

# Synthetic Data Generation Techniques for Enhancing Foundation Model Robustness in TabML

Assignee Research

June 8, 2026

## Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How do different synthetic data generation techniques (e.g., GANs, VAEs, diffusion models) compare in their ability to improve foundation model robustness when mixed with real data, as evaluated by. 14 claims were extracted from source literature; 2 were 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: A survey of synthetic data augmentation methods in computer vision. Research question: How do different synthetic data generation techniques (e.g., GANs, VAEs, diffusion models) compare in their ability to improve foundation model robustness when mixed with real data, as evaluated by TabML benchmark scores?.

## 2 Methodology

Systematic literature search across multiple databases yielded 15 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

15 papers retrieved. 14 claims extracted; 2 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
The FlyingThings3D dataset has proven effective in training deep learning models for optical flow and scene flow tasks.	×	0.04
Neural rendering aims to realize the scene rendering process using deep learning models.	×	0.06
Unlike traditional scene rendering based on 3D graphical modeling, the neural rendering process can be accomplished in b	×	0.06
In the forward direction of neural rendering, 2D images are generated from 3D scenes and additional scene parameters.	×	0.04
In the backward direction of neural rendering, the pixel image is translated into a realistic 3D scene.	×	0.08
The rendering process is inherently non-differentiable.	×	0.02
The non-differentiable nature of the rendering process severely constrains its incorporation in deep neural networks.	×	0.04
Point cloud representation has low memory requirements but low accuracy of scene topology information.	×	0.02
Voxel representation is more accurate with less processing and offers simplicity, but has a high memory footprint.	×	0.01
Mesh representation provides more grounding (physics-aware scene representation) but has high computational cost and dif	×	0.04
Multimodal representation offers high resolution and is more robust to visual artifacts, but is more complex and has hig	×	0.02
Implicit (NN) representation is naturally differentiable and has low memory requirements, but lacks grounding.	×	0.01
Data augmentation is a method used to mitigate the challenge of obtaining sufficient image data for target tasks in comp	✓	0.23
Synthetic data augmentation involves synthesizing training data from scratch when data for the target domain is not acce	✓	0.24

## References

- <http://arxiv.org/abs/2411.15497v3>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2403.10075v2>