

Diffusion-Based Synthetic Tabular Data Fidelity and Classifier Performance on Imbalanced OpenML-CC18 Tasks

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does the fidelity of synthetic tabular data generated by diffusion models correlate with downstream classifier performance on imbalanced OpenML-CC18 tasks as a function of training sample size. 13 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Robust Tabular Foundation Models. Research question: How does the fidelity of synthetic tabular data generated by diffusion models correlate with downstream classifier performance on imbalanced OpenML-CC18 tasks as a function of training sample size?.

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 3.5/10.

3 Results

14 papers retrieved. 13 claims extracted; 1 independently verified. Quality review score: 3.5/10.

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
Tabular foundation models (TFMs) rely on in-context learning (ICL) for classification and regression tasks with structural	×	0.13
TFMs can produce high-quality predictions on new datasets in milliseconds when GPU-accelerated.	×	0.07
Training TFMs relies on generating diverse synthetic datasets constructed from structural causal models (SCMs).	×	0.07
All current publicly available, competitive TFMs have been pretrained on datasets generated from a fixed prior distribution	×	0.06
Fixed priors in TFM training underrepresent certain regions of the parameter space, potentially degrading performance on	×	0.06
State-of-the-art TFMs lag behind tree-based methods on some benchmarks.	×	0.06
The proposed method formalizes adversarial training over the SCM parameter space to adapt to challenging regions such as	×	0.06
The proposed algorithm is named ROBUST TABULAR FOUNDATION MODELS (RTFM).	✓	0.20
Applying RTFM to TabPFN V2 with only 90k additional training datasets significantly improves its ranking on several real	×	0.10
The maximization stage of the methodology uses a black-box optimization algorithm to search the parameter space for para	×	0.02
In the described implementation, estimating the optimality gap with $nds=20$ and $e=7$ can be computed in a matter of second	×	0.03
The benchmark table includes synthetic dataset configurations with activation functions such as tanh, identity, elu, and	×	0.02
The benchmark table includes synthetic dataset configurations with noise distributions including uniform, exponential, a	×	0.03

References

- <http://arxiv.org/abs/2411.15497v3>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2308.02966v1>