

CausalMixFT vs. Diffusion-Based Augmentation for Few-Shot Learning in High-Cardinality Tabular Data

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: How does CausalMixFT’s structural consistency preservation compare to diffusion-based tabular augmentation in improving few-shot accuracy on OpenML-CC18 datasets with high cardinality categorical. 8 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Diffusion and Flow Matching Models for Tabular Data: A Survey. Research question: How does CausalMixFT’s structural consistency preservation compare to diffusion-based tabular augmentation in improving few-shot accuracy on OpenML-CC18 datasets with high cardinality categorical features?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.0/10.

3 Results

11 papers retrieved. 8 claims extracted; 0 independently verified. Quality review score: 3.0/10.

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
Future benchmarks should include missing-data settings, imbalanced classes, high-cardinality categorical features, relat	×	0.06
Formal differential privacy has begun to appear in tabular diffusion models [116], [117].	×	0.06
Recent memorization studies [118], [119] indicate that a small subset of records may dominate memorized generations.	×	0.01
Anomaly detection methods such as TCCM [71] show that feature-level residuals can support interpretability.	×	0.05
Healthcare models such as FlexGen-EHR [85] and PatientFlow [41] point toward multimodal and longitudinal tabular generat	×	0.05
Several recent methods combine diffusion or flow matching with autoencoders, transformers, tree models, or feature-token	×	0.12
Data augmentation for tabular data can be divided into two different tasks: 1) data synthesis and 2) over-sampling.	×	0.12
Over-sampling can be considered as a special case of single table synthesis where we only generate a part of the table.	×	0.07

References

- <http://arxiv.org/abs/2407.21523v1>
- <http://arxiv.org/abs/2206.00092v1>
- <http://arxiv.org/abs/2502.17119v2>