

# CausalMixFT’s Generalization Performance on TABFACT Under Data Scarcity

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 CausalMixFT on the generalization performance of tabular foundation models when evaluated on the TABFACT dataset under severe data scarcity?.

## 2 Methodology

Systematic literature search across multiple databases yielded 13 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

13 papers retrieved. 15 claims extracted; 10 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 with 10 folds each from the TabArena ben  | ✓        | 0.21       |
| The study totaled 2,310 fine-tuning runs.   | ×        | 0.09       |
| Model performance is reported as normalized ROC-AUC relative to the pre-trained model.  | ✓        | 0.27       |
| CausalMixFT achieves a median improvement of $+0.12 \pm 0.63$ over the pre-trained model.   | ✓        | 0.20       |
| The default fine-tuning baseline achieves a median improvement of $+0.10 \pm 0.98$ over the pre-trained model.  | ✓        | 0.24       |
| Purely synthetic augmentation methods (CTGAN, SCM, TabEBM, TableAugment, and MixedModel) show negative median improvement   | ✓        | 0.22       |
| CausalMixFT has a performance variability of $\pm 0.63$ , while default fine-tuning has a variability of $\pm 0.98$ .   | ×        | 0.15       |
| In average rank analysis across datasets, CausalMixFT ranks first overall.  | ×        | 0.14       |
| In average rank analysis, the default fine-tuning baseline ranks second, followed by purely synthetic generators.   | ✓        | 0.17       |
| The normalization strategy uses the base model’s (Mitra’s) zero-shot performance as the baseline.   | ×        | 0.06       |
| The normalization formula is defined as: $\text{score\_normalized} = \text{metricsign} \times (\text{score\_method} / \text{score\_baseline} - 1) \times 100\%$ . | ×        | 0.00       |
| The method generates synthetic data using SCMs fitted to the target dataset.  | ✓        | 0.15       |
| Structural relations between features are estimated using the PC and FCI algorithms.  | ✓        | 0.16       |
| 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 within the SCM.  | ✓        | 0.18       |

## References

- <http://arxiv.org/abs/2601.04110v2>
- <http://arxiv.org/abs/2411.15497v3>
- <http://arxiv.org/abs/2311.14544v1>