

Synthetic Pretraining Dimensionality and Sample Efficiency in Tabular Foundation Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: Does increasing feature dimensionality in synthetic pretraining data improve the sample efficiency of tabular foundation models when fine-tuned on low-resource benchmarks. 15 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.8/10. 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: Does increasing feature dimensionality in synthetic pretraining data improve the sample efficiency of tabular foundation models when fine-tuned on low-resource benchmarks?.

2 Methodology

Systematic literature search across multiple databases yielded 14 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

14 papers retrieved. 15 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 on 33 classification data	×	0.10
CausalMixFT outperforms the default fine-tuning baseline $(+0.10 \pm 0.98)$ and all purely synthetic augmentation methods, in	×	0.11
Default fine-tuning has a variability of ± 0.98 , while CausalMixFT has a variability of ± 0.63 , indicating greater instability	×	0.08
CausalMixFT ranks first overall in average ranks across datasets, followed by the default fine-tuning baseline, with pur	×	0.07
Early stopping based on limited validation data leads to significant validation set overfitting depending on the fine-tu	✓	0.16
The normalization strategy used to compare performance across different data generators is based on Gorishniy et al. [12]	×	0.06
The normalized performance is computed as $\text{score}_{\text{normalized}} = \text{metric}_{\text{sign}} \times (\text{score}_{\text{method}} / \text{score}_{\text{baseline}} - 1) \times 100\%$, where	×	0.04
CausalMixFT extends the fine-tuning framework of Bhlér et al. [5] by mixing real and causally grounded synthetic sample	×	0.11
SCM-Based Synthetic Augmentation (CausalMixFT) uses SCMs fitted to the target dataset, enabling the model to learn joint	×	0.11
SCMs explicitly encode causal dependencies among features through a directed acyclic graph (DAG) and a set of structural	×	0.05
The structural relations between the features are estimated using the PC and FCI algorithms, producing a probabilistic a	×	0.02
DAGs are sampled and fitted using DoWhy’s SCM framework with additive noise models.	×	0.03
Numerical features are modeled with regressors, and categorical features with classifiers in the SCM framework.	×	0.04
The complexity of the internally used model types can be controlled through a quality hyperparameter.	×	0.03
Synthetic samples are generated by sampling exogenous noise and propagating it through the fitted SCM.	×	0.04

References

- <http://arxiv.org/abs/2601.04110v2>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2602.09439v1>