

To what extent does replacing GAN-based layout priors with denoising diffusion processes improve the alignment

Assignee Research

June 10, 2026

Abstract

Layout Generation aims to synthesize plausible arrangements from given elements. Currently, the predominant methods in layout generation are Generative Adversarial Networks (GANs) and diffusion models, each presenting its own set of challenges. GANs typically struggle with handling discrete data due to their requirement for differentiable generated samples and have historically circumvented the direct generation of discrete labels by treating them as fixed conditions. Conversely, diffusion-based models, despite achieving state-of-the-art performance across several metrics, require extensive sa

1 Introduction

This paper examines: DogLayout: Denoising Diffusion GAN for Discrete and Continuous Layout Generation. Research question: To what extent does replacing GAN-based layout priors with denoising diffusion processes improve the alignment stability of instruction-tuned models in visual grounding tasks?.

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 3.0/10.

3 Results

8 papers retrieved. 0 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.

References

- <http://arxiv.org/abs/2506.05199v3>
- <http://arxiv.org/abs/2402.11690v1>
- <http://arxiv.org/abs/2412.00381v1>