

Diffusion Models vs. VAEs and GANs in Scalable Tabular Data Generation

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: How does the scalability of diffusion-based tabular data generation compare to VAEs and GANs in terms of training time and memory efficiency (measured in GFLOPs and peak GPU memory usage) when. 10 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: DogLayout: Denoising Diffusion GAN for Discrete and Continuous Layout Generation. Research question: How does the scalability of diffusion-based tabular data generation compare to VAEs and GANs in terms of training time and memory efficiency (measured in GFLOPs and peak GPU memory usage) when applied to large-scale tabular datasets like Criteo or Amazon Reviews?.

2 Methodology

Systematic literature search across multiple databases yielded 8 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.8/10.

3 Results

8 papers retrieved. 10 claims extracted; 0 independently verified. Quality review score: 3.8/10.

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
The Rico dataset contains 72K mobile app screens collected from over 97K Android apps with 25 label types.	×	0.01
The PubLayNet dataset contains 330K document layout annotations sourced from PubMed Central with 5 label types.	×	0.02
The dataset split of training, validation and test is 35,851/2,109/4,218 for Rico and 315,757/16,619/11,142 for PubLayNe	×	0.02
The evaluation metrics include Frchet Inception Distance (FID), Alignment, Maximum Intersection over Union (IoU), Overl	×	0.05
The research explores multiple tasks in layout generation: Conditional Generation by Class ($C \rightarrow S+P$), Conditional Generati	×	0.07
DogLayout is compared against models such as Layout-VAE, NDN-none, LayoutGAN++, BLT, MaskGIT, and LayoutDM.	×	0.07
The transformer-encoder in the generator and decoder has 4 layers, 8 attention heads, and a feed-forward network width o	×	0.01
Using LayoutGAN++ directly to model discrete label data results in gradient vanishing.	×	0.07
A layout l with M elements is defined as $\{(c_1, b_1), \dots, (c_M, b_M)\}$, where (c_i, b_i) represents the i -th elements in l ,	×	0.05
A standard diffusion model contains a forward diffusion process and reverse diffusion process.	×	0.09

References

- <http://arxiv.org/abs/2412.00381v1>

- <http://arxiv.org/abs/2504.20900v1>
- <http://arxiv.org/abs/2502.17119v2>