Floorplan Generation with Graph Beta Diffusion

Yusuke Takeuchi
Tetraz, Inc.
The University of Tokyo
Japan
y-takeuchi@robot.t.u-tokyo.ac.jp

Qi An
The University of Tokyo
Japan
anqi@robot.t.u-tokyo.ac.jp

Atsushi Yamashita
The University of Tokyo
Japan
yamashita@robot.t.u-tokyo.ac.jp

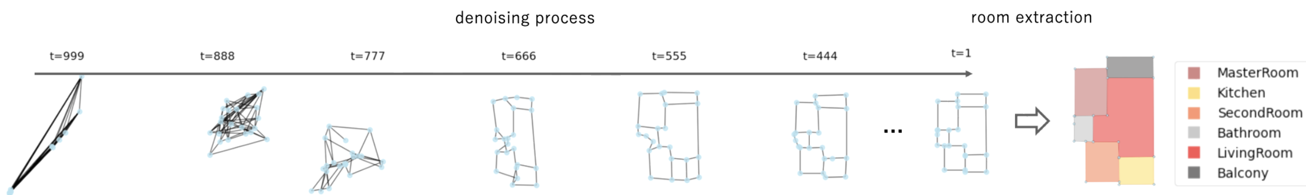

Figure 1: From a beta-noised graph to a structured floorplan, then rendered as layouts.

Abstract

We present a single-stage graph beta diffusion model for residential floorplan generation. In contrast to multi-stage pipelines, our formulation is end-to-end, avoiding error accumulation across modules and reducing the number of model parameters and hyperparameters. We further introduce an unsupervised Manhattan alignment loss that encourages axis-aligned walls. Our method substantially improves Fréchet Inception Distance over a parameter-efficient GAN baseline.

ACM Reference Format:
Yusuke Takeuchi, Qi An, and Atsushi Yamashita. 2025. Floorplan Generation with Graph Beta Diffusion. In *SIGGRAPH Asia 2025 Posters (SA Posters '25), December 15-18, 2025.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3757374.3771518

1 INTRODUCTION

Architectural floorplan design is labor-intensive, especially in the early stages that require exploring many layout alternatives. Generative models can automate this process, improving efficiency and providing diverse, well-structured options that enhance communication and client satisfaction.

Since the RPLAN dataset [Wu et al. 2019], deep learning has become central to data-driven floorplan generation. GSDiff [Hu et al. 2025] models layouts as graphs but uses a two-stage node–edge pipeline, where upstream errors propagate and multiple models enlarge both parameter count and hyperparameter space.

We propose a single-stage beta diffusion framework for structural graphs, jointly generating room types, connectivity, and geometry. Building on Graph Beta Diffusion (GBD) [Liu et al. 2024], our method simplifies the pipeline and introduces an unsupervised Manhattan alignment loss encouraging axis-aligned layouts. Our contributions are: (i) a single-model end-to-end framework, (ii) the first application of GBD to real-world structural data, and (iii) a reconstruction-free unsupervised Manhattan alignment loss that promotes axis-aligned geometry.

SA Posters '25, December 15-18, 2025, Hong Kong, Hong Kong

ACM ISBN 979-8-4007-2134-2/2025/12
https://doi.org/10.1145/3757374.3771518

2 METHOD

2.1 Floorplan Representation

A floorplan is modeled as an undirected graph $G = (V, A)$ with N junction nodes following [Hu et al. 2025]. Each node $v_i = (x_i, y_i, r_i)$ encodes normalized coordinates $(x_i, y_i) \in [0, 1]^2$ and a multi-hot room-type vector $r_i \in \{0, 1\}^C$ over C room categories. Wall segments are represented by a binary adjacency matrix $A \in \{0, 1\}^{N\times N}$ with $A_{ij} = 1$ if nodes i and j are connected by a wall.

