

Fine-Tuning Causal Synthetic Pretrained Tabular Foundation Models on the Shifted Tabular Dataset Benchmark

Assignee Research

June 11, 2026

Abstract

Fine-tuning tabular foundation models (TFMs) under data scarcity is challenging, as early stopping on even scarcer validation data often fails to capture true generalization performance. We propose CausalMixFT, a method that enhances fine-tuning robustness and downstream performance by generating structurally consistent synthetic samples using Structural Causal Models (SCMs) fitted on the target dataset. This approach augments limited real data with causally informed synthetic examples, preserving feature dependencies while expanding training diversity. Evaluated across 33 classification datasets

1 Introduction

This paper examines: Causal Data Augmentation for Robust Fine-Tuning of Tabular Foundation Models. Research question: What is the effect of fine-tuning tabular foundation models pretrained on causal synthetic data on their performance across different domains in the Shifted Tabular Dataset (STD) benchmark, evaluated using accuracy and calibration metrics?.

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 7.2/10.

3 Results

9 papers retrieved. 15 claims extracted; 11 independently verified. Quality review score: 7.2/10.

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
Experiments were conducted on the Mitra model across 33 classification datasets from the TabArena benchmark suite.	✓	0.18
The experimental setup involved 10 folds for each of the 33 datasets, totaling 2,310 fine-tuning runs.	×	0.12
Model performance was reported as normalized ROC-AUC relative to the pre-trained model.	✓	0.27
CausalMixFT achieved a median improvement of $+0.12 \pm 0.63$ over the pre-trained model.	✓	0.18
The default fine-tuning baseline achieved a median improvement of $+0.10 \pm 0.98$ over the pre-trained model.	✓	0.22
Purely synthetic augmentation methods (CTGAN, SCM, TabEBM, TableAugment, Mixed-Model) showed negative median improvements	✓	0.20
CausalMixFT demonstrated lower performance variability (± 0.63) compared to default fine-tuning (± 0.98).	×	0.13
In average rank analysis across datasets, CausalMixFT ranked first, followed by the default fine-tuning baseline.	✓	0.17
The normalization strategy used computes $\text{score_normalized} = \text{metricsign} \times (\text{score_method} / \text{score_baseline} - 1) \times 100\%$.	×	0.01
The normalization strategy suggested by Gorishniy et al. [12] was applied to compare performance across different data g	×	0.06
The method generates synthetic data using Structural Causal Models (SCMs) fitted to the target dataset.	✓	0.19
Structural relations between features were estimated using the PC and FCI algorithms.	✓	0.16
DAGs were sampled and fitted using DoWhy’s SCM framework with additive noise models.	✓	0.23
In the SCM framework, numerical features are modeled with regressors and categorical features with classifiers.	✓	0.19
Synthetic samples are generated by sampling exogenous noise and propagating it through the fitted SCM.	✓	0.22

References

- <http://arxiv.org/abs/2601.04110v2>
- <http://arxiv.org/abs/2504.20900v1>
- <http://arxiv.org/abs/2512.03307v1>