

Causal Data Augmentation for Fine-Tuning Tabular Foundation Models in Few-Shot Classification

Assignee Research

June 14, 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: How does the use of causal data augmentation techniques like CausalMixFT compare to traditional data augmentation methods (e.g., SMOTE, GAN-based) in improving the fine-tuning performance of tabular foundation models (TFMs) on few-shot classification tasks, measured by accuracy on benchmarks like TabMNAR or TabPFN?

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 8.5/10.

3 Results

10 papers retrieved. 10 claims extracted; 9 independently verified. Quality review score: 8.5/10.

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	✓	0.27
Default fine-tuning has a variability of ± 0.98 , while CausalMixFT has a variability of ± 0.63 , indicating greater instability	✓	0.18
CausalMixFT ranks first overall in average ranks across datasets, followed by the default fine-tuning baseline, with pure	✓	0.25
Early stopping based on limited validation data leads to significant validation set overfitting depending on the fine-tuning	✓	0.26
The normalization strategy used to compare performance across different data generators is based on the zero-shot performance	×	0.07
CausalMixFT extends the fine-tuning framework by mixing real and causally grounded synthetic samples into the fine-tuning	✓	0.22
SCM-Based Synthetic Augmentation (CausalMixFT) uses SCMs fitted to the target dataset to generate synthetic data that resemble	✓	0.20
The PC and FCI algorithms are used to estimate the structural relations between features in the SCM framework.	✓	0.18
DoWhy’s SCM framework with additive noise models is used to sample and fit DAGs for generating synthetic samples.	✓	0.17
Numerical features are modeled with regressors, and categorical features with classifiers in the SCM framework.	✓	0.19

References

- <http://arxiv.org/abs/2601.04110v2>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2511.08667v2>