

How does the integration of causal inference techniques into tabular foundation models affect their robustness

Assignee Research

June 10, 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 datas

1 Introduction

This paper examines: Causal Data Augmentation for Robust Fine-Tuning of Tabular Foundation Models. Research question: How does the integration of causal inference techniques into tabular foundation models affect their robustness to distribution shifts, as evaluated by accuracy degradation on adversarially perturbed datasets in benchmarks like CIFAR-10-Tab or Tabular Shift?.

2 Methodology

Systematic literature search across multiple databases yielded 8 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.8/10.

3 Results

8 papers retrieved. 13 claims extracted; 1 independently verified. Quality review score: 4.8/10.

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 in normalized ROC-AUC	×	0.11
Default fine-tuning has a variability of ± 0.98 in normalized ROC-AUC, while CausalMixFT has a variability of ± 0.63 .	×	0.10
CausalMixFT ranks first overall in average ranks across datasets, followed by the default fine-tuning baseline.	×	0.07
Purely synthetic generators (CTGAN, SCM, TabEBM, TableAugment, MixedModel) show negative median improvements in normalized	×	0.12
Early stopping based on limited validation data leads to significant validation set overfitting depending on the fine-tu	✓	0.16
The normalization strategy for performance comparison is based on the zero-shot performance of the pre-trained Mitra mod	×	0.04
CausalMixFT extends the fine-tuning framework by mixing real and causally grounded synthetic samples.	×	0.12
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	×	0.04
The PC and FCI algorithms are used to estimate the structural relations between the features in CausalMixFT.	×	0.03
DoWhy’s SCM framework with additive noise models is used to sample and fit DAGs in CausalMixFT.	×	0.02
Numerical features in CausalMixFT are modeled with regressors, and categorical features with classifiers.	×	0.04
Synthetic samples in CausalMixFT are generated by sampling exogenous noise and propagating it through the fitted SCM.	×	0.05

References

- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/1905.11374v5>
- <http://arxiv.org/abs/2601.04110v2>