

Class-Partitioned VQ-VAE and Latent Flow Matching for Point Cloud Scene Generation

Dasith de Silva Edirimuni, Ajmal Saeed Mian

The University of Western Australia
35 Stirling Highway,
Perth, WA 6009 Australia
{dasith.desilva, ajmal.mian}@uwa.edu.au

Abstract

Most 3D scene generation methods are limited to only generating object bounding box parameters while newer diffusion methods also generate class labels and latent features. Using object size or latent feature, they then retrieve objects from a predefined database. For complex scenes of varied, multi-categorical objects, diffusion-based latents cannot be effectively decoded by current autoencoders into the correct point cloud objects which agree with target classes. We introduce a Class-Partitioned Vector Quantized Variational Autoencoder (CPVQ-VAE) that is trained to effectively decode object latent features, by employing a pioneering *class-partitioned codebook* where codevectors are labeled by class. To address the problem of *codebook collapse*, we propose a *class-aware* running average update which reinitializes dead codevectors within each partition. During inference, object features and class labels, both generated by a Latent-space Flow Matching Model (LFMM) designed specifically for scene generation, are consumed by the CPVQ-VAE. The CPVQ-VAE’s class-aware inverse look-up then maps generated latents to codebook entries that are decoded to class-specific point cloud shapes. Thereby, we achieve pure point cloud generation without relying on an external objects database for retrieval. Extensive experiments reveal that our method reliably recovers plausible point cloud scenes, with up to 70.4% and 72.3% reduction in Chamfer and Point2Mesh errors on complex living room scenes.

Code and extended paper —

<https://github.com/ddsediri/CPVQ-VAE-LFMM>

1 Introduction

Architectural advancements to neural networks such as the U-Net (Ronneberger, Fischer, and Brox 2015) and Transformer (Vaswani et al. 2017) have driven rapid progress in text and image generation. This was accelerated by the introduction of ChatGPT (Radford et al. 2018) and similar variants for natural language processing while latent diffusion models such as StableDiffusion, inspired by Rombach et al. (2022), led the surge in image generation. By contrast, the maturity of 3D generation methods lags behind significantly and they fall under two main categories: 1) object generation and 2) scene generation. Of particular interest are point

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

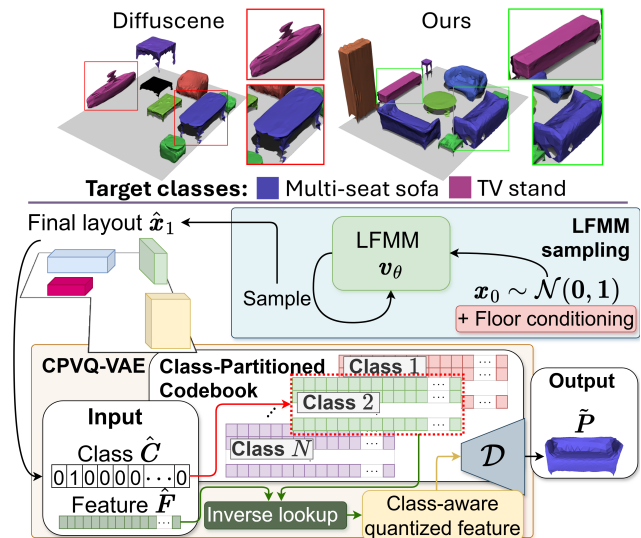


Figure 1: **Top:** Diffuscene, together with a pretrained VAE, generates object latents resulting in poorly decoded point clouds which *frequently disagree* with the object’s target class. Our Latent Flow Matching Model (LFMM) produces more robust latents which are *correctly decoded* by the Class-Partitioned VQ-VAE (CPVQ-VAE), that exploits class labels from the LFMM to look-up the corresponding codevectors. **Bottom:** Overview of our proposed point cloud scene generation framework. The LFMM generates object bounding box parameters, classes and respective features for the entire scene. Thereafter, the CPVQ-VAE decodes object features, given the generated object class and feature.

clouds as they are a compact representation of 3D data that can be efficiently processed compared to other representations such as voxel grids (Qi et al. 2017a,b). Contemporary object generation methods have mainly employed Continuous Normalizing Flows (Chen et al. 2018; Yang et al. 2019) Diffusion (Song and Ermon 2019; Luo and Hu 2021a) and Rectified Flows (Liu, Gong, and Liu 2023; Wu et al. 2023; Li et al. 2025) as the underlying generative mechanisms. Meanwhile, 3D scene generation methods seek to generate multiple objects within a single, plausible, scene. However, current point cloud scene generation methods fall short of

producing target point clouds. Instead, they produce the latent shape code of each object point cloud and use this latent to retrieve the corresponding object mesh from a database by minimizing the L_2 norm between generated features and database features, the latter must be filtered by class as part of a post-processing step.

Notably, Diffuscene (Tang et al. 2024) treats each scene as a stack of object vectors. Each vector is a concatenation of an object’s bounding box parameters (centroid translation, rotation, size), class vector and latent code which are generated via a diffusion model. After applying a post-processing step that filters the objects database by generated class, the latent feature is used to retrieve the closest object from the filtered database. This method can be extended to a pure point cloud generation method by directly decoding latents using the pre-trained variational autoencoder (VAE), which is the same used to produce target latents for training. However, we identify one key shortcoming: for complex scene types, such as living and dining rooms, object features generated directly by the diffusion model cannot be effectively decoded by the VAE to produce valid shapes. In fact, Diffuscene generates inexact latents which frequently disagree with the target classes, leading to incorrectly decoded shapes as shown in Fig. 1.

