

CausalMixFT Synthetic Dataset Size and Tabular Foundation Model Fine-Tuning Performance

Assignee Research

June 13, 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: What is the impact of the size of the synthetic dataset generated by CausalMixFT on the fine-tuning convergence speed and downstream performance of tabular foundation models, evaluated using validation accuracy across varying sizes of synthetic data?.

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 7.5/10.

3 Results

9 papers retrieved. 13 claims extracted; 10 independently verified. Quality review score: 7.5/10.

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	✓	0.26
Default fine-tuning has a variability of ± 0.98 , while CausalMixFT has a variability of ± 0.63 .	×	0.14
CausalMixFT ranks first overall in average ranks across datasets, followed by the default fine-tuning baseline.	✓	0.23
Purely synthetic generators, including CTGAN, SCM, TabEBM, TableAugment, and Mixed-Model, show negative median improvement	✓	0.22
Early stopping based on limited validation data leads to significant validation set overfitting depending on the fine-tuning	✓	0.26
The normalization strategy suggested by Gorishniy et al. [12] is used to compare the performance across different data groups	×	0.06
The base model’s (Mitra’s) zero-shot performance is used as the performance baseline to measure the improvement after fine-tuning	×	0.11
CausalMixFT extends the fine-tuning framework of Böhler et al. [5] by mixing real and causally grounded synthetic samples	✓	0.28
SCMs explicitly encode causal dependencies among features through a directed acyclic graph (DAG) and a set of structural equations	✓	0.27
The PC and FCI algorithms are used to estimate the structural relations between the features, producing a probabilistic DAG	✓	0.26
DAGs are sampled and fitted using DoWhy’s SCM framework with additive noise models.	✓	0.24
Numerical features are modeled with regressors, and categorical features with classifiers in the SCM framework.	✓	0.20
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/2410.02152v1>
- <http://arxiv.org/abs/2512.03307v1>