SCM-Based Synthetic Augmentation (CausalMixFT) uses SCMs fitted to the target dataset to generate synthetic data.	×	0.13
SCMs explicitly encode causal dependencies among features through a directed acyclic graph (DAG) and a set of structural	✓	0.15
The PC and FCI algorithms are used to estimate the structural relations between the features.	×	0.02
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 in the SCM framework.	×	0.04
Synthetic samples are generated by sampling exogenous noise and propagating it through the fitted SCM.	×	0.04

References

- <http://arxiv.org/abs/2601.04110v2>
- <http://arxiv.org/abs/2603.10254v1>
- <http://arxiv.org/abs/2504.20900v1>