Conversely, methods such as ATISS (Paschalidou et al. 2021) and DeBaRa (Maillard et al. 2024) completely forgo generating object latents and utilize the regressed object sizes to filter the closest object from the database. Based on these observations, we surmise that shape latents generated from a diffusion process must follow the correct class conditioning while also producing a reliable enough latent that can be decoded to a sensible shape. Additionally, Diffuscene takes a comparatively large time to generate the scene layout. In contrast to point cloud methods, signed-distance field based methods require class conditioning and object-object relational information (Zhai et al. 2023, 2025; Wu et al. 2024). Moreover, the signed distance field estimation is performed on voxel grids which is less efficient (Ju et al. 2024). To address these limitations, we propose:

1. A *pure* point cloud scene generation method that simultaneously generates object bounding box parameters, class labels and latent features.
2. A novel *Class-Partitioned* VQ-VAE (CPVQ-VAE) with a labeled codebook where each vector belongs to a specific class. During inference, the CPVQ-VAE relies on labels, generated by latent space flow matching, when decoding codevectors into class-consistent point clouds.
3. A class-aware running average update which minimizes the number of *dead* codebook entries and addresses the problem of codebook collapse within the CPVQ-VAE.
4. A 3D latent space flow matching model (LFMM) that takes significantly fewer sampling steps to converge to plausible scene layouts, while generating object features which can successfully be decoded by the CPVQ-VAE.

Experiments show our method’s superiority in generating complex point cloud scenes, *with up to 70.4% and 72.3% reduction* in Chamfer and Point2Mesh distance errors, respectively, on the task of living room scene generation.

2 Related Works

Latent representation learning using VQ-VAEs: Hinton and Zemel (1993) were among the first to introduce autoencoders which were later exploited in tasks such as dimensionality reduction (Hinton and Salakhutdinov 2006). Variational autoencoders were then proposed (Kingma and Welling 2014; Kingma et al. 2014) that took advantage of variational inference and facilitated generative methods. Learning discrete representations of latent features using vector quantization was introduced by van den Oord, Vinyals, and Kavukcuoglu (2017) and extended to consider hierarchical codebook quantization (Razavi, van den Oord, and Vinyals 2019) with conditioned PixelCNN priors (van den Oord et al. 2016b). Furthermore, a categorical decoder was proposed for VQ-VAE architectures by Chorowski et al. (2019). While a powerful representation learning technique, naive vector quantization methods suffer from the problem of *codebook collapse* (Takida et al. 2022) which was addressed using techniques such as residual quantization (Lee et al. 2022), hierarchical quantization (Takida et al. 2024) and online clustering (Zheng and Vedaldi 2023). Specifically, the online clustering method proposes a simple framework for reinitializing dead codevectors via a running average update.

Generative processes: Score matching of Hyvärinen (2005), based on physical diffusion, underpins the revolution in image generation with foundational work being done by Song and Ermon (2019). Diffusion models were further refined by Ho, Jain, and Abbeel (2020) and Song, Meng, and Ermon (2021) who introduced the denoising diffusion probabilistic and implicit models (DDPM and DDIM), respectively, while Rombach et al. (2022) proposed to denoise images in latent space. Building on the work of Maoutsa, Reich, and Opper (2020), Song et al. (2021) obtained a deterministic process, named the *probability-flow* (PF-ODE) equation which facilitates the distillation of diffusion models into consistency models (Song et al. 2023) yielding straightened paths. Similar to consistency models, rectified flows were proposed by Liu, Gong, and Liu (2023) to regress optimal transport paths from noisy to clean samples by inferring a constant velocity field at each linearly interpolated state between samples. The stochastic interpolants of Albergo and Vanden-Eijnden (2023) and flow matching method of Lipman et al. (2023) are two related methods that subsume rectified flows where the velocity field at interpolated points between samples need not be constant.

Point cloud scene generation: Point cloud shape generation using score matching methods was first proposed by Cai et al. (2020) and further extended to use to DDPM sampling along with better shape conditioning by the work of Luo and Hu (2021a). However, these methods focus on single object generation and do not handle multi-categorical data. Score matching techniques have also been adopted for point cloud denoising (Luo and Hu 2021b; de Silva Edirimuni et al. 2023). Similarly, optimal transport based flow matching has been adopted for point cloud shape generation (Wu et al. 2023) and denoising (de Silva Edirimuni et al. 2024). Meanwhile, early scene generation methods focused

on autoregressive retrieval (Ritchie, Wang, and Lin 2019; Wang, Yeshwanth, and Nießner 2021; Paschalidou et al. 2021). Paschalidou et al. (2021), inspired by previous works (van den Oord et al. 2016a; Salimans et al. 2017), modeled object attributes with a mixture of logistic distributions. Thereafter, LEGO-Net denoised noisy arrangements of given objects to generate rearranged scenes (Wei et al. 2023), while Tang et al. (2024) introduced Diffuscene to generate all object parameters (including class labels and features) via diffusion. Diffuscene specifically uses these latents to improve the object retrieval process. Similarly, building on diffusion based methods, Maillard et al. (2024) proposed an EDM sampling driven method to retrieve scene layouts, which requires class conditioning to aid bounding box generation. Several other methods (Zhai et al. 2023, 2025; Wu et al. 2024) reverse diffuse Signed Distance Field representations of 3D data but require class labels, and object-to-object relations (in the form of scene graphs) as prior conditioning. Furthermore, Ju et al. (2024) proposed to reverse diffuse voxel occupancy to generate scenes. Finally, Feng et al. (2025) proposed CASAGPT which generates cuboid representations of 3D objects which are then used for retrieval.

