

Causal Structure Integration in Synthetic Data Generation for Low-Resource Tabular Foundation Models

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: How does the integration of causal structure complexity in synthetic data generation compare to non-causal augmentation techniques in improving the cross-domain F1 scores of tabular foundation models. 17 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.4/10. 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: How does the integration of causal structure complexity in synthetic data generation compare to non-causal augmentation techniques in improving the cross-domain F1 scores of tabular foundation models on low-resource datasets such as OpenML?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.4/10.

3 Results

11 papers retrieved. 17 claims extracted; 1 independently verified. Quality review score: 4.4/10.

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 achieved a median improvement of $+0.12 \pm 0.63$ over the pre-trained model.	×	0.05
The default fine-tuning baseline achieved 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) showed negative median improvement	×	0.08
CausalMixFT has a performance variability of ± 0.63 , while default fine-tuning has a variability of ± 0.98 .	×	0.10
In average rank analysis across datasets, CausalMixFT ranks first overall.	×	0.03
The normalization strategy uses the base model's (Mitra's) zero-shot performance as the baseline.	×	0.05
The normalization formula is defined as: $\text{score_normalized} = \text{metric_sign} \times ((\text{score_method} / \text{score_baseline}) - 1) \times 100\%$.	×	0.01
The metric_sign is set to 1 for metrics where higher is better (e.g., ROC-AUC) and -1 for metrics where lower is better	×	0.03
The method generates synthetic data using Structural Causal Models (SCMs) fitted to the target dataset.	✓	0.22
Structural relations between features are estimated using the PC and FCI algorithms.	×	0.03
The estimation process produces a probabilistic adjacency matrix encoding edge strengths between variables.	×	0.01
DAGs are sampled and fitted using DoWhy's SCM framework with additive noise models.	×	0.04
Numerical features are modeled with regressors and categorical features with classifiers within the SCM.	×	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/2512.03307v1>
- <http://arxiv.org/abs/2507.03528v1>