CausalMixFT vs. GAN and Diffusion-Based Augmentation for Tabular Foundation Model Generalization

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: How does CausalMixFT compare to other data augmentation techniques (e.g., GAN-based or diffusion-based synthetic data) in terms of improving the generalization of tabular foundation models on. 14 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Navigating Simply, Aligning Deeply: Winning Solutions for Mouse vs. AI 2025. Research question: How does CausalMixFT compare to other data augmentation techniques (e.g., GAN-based or diffusion-based synthetic data) in terms of improving the generalization of tabular foundation models on cross-domain tasks, as evaluated by F1 scores in low-resource settings?.

2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.3/10.

3 Results

4 papers retrieved. 14 claims extracted; 0 independently verified. Quality review score: 3.3/10.

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
The final model combining SimpleCNN, GLU, and observation normalization achieved a final score of 95.4% in Track 1.	×	0.10
The IMPALA ResNet baseline with 4 layers achieved a final score of 87.7% in Track 1.	×	0.05
The 24-layer ResNet achieved a final score of 65.98% in Track 1.	×	0.05
The 24-layer ResNet exhibited a 30 percentage point gap between ASR (80.96%) and MSR (51.00%).	×	0.01
The SimpleCNN model maintained a 2.8 percentage point gap between ASR (96.80%) and MSR (94.00%).	×	0.01
Adding data augmentation to ResNet decreased performance in Track 1.	×	0.03
All deep models in Track 2 used the ResNet + GLU architecture with 17.8M parameters.	×	0.07
The Track 1 visual encoder consists of two convolutional layers.	×	0.08
The first convolutional layer in the Track 1 solution uses an 8x8 kernel with stride 4.	×	0.02
The first convolutional layer produces 16 feature channels from an 86x155x1 input.	×	0.02
The second convolutional layer in the Track 1 solution uses a 4x4 kernel with stride 2.	×	0.02
The second convolutional layer expands to 32 feature channels.	×	0.01
Both convolutional layers in the Track 1 solution employ LeakyReLU activation with a negative slope of 0.2.	×	0.03
The feature maps from the visual encoder are flattened and projected to 256 dimensions via a fully-connected layer.	×	0.03

References

- http://arxiv.org/abs/2602.00982v1
- http://arxiv.org/abs/2504.07569v2
- http://arxiv.org/abs/2105.10886v1