3 Methodology

In this section, we present a Class-Partitioned Vector Quantized Autoencoder (CPVQ-VAE) and Latent space Flow Matching Model (LFMM) for point cloud scene generation. Our goal is to generate a plausible scene $\hat{\mathbf{X}} = \{\hat{\mathbf{x}}^m \mid 1 \leq m \leq M\}$, which consists of a maximum M objects, positioned and oriented consistently to their target classes. Denoised object bounding box attributes are characterized by a center coordinate translation $\hat{\mathbf{T}}^m \in \mathbb{R}^3$, a rotation $\hat{\mathbf{R}}^m = (\cos(\gamma), \sin(\gamma)) \in \mathbb{R}^2$, w.r.t. the scene’s vertical axis where γ is the relative angle, and the object size $\hat{\mathbf{S}}^m \in \mathbb{R}^3$. Moreover, the class attribute vector $\hat{\mathbf{C}}^m \in \mathbb{R}^{N_c}$, with N_c the number of classes for a given scene type, is used to obtain the object class $\hat{c}^m = \operatorname{argmax}_{n \in [N_c]} \hat{\mathbf{C}}_n^m$. Finally, latent features are represented by $\hat{\mathbf{F}}^m \in \mathbb{R}^{32}$. All attributes form a tuple that describes each object within the scene, $\hat{\mathbf{x}}^m = (\hat{\mathbf{T}}^m, \hat{\mathbf{R}}^m, \hat{\mathbf{S}}^m, \hat{\mathbf{C}}^m, \hat{\mathbf{F}}^m)$.

The object feature $\hat{\mathbf{F}}^m$ is subsequently decoded into a point cloud $\tilde{\mathbf{P}}^m$ using the CPVQ-VAE decoder \mathcal{D} and $\mathbf{L}(\hat{\mathbf{F}}^m, \mathcal{C}, \hat{c}^m)$, an inverse look-up function that we introduce to map $\hat{\mathbf{F}}^m$ to codebook vectors within the CPVQ-VAE codebook \mathcal{C} . The CPVQ-VAE architecture ensures that generated object point clouds are *class-consistent*, i.e., latent features are decoded to point clouds of the target class \hat{c}^m .

3.1 Class-Partitioned VQ-VAE (CPVQ-VAE)

Previous methods, such as Diffuscene, generate unreliable object latents that frequently produce incorrectly decoded point clouds that do not agree with the target class, or other object properties such as size. We propose a Class-Partitioned VQ-VAE (CPVQ-VAE) to ensure that object latents are mapped to labeled codebook entries in a consistent

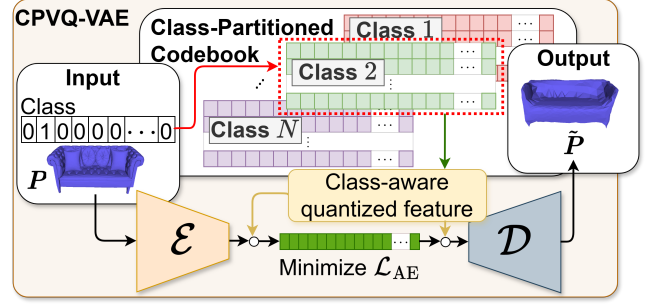


Figure 2: The Class-Partitioned VQ-VAE uses class label inputs to partition the codebook, allowing it to learn class-specific point cloud representations.

manner. The key idea is that object classes can be generated by a flow matching process, along with other object parameters such as the object feature. Thereafter, given the CPVQ-VAE’s partitioned codebook, the generated class and object feature can be decoded to produce the final 3D point cloud. The proposed CPVQ-VAE architecture is shown in Fig. 2.

A traditional VQ-VAE consists of an encoder \mathcal{E} and decoder \mathcal{D} along with a codebook $\mathcal{C} = \{(k, \mathbf{e}^k)\} \in \mathbb{R}^{N_K \times D_K}$ where k is a given code. Moreover, \mathbf{e}^k is the k -th codevector, $D_K = 128$ is the codevector dimension and N_K is the number of total codevectors. Given a point cloud $\mathbf{P} \in \mathbb{R}^{N_P \times 3}$ of N_P points, a VQ-VAE first encodes the point cloud into latent representation $\mathcal{E}(\mathbf{P}) \in \mathbb{R}^{D_K}$. The quantization process $\mathcal{Q}(\mathcal{E}(\mathbf{P}); \mathcal{C})$ maps this encoding to the nearest codevector in \mathcal{C} such that,

$$\mathcal{Q}(\mathcal{E}(\mathbf{P}); \mathcal{C}) = k^* = \operatorname{argmin}_{k \in [N_K]} \|\mathcal{E}(\mathbf{P}) - \mathbf{e}^k\|_2^2. \quad (1)$$

The quantized counterpart of $\mathcal{E}(\mathbf{P})$ is then $\mathbf{z}^q = \mathbf{e}^{k^*}$.

By contrast, the CPVQ-VAE partitions the codebook by class and the number of total codevectors becomes $N_K = N_c \times N_q$. Here, N_c is the number of classes and N_q is the number of codevectors assigned to each class. The quantization process is modified in the following way:

$$\mathcal{Q}(\mathcal{E}(\mathbf{P}); \mathcal{C}, c) = k_c^* = \operatorname{argmin}_{k \in [N_K]} \|\mathcal{E}(\mathbf{P}) - \mathbf{1}(c, k) \times \mathbf{e}^k\|_2^2, \quad (2)$$

where the indicator function $\mathbf{1}(c, k)$ is used to partition the codebook \mathcal{C} by object class $c \in [N_c]$. It is given by,

$$\mathbf{1}(c, k) := \begin{cases} 1 & \text{if } (c-1)N_q \leq k < cN_q, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The *class-aware* quantized feature is now $\mathbf{z}^{qc} = \mathbf{e}^{k_c^*}$. Thereby, we ensure that codevectors \mathbf{z}^{qc} belong to the appropriate class, c . The CPVQ-VAE is trained with the standard training objective:

$$\mathcal{L}_{\text{AE}} = \lambda_{\text{CD}} \mathcal{L}_{\text{CD}}(\mathbf{P}, \tilde{\mathbf{P}}) + \|\operatorname{sg}(\mathcal{E}(\mathbf{P})) - \mathbf{z}^{qc}\|_2^2 + \|\mathcal{E}(\mathbf{P}) - \operatorname{sg}(\mathbf{z}^{qc})\|_2^2, \quad (4)$$