2.2 Beta Diffusion

Let $G_0 = (V_0, A_0)$ denote the clean graph. The forward process applies element-wise Beta noise parameterized by G_0; its closed-form marginal is

$$q(G_t \mid G_0) = \mathrm{Beta}\big(\eta\, \alpha_t G_0,\ \eta(1 - \alpha_t G_0)\big), \tag{1}$$

where $\alpha_t \in (0, 1]$ decreases as t increases and $\eta > 0$ controls concentration (large η lowers variance).

The reverse process time-reverses (1) and learns per-step denoising transitions. The one-step reverse posterior is

$$q(G_{t-1} | G_t, G_0) = (1 - G_t)^{-1}\, \mathrm{Beta}(u_t;\, a_t, b_t)\,, \tag{2}$$

where

$$u_t = \frac{G_{t-1} - G_t}{1 - G_t}, \tag{3}$$

$$a_t = \eta(\alpha_{t-1} - \alpha_t)G_0, \qquad b_t = \eta\big(1 - \alpha_{t-1}G_0\big). \tag{4}$$

At inference we predict $\hat{G}_0 = f_\theta(G_t, t)$ and define

$$p_\theta(G_{t-1} | G_t) = q(G_{t-1} \mid G_t, \hat{G}_0). \tag{5}$$

2.3 Training Loss

We minimize the sum of a diffusion surrogate loss and the proposed Manhattan alignment loss:

$$\mathcal{L} = \mathcal{L}_{\text{GBD}} + \lambda_{\text{man}}\, \mathcal{L}_{\text{man}}^{\text{scheduled}}. \tag{6}$$

Manhattan alignment loss. To promote orthogonal geometry, we impose an unsupervised alignment loss that penalizes off-axis edges:

$$\mathcal{L}_{\text{man}} = \frac{1}{\sum_{i,j} \hat{E}_{ij}} \sum_{i,j} \hat{E}_{ij}\, \log\Big(1 + \frac{\min\{|x_i - x_j|^2, |y_i - y_j|^2\}}{\sigma^2}\Big), \tag{7}$$

where $\hat{E}_{ij} = \text{sigmoid}\big(\kappa(E_{ij} - \tau)\big)$ is a differentiable edge derived from the soft adjacency E. Here, τ is a confidence threshold, κ controls the steepness, and σ sets the tolerance for near-horizontal/vertical alignment. We apply noise-aware scheduling,

$$\mathcal{L}_{\text{man}}^{\text{scheduled}} = \sum_{t=2}^{T} \alpha_t^{\gamma}\, \mathcal{L}_{\text{man}}^{(t)}, \tag{8}$$

where $\mathcal{L}_{\text{man}}^{(t)}$ is computed on the step-t predictions, T is the total number of diffusion steps, and $\gamma > 0$ controls the schedule strength. This keeps the penalty small under high noise and larger at lower noise, encouraging axis-aligned walls when edges are reliable.

GBD surrogate loss. Following GBD [Liu et al. 2024], we use a two-term KL surrogate that upper-bounds the divergence between the learned reverse chain and the forward-defined target with trade-off weight ω.

$$\mathcal{L}_{\text{GBD}} = \sum_{t=2}^{T} \Big((1-\omega)\, \mathcal{L}_{\text{sampling}}(t) + \omega\, \mathcal{L}_{\text{correction}}(t) \Big), \tag{9}$$

$$\mathcal{L}_{\text{sampling}}(t) = \mathbb{E}_{q(G_t, G_0)} \left[\text{KL}\big(p_\theta(G_{t-1} \mid G_t) \,\|\, q(G_{t-1} \mid G_t, G_0)\big) \right], \tag{10}$$

$$\mathcal{L}_{\text{correction}}(t) = \mathbb{E}_{q(G_t, G_0)} \left[\text{KL}\big(q_t(\cdot \mid \hat{G}_0(G_t, t)) \,\|\, q_t(\cdot \mid G_0)\big) \right]. \tag{11}$$

3 EXPERIMENTS

3.1 Implementation and Data

We implement the model in PyTorch with a 6-layer Graph Transformer and a Beta-diffusion schedule of T=1000 steps. The forward Beta uses concentration η=10,000 for active nodes/edges and 10 for background. Training runs for 1M steps on a single RTX 3090 using AdamW with $lr = 10^{-3}$ and batch size is 128. Loss hyperparameters are λ_{man}=20, σ=0.01, κ=30, τ=0.9, γ=3, and ω=0.01. We evaluate on RPLAN [Wu et al. 2019] (~80k floorplans), following GSDiff for preprocessing to structural graphs and for the train/validation/test split.

