

# CausalMixFT Synthetic Data Augmentation and F1 Performance in Imbalanced Tabular Classification

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 9 peer-reviewed papers addressing the following research question: How does the F1-score performance of tabular foundation models trained on diffusion-generated synthetic data compare to those trained on real data in imbalanced classification tasks when evaluated. 10 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Causal Data Augmentation for Robust Fine-Tuning of Tabular Foundation Models. Research question: How does the F1-score performance of tabular foundation models trained on diffusion-generated synthetic data compare to those trained on real data in imbalanced classification tasks when evaluated across datasets with varying degrees of class skew?.

## 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 4.5/10.

## 3 Results

9 papers retrieved. 10 claims extracted; 1 independently verified. Quality review score: 4.5/10.

## 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	×	0.09
Default fine-tuning has a variability of $\pm 0.98$ , while CausalMixFT has a variability of $\pm 0.63$ , indicating greater instability	×	0.09
CausalMixFT ranks first overall in average ranks across datasets, followed by the default fine-tuning baseline, while pu	×	0.08
The normalization strategy used to compare performance across different data generators is based on the zero-shot perfor	×	0.06
CausalMixFT extends the fine-tuning framework by mixing real and causally grounded synthetic samples into the fine-tunin	×	0.14
SCM-Based Synthetic Augmentation (CausalMixFT) uses Structural Causal Models (SCMs) fitted to the target dataset to gene	✓	0.18
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.	×	0.02
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

## References

- <http://arxiv.org/abs/2601.04110v2>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2502.17119v2>