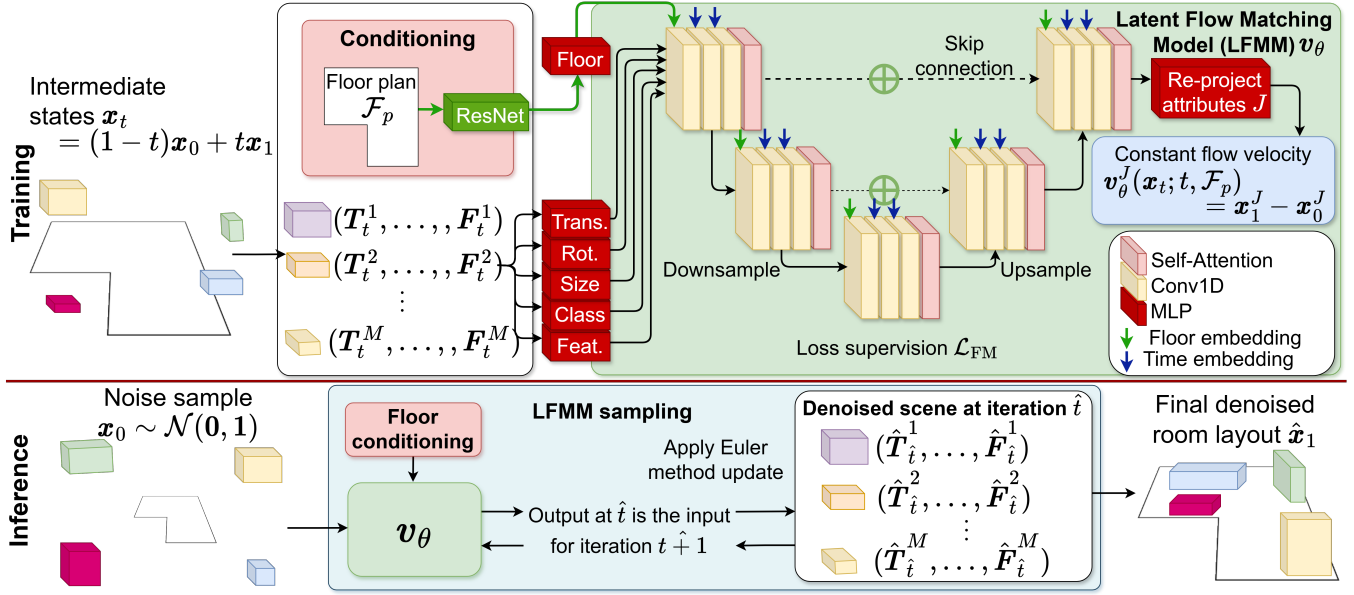


Figure 3: U-Net architecture of the Latent space Flow Matching Model.

where \mathcal{L}_{CD} is the reconstruction loss given by the Chamfer distance between the input, \mathbf{P} , and its reconstruction, $\tilde{\mathbf{P}}$. The weight $\lambda_{CD} = 10$ ensures all loss terms are of the same order and $\text{sg}(\cdot)$ represents the stop-gradient operator.

3.2 Class-Aware Running Average Update

The training objective in Eq. (4) is susceptible to codebook collapse, a common problem in VQ-VAE architectures. At the start of training, the network optimizes a few codevectors within the codebook while a large number of codevectors are ignored. Codevectors which are favored for optimization at earlier times during training remain *active* during training while those codevectors which do not receive substantial updates end up as *dead* codevectors. This further leads to a degradation in the VQ-VAE’s decoding ability as fewer viable codevectors result in fewer learned point cloud shapes.

To mitigate this issue, we reinitialize dead codevectors dynamically during training. Our approach is inspired by Zheng and Vedaldi (2023), however, we introduce a running average update strategy that reinitializes dead codevectors in a *class-aware* manner whereas Zheng and Vedaldi (2023) do not label codevectors by class and their proposed reinitialization is also class-agnostic. To achieve our dynamic, class-aware, reinitialization we first track average usage of codevectors, U_s^k , during each training step:

$$U_s^k = \gamma U_{s-1}^k + \frac{1-\gamma}{B} u_s^k, \quad (5)$$

where s is the training step, u_s^k is the number of encodings in the mini-batch which map to the codevector e^k , $U_0^k = 0$, B is the batch size, and $\gamma = 0.99$ is a decay hyperparameter. Next, we select the closest features $\mathcal{E}^{i_c^*}(\mathbf{P})$ in the mini-batch, to the codevectors e^k . These features become the an-

chors for the running average update, where we have,

$$i_c^* = \underset{i \in [B]}{\text{argmin}} \left\| \mathcal{E}^i(\mathbf{P}) - \mathbf{1}(c, k) \times e^k \right\|_2^2, \quad (6)$$

with the index i iterating over the entries of the mini-batch. Unlike previous works, we ensure that the method of anchor selection is *class-aware* by using the indicator function to partition the codebook by class. Now, the anchor used to update the codevector e^k is given by $\mathcal{E}^{i_c^*}(\mathbf{P})$.