3.2 Evaluation Metrics

We evaluate generation quality using the Fréchet Inception Distance (FID), which measures distributional similarity between the test set and generated floorplans. As baselines we include HouseGAN++ [Nauata et al. 2021], a parameter efficient GAN whose inference uses a single generator and thus matches our small model setting, and GSDiff [Hu et al. 2025]. We also report an ablation that removes the proposed unsupervised Manhattan alignment loss.

Table 1: Quantitative comparison on RPLAN.

Method	FID↓	#Param (M)
HouseGAN++ [Nauata et al. 2021]	48.40	2
GSDiff [Hu et al. 2025]	4.83	125
Ours (with Manhattan Loss)	19.49	5
Ours (without Manhattan Loss)	149.86	5

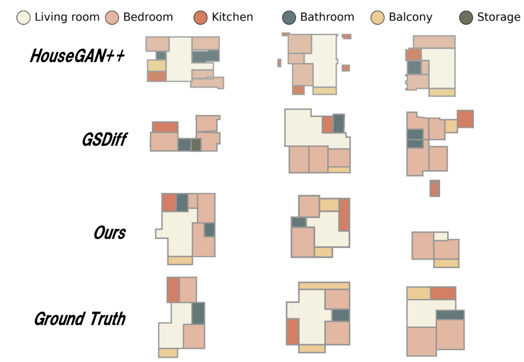

Figure 2: Qualitative comparison on RPLAN.

4 RESULTS

Table 1 summarizes FID and model size: our single-stage model improves FID over HouseGAN++ but does not yet match two-stage GSDiff; removing the Manhattan alignment term severely degrades FID, confirming its utility. Numbers for HouseGAN++ and GSDiff are taken from [Hu et al. 2025]; unlike HouseGAN++, our method does not require bubble-diagram conditioning.

In Fig. 2, HouseGAN++ examples are reproduced from [Hu et al. 2025], and GSDiff outputs are generated in our environment using the authors' released model; HouseGAN++ often yields non-straight walls, GSDiff is more orthogonal yet shows occasional failures, and ours is typically clean and close to ground truth with rare failures (e.g., the rightmost *Ours*).

5 CONCLUSION

We introduced a single-stage graph beta diffusion framework with an unsupervised Manhattan alignment loss, achieving clear FID gains over HouseGAN++ while retaining a simpler pipeline. We will scale model capacity and conditioning to probe the FID-capacity trade-off against two-stage pipelines such as GSDiff.

References

Sizhe Hu, Wenming Wu, Yuntao Wang, Benzhu Xu, and Liping Zheng. 2025. GSDiff: Synthesizing Vector Floorplans via Geometry-enhanced Structural Graph Generation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 39. AAAI Press, 17323–17332.

Xinyang Liu, Yilin He, Bo Chen, and Mingyuan Zhou. 2024. Advancing Graph Generation through Beta Diffusion. (2024). arXiv:2406.09357 [cs.LG]

Nelson Nauata, Sepidehsadat Hosseini, Kai-Hung Chang, Hang Chu, Chin-Yi Cheng, and Yasutaka Furukawa. 2021. HouseGAN++: Generative Adversarial Layout Refinement Network towards Intelligent Computational Agent for Professional Architects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 13632–13641.

Wenming Wu, Xiao Ming Fu, Rui Tang, Yuhan Wang, Yu Hao Qi, and Ligang Liu. 2019. Data-driven interior plan generation for residential buildings. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–13. https://doi.org/10.1145/3355089.3356556