

Causal Consistency in Synthetic Data and Robustness of Fine-Tuned Tabular Foundation Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 9 peer-reviewed papers addressing the following research question: To what extent does the causal consistency of synthetic data generated by SCMs versus CT-GAN/TVAE impact the robustness of fine-tuned tabular foundation models under distribution shifts, as measured. 13 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: To what extent does the causal consistency of synthetic data generated by SCMs versus CT-GAN/TVAE impact the robustness of fine-tuned tabular foundation models under distribution shifts, as measured by accuracy degradation on adversarial and out-of-distribution subsets of OpenML-CC18?.

2 Methodology

Systematic literature search across multiple databases yielded 9 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

9 papers retrieved. 13 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 .	×	0.08
CausalMixFT ranks first overall in average ranks across datasets, followed by the default fine-tuning baseline.	×	0.07
Purely synthetic generators, including CTGAN, SCM, TabEBM, TableAugment, and Mixed-Model, show negative median improvement	×	0.08
Early stopping based on limited validation data leads to significant validation set overfitting depending on the fine-tuning	✓	0.16
The normalization strategy suggested by Gorishniy et al. [12] is used to compare the performance across different data groups	×	0.04
CausalMixFT extends the fine-tuning framework of Böhler et al. [5] by mixing real and causally grounded synthetic samples	×	0.10
SCM-Based Synthetic Augmentation (CausalMixFT) uses SCMs fitted to the target dataset to generate synthetic data.	×	0.14
SCMs explicitly encode causal dependencies among features through a directed acyclic graph (DAG) and a set of structural equations	×	0.05
The PC and FCI algorithms are used to estimate the structural relations between the features.	×	0.03
DoWhy’s SCM framework with additive noise models is used to sample and fit DAGs.	×	0.02
Numerical features are modeled with regressors, and categorical features with classifiers in the SCM framework.	×	0.01
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/2512.21798v2>
- <http://arxiv.org/abs/2312.05435v1>