Finally, we ensure that dead codevectors are reinitialized while active ones remain unchanged. To do so, we reinitialize each codevector e_s^k based on a decay value α_s^k . We calculate α_s^k and update the codevector in the following manner:

$$\alpha_s^k = \exp\left(-\frac{10U_s^k N_q}{1-\gamma} - \epsilon\right), \quad (7)$$

$$e_s^k = (1 - \alpha_s^k) e_{s-1}^k + \alpha_s^k \mathcal{E}^{i_c^*}(\mathbf{P}). \quad (8)$$

This *class-aware* running average update improves the CPVQ-VAE’s learning ability as unused codevectors are more frequently reinitialized, which allows them to be better optimized during training.

Given the trained CPVQ-VAE, we generate quantized latents z^{q_c} of point cloud objects which are used for training the latent flow model described in Sec. 3.3. As z^{q_c} is a 128 dimensional vector, we truncate it to only the first 32 entries to form the object latent such that $\mathbf{F}^m = z_{n < 32}^{q_c}$. We found that using the 32 dimensional truncated latent was more efficient for training.

3.3 Flow Matching on Optimal Transport Paths

Next, we introduce flow matching in latent space to evolve noisy inputs towards clean scene arrangements $\mathbf{X} = \{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^M, \mathcal{F}_p\}$ with M objects and floor plan \mathcal{F}_p . The object attribute vector, with dimensionality D_x , is

formed by concatenating all object parameters such that $\mathbf{x}^m = (\mathbf{T}^m; \mathbf{R}^m; \mathbf{S}^m; \mathbf{C}^m; \mathbf{F}^m) \in \mathbb{R}^{D_x}$. The sampling objective of generative methods can be expressed as a transport problem where samples from a noise distribution p_0 , given by $\mathbf{x}_0^m \sim p_0$, at time $t = 0$, are transported towards clean samples in the data distribution $\mathbf{x}_1^m \sim p_1(\mathbf{x}|\mathcal{F}_p)$ at time $t = 1$ where $p_1 = p_{\text{data}}$. For simplicity, we remove the superscript m signifying the m -th object of the scene as the flow matching process is identical for all objects and only keep the subscript t for time spent along the denoising trajectory.

In particular, we choose $p_0 = \mathcal{N}(\mathbf{0}, \mathbf{1})$ as a simple starting distribution. For an optimal transport-based flow matching model, we aim to infer constant paths between samples, leading to smaller discretization errors and faster sampling. This can be accomplished if each intermediate state \mathbf{x}_t at time t is a linear interpolation between \mathbf{x}_0 and \mathbf{x}_1 ,

$$\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1. \quad (9)$$

Thereby, we deduce the velocity $\mathbf{v}(\mathbf{x}_t, t)$ to be constant for the above linearly interpolated states:

$$\frac{d\mathbf{x}_t}{dt} = \mathbf{v}(\mathbf{x}_t; t, \mathcal{F}_p) = \mathbf{x}_1 - \mathbf{x}_0, \quad (10)$$

A constant flow velocity can be obtained by optimizing the following expected value:

$$\min_v \int_0^1 \mathbb{E}_{t \sim \mathcal{U}(0,1)} \left[\|\mathbf{v}(\mathbf{x}_t; t, \mathcal{F}_p) - (\mathbf{x}_1 - \mathbf{x}_0)\|_2^2 \right] dt, \quad (11)$$

This flow velocity can be approximated by a deep neural network \mathbf{v}_θ with network parameters θ . Such a network aims to infer a constant flow $\mathbf{v}_\theta(\mathbf{x}_t; t, \mathcal{F}_p) = \mathbf{x}_1 - \mathbf{x}_0$ at each interpolated state \mathbf{x}_t . We train this neural network using the following training objective:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t \sim \mathcal{U}(0,1)} \left[\lambda_J \sum_J \left(\|\mathbf{v}_\theta^J(\mathbf{x}_t; t, \mathcal{F}_p) - (\mathbf{x}_1^J - \mathbf{x}_0^J)\|_2^2 \right) \right], \quad (12)$$

where the λ_J is a weight that accounts for the loss contribution of each object parameter type, $J = T, R, S, C, F$, and $\mathbf{v}_\theta^J(\mathbf{x}_t; t, \mathcal{F}_p)$ are the entries of $\mathbf{v}_\theta(\mathbf{x}_t; t, \mathcal{F}_p)$ that correspond to the J -th parameter. Fig. 3 provides an overview of the our latent flow matching model.

3.4 Velocity Flow Sampling

A trained velocity flow model can be used to generate plausible scene layouts using the Euler method for sampling, as it provides a numerical solution to Eq. (10). Consequently, the denoised object parameters, at each iteration, are given by $\hat{\mathbf{x}}_{(\hat{t}+1)/N_{\hat{t}}} = \hat{\mathbf{x}}_{\hat{t}/N_{\hat{t}}} + \frac{1}{N_{\hat{t}}} \mathbf{v}_\theta(\hat{\mathbf{x}}_{\hat{t}/N_{\hat{t}}}; \hat{t}/N_{\hat{t}}, \mathcal{F}_p)$ where $\hat{t} = \{0, 1, 2, \dots, N_{\hat{t}}\}$ are integer time steps and $N_{\hat{t}} = 100$ is the total number of steps.

3.5 Class-Aware Inverse Look-up

The denoised feature $\hat{\mathbf{F}}$ is a 32 dimensional vector while codebook entries \mathbf{e}^k are 128 dimensional. We must look-up the correct codebook entry related to $\hat{\mathbf{F}}$ which also agrees

with the generated class $\hat{c} = \text{argmax}_{n \in [N_c]} \hat{\mathbf{C}}_n$ (where $\hat{\mathbf{C}}_n$ is the n -th element of $\hat{\mathbf{C}} \in \mathbb{R}^{N_c}$). Therefore, we propose the look-up function \mathbf{L} w.r.t truncated codevectors $\mathbf{e}_{n < 32}^k$ (consisting of the first 32 elements of \mathbf{e}^k),

$$\mathbf{L}(\hat{\mathbf{F}}, \mathcal{C}, \hat{c}) = k_{\hat{c}}^* = \text{argmax}_{k \in [N_K]} \left(\hat{\mathbf{F}} \cdot [\mathbf{1}(\hat{c}, k) \times \mathbf{e}_{n < 32}^k] \right), \quad (13)$$

that maximizes the cosine similarity between $\mathbf{e}_{n < 32}^k$ and denoised object features, in a class-aware manner.

Finally we recover the quantized feature $\mathbf{z}^{q\hat{c}} = \mathbf{e}^{k_{\hat{c}}^*}$, which is subsequently decoded by the CPVQ-VAE to give the generated point cloud $\tilde{\mathbf{P}} = \mathcal{D}(\mathbf{z}^{q\hat{c}})$.

4 Experiments

We follow the work of Paschalidou et al. (2021) and train and test our models on the 3D-FRONT dataset (Fu et al. 2021a,b) containing living, dining and bed rooms. Specifically, we consider the version of the 3D-FRONT used by Wei et al. (2023) and Maillard et al. (2024). The pre-processing of ATISS (Paschalidou et al. 2021) gives training and testing splits of 2338/587, 2071/516 and 5668/224 for living room, dining room and bedroom scenes, respectively. Additional details are provided in the *supplementary*.

Implementation Separate networks are trained on the living and dining room training sets for 100K epochs and on the bedroom training set for 40K epochs, respectively. Training and inference is performed on an NVIDIA RTX 4090 graphics card, using PyTorch for the implementation code.

Benchmark Comparisons We compare generation results with Diffuscene (Tang et al. 2024), and a variant of our LFMM which uses the same pre-trained VAE as Diffuscene. For retrieval, we additionally compare with ATISS (Paschalidou et al. 2021). We do not compare with DeBaRa (Maillard et al. 2024) or CASAGPT (Feng et al. 2025) as no code was available, and all generation methods are re-trained to ensure a fair comparison across all settings. For visualizing denoised point clouds, we employ alpha wrapping to reconstruct meshes (Edelsbrunner and Mücke 1994). For point cloud generation, decoded point clouds are evaluated w.r.t their counterparts retrieved from the predefined objects database. We treat these retrieved objects as the ground truth and compute Chamfer Distance ($\text{CD} \times 10^3$) and Point2Mesh Distance ($\text{P2M} \times 10^3$) metrics using PyTorch3D (Ravi et al. 2020). To evaluate retrieval results, Fréchet Inception Distance (FID), Kernel Inception Distance (KID $\times 10^3$), Scene Classification Accuracy and Categorical KL-Divergence ($\text{CKL} \times 10^2$) are used. The *supplementary document* contains more details regarding comparisons.

4.1 Point Cloud Generation

The top half of Table 1 and the *generation* columns of Fig. 4 indicate the performance of different methods on the point cloud generation task. We observe that Diffuscene regularly decodes incorrect point cloud objects which are not coherent in class labels and bounding box parameters on both

Generation Method	Living room				Dining room				Bedroom			
	CD ↓		P2M ↓		CD ↓		P2M ↓		CD ↓		P2M ↓	
Diffuscene	30.63		29.87		30.60		29.49		45.01		44.88	
Ours (LFMM + VAE)	<u>24.65</u>		<u>23.41</u>		<u>2.66</u>		<u>2.62</u>		<u>4.24</u>		<u>3.63</u>	
Ours (LFMM + CPVQ-VAE)	9.06		8.27		2.38		2.17		2.46		2.06	
Retrieval Method	FID ↓	KID ↓	SCA %	CKL ↓	FID ↓	KID ↓	SCA %	CKL ↓	FID ↓	KID ↓	SCA %	CKL ↓
ATISS	27.78	5.869	0.617	<u>0.560</u>	27.82	3.470	0.557	<u>0.633</u>	<u>23.02</u>	1.720	0.498	2.568
Diffuscene	27.39	6.169	0.615	0.455	30.23	5.697	0.562	0.546	25.13	2.459	0.606	6.184
Ours (LFMM + VAE)	24.46	4.020	0.608	1.063	<u>25.89</u>	<u>2.945</u>	<u>0.552</u>	1.309	21.86	0.723	0.545	3.095
Ours (LFMM + CPVQ-VAE)	<u>26.10</u>	<u>5.818</u>	<u>0.614</u>	1.129	25.58	2.708	0.484	1.176	21.34	0.794	<u>0.534</u>	<u>3.837</u>

Table 1: Comparison on point cloud generation and scene retrieval. Top results are in bold while second best are underlined. CD, P2M and KID results are multiplied by 10^3 and CKL by 10^2 . SCA scores closer to 0.5 are better

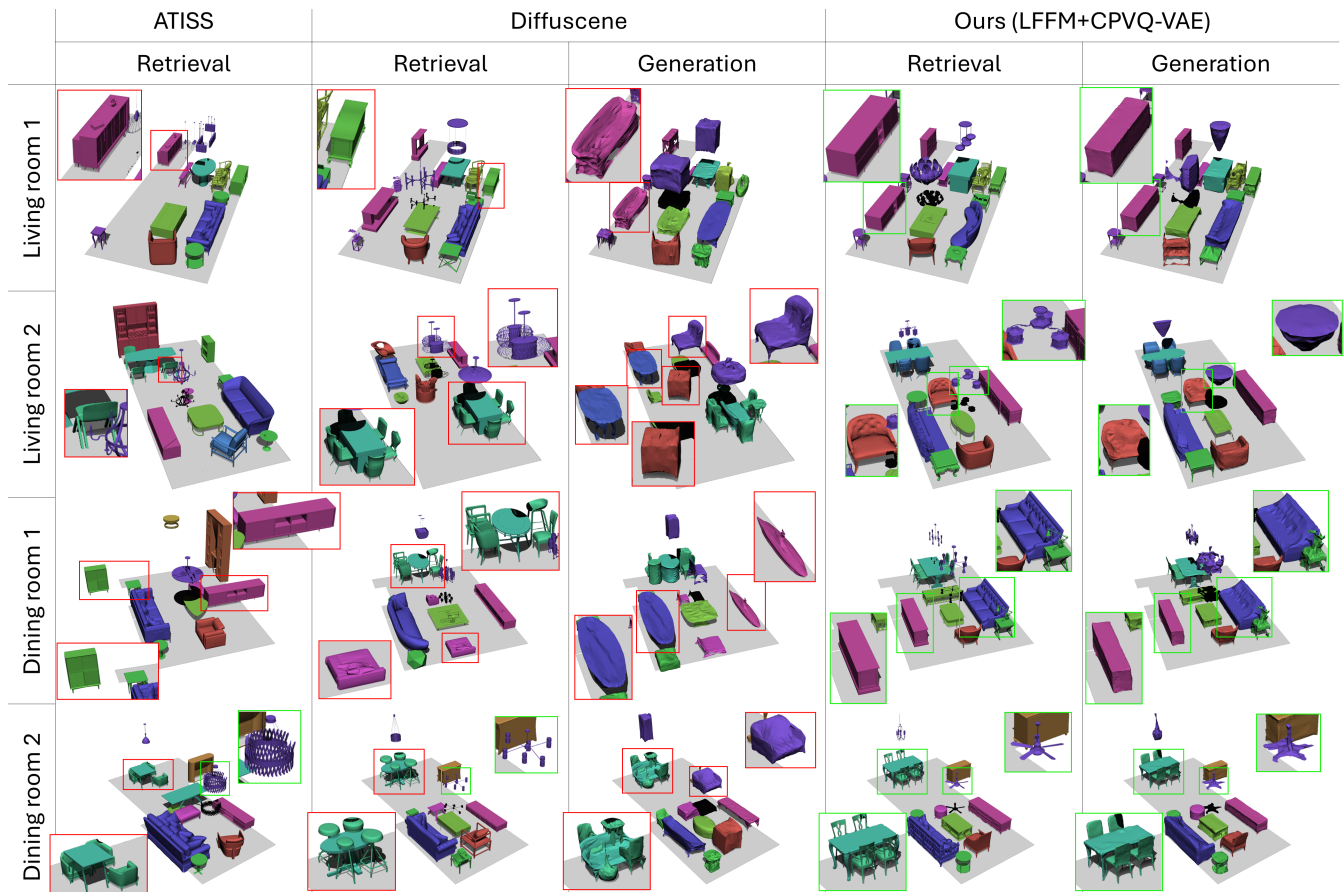


Figure 4: Visualization of retrieved and generated scenes. Green close-ups indicate correctly retrieved/decoded objects while incorrect results are in red. Notably, unlike the VAE decoder used by Diffuscene, our CPVQ-VAE decoder accurately decodes object features into class-consistent point clouds.

living and dining room sets (see red close-ups in Fig. 4). Table 1 shows that our LFMM with a standard VAE decoder suffers from the same drawback on more complex living room scenes while producing much better object features (w.r.t Diffuscene) for dining room scenes. By contrast, our LFMM together with the CPVQ-VAE, successfully de-

codes point cloud objects that agree with their target class labels across all scene types. Notably, on complex living room data, we achieve up to 70.4% and 72.3% reduction in CD and P2M errors, compared to Diffuscene. With respect to our LFMM+VAE variant, our full model showcases a 63.2% and 64.7% reduction in CD and P2M errors, respectively.

Method	Runtime (s) ↓	Avg. FID ↓	Avg. KID ↓
ATISS	0.024	<u>26.21</u>	<u>3.686</u>
Diffuscene	9.153	27.58	4.775
Ours	<u>0.892</u>	24.34	3.107

Table 2: Runtimes of different methods along with average FID and KID performance across all scene types.

4.2 Point Cloud Retrieval

The bottom half of Table 1 and the *retrieval* columns of Fig. 4 demonstrate each method’s retrieval performance. In general, we see that retrieval using flow matching is superior to scenes retrieved via diffusion (Diffuscene) or autoregression (ATISS). Although ATISS has a simpler training objective, which does not require the network to regress latent features, its performance on retrieval metrics is sub-optimal compared to our LFMM variants. Moreover, both ATISS, and Diffuscene, frequently generate incorrect bounding box parameters which lead to object collisions or object placements beyond the floor plan (see Fig. 4 for visualizations).

Runtime efficiency Table 2 provides runtime and average performance results. ATISS is the most efficient among all methods but remains sub-optimal in retrieval. Furthermore, among methods which generate the complete set of object attributes, our method has a clear advantage in efficiency. Concretely, our method is 90.3% faster than Diffuscene.

5 Ablations

To showcase the effectiveness of our proposed contributions, we investigate the following: (1) the role of autoencoder design in point cloud generation and (2) the impact of sampling iteration number on scene retrieval. For (1), we propose four autoencoder variants (V1 to V4) and (2), we consider our preferred CPVQ-VAE+LFMM model with different sampling iteration numbers $N_{\hat{t}}$.

Autoencoder design for point cloud generation. Table 3 shows our full CPVQ-VAE (V4) significantly improves upon the traditional VAE (V1), with a 41.8% and 43.1% reduction in CD and P2M error. On more complex scenes such as living rooms, the gap in CD and P2M results are wider as evidenced by table 1. We also see significant differences in CD and P2M results between a simple VQ-VAE (V2) and a class-partitioned VQ-VAE (V3), as class partitioning enforces class-consistent decoding of latents. However, both (V2) and (V3) suffer from the issue of codebook collapse and, in turn, perform worse than a standard VAE (V1). The full CPVQ-VAE (V4) reinforces the importance of partitioning the codebook by class *and* applying the class-aware running average update to reinitialize dead codevectors during training, which leads to impressive gains over (V1).

Sampling iteration number and scene retrieval. While the CPVQ-VAE+LFMM model infers plausible scene arrangements at low $N_{\hat{t}}$ (table 4), it produces sub-optimal performance. In fact, a higher number of sampling iterations is necessary to obtain consistent bounding box parameters to

Variant	Module				Bedroom	
	VAE	VQ-VAE	CP	RAU	CD ↓	P2M ↓
V1	✓	-	-	-	<u>4.24</u>	<u>3.63</u>
V2	-	✓	-	-	36.27	33.93
V3	-	✓	✓	-	5.00	4.17
V4	-	✓	✓	✓	2.46	2.06

Table 3: Ablation on autoencoder variants. CP denotes class partitioning and RAU is the running average update.

$N_{\hat{t}}$	Bedroom				
	FID ↓	KID ↓	SCA %	CKL ↓	Runtime (s) ↓
10	21.79	1.324	0.537	3.954	0.094
100	21.34	0.794	<u>0.534</u>	<u>3.837</u>	0.892
1000	<u>21.75</u>	<u>0.967</u>	0.517	2.411	8.963

Table 4: Ablation results on flow matching iterations $N_{\hat{t}}$ along with respective runtimes in seconds.

correctly position scene objects. However, this comes at the price of higher latency. We find $N_{\hat{t}} = 100$ is the optimal number of sampling iterations that balances performance and speed when generating object bounding box parameters.

6 Limitations and Future Work

Our method makes significant contributions to point cloud scene generation. However, current autoencoders cannot process dense point clouds and we limit ourselves to objects of $N_P = 2025$ points. Generating denser point clouds would yield better mesh reconstructions but requires improving the learning capacity of autoencoders. Furthermore, vector quantization introduces quantization error which marginally deteriorates the variety of *retrieved* objects. Therefore, scaling the CPVQ-VAE learning capacity and reducing quantization error are directions for future research.

7 Conclusion

In this paper, we introduced the first scene-level point cloud generation mechanism, capable of handling complex scenes of varied, multi-categorical objects. We introduced a class-partitioned vector quantized variational autoencoder (CPVQ-VAE) with labeled codebook entries that, given class labels and using a class-aware inverse look-up, successfully decodes object latents into the correct point cloud shapes which agree with target classes. We alleviate codebook collapse within the CPVQ-VAE via a class-aware running average update mechanism. Both object latents and class labels, along with object bounding box parameters, are generated by a latent space flow matching model designed specifically to handle scene-level generation. Extensive experiments across multiple room types reveal that our method reliably recovers point cloud scenes with consistent object properties, unlike previous methods, and promotes a new direction for future 3D scene generation.

Acknowledgments

This research was supported by the Australian Government through the Australian Research Council’s Discovery Projects funding scheme (project DP240101926).

References

- Albergo, M. S.; and Vanden-Eijnden, E. 2023. Building Normalizing Flows with Stochastic Interpolants. In *the 11th International Conference on Learning Representations (ICLR)*.
- Cai, R.; Yang, G.; Averbuch-Elor, H.; Hao, Z.; Belongie, S.; Snavely, N.; and Hariharan, B. 2020. Learning Gradient Fields for Shape Generation. In *Computer Vision – ECCV 2020*, 364–381.
- Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. 2018. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 6572–6583.
- Chorowski, J.; Weiss, R. J.; Bengio, S.; and van den Oord, A. 2019. Unsupervised Speech Representation Learning Using WaveNet Autoencoders. *IEEE ACM Trans. Audio Speech Lang. Process.*, 27(12): 2041–2053.
- de Silva Edirimuni, D.; Lu, X.; Li, G.; Wei, L.; Robles-Kelly, A.; and Li, H. 2024. StraightPCF: Straight Point Cloud Filtering. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 20721–20730.
- de Silva Edirimuni, D.; Lu, X.; Shao, Z.; Li, G.; Robles-Kelly, A.; and He, Y. 2023. IterativePFN: True Iterative Point Cloud Filtering. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 13530–13539.
- Edelsbrunner, H.; and Mücke, E. P. 1994. Three-dimensional alpha shapes. *ACM Trans. Graph.*, 13(1): 43–72.
- Feng, W.; Zhou, H.; Liao, J.; Cheng, L.; and Zhou, W. 2025. CASAGPT: Cuboid Arrangement and Scene Assembly for Interior Design. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 29173–29182.
- Fu, H.; Cai, B.; Gao, L.; Zhang, L.-X.; Wang, J.; Li, C.; Zeng, Q.; Sun, C.; Jia, R.; Zhao, B.; et al. 2021a. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10933–10942.
- Fu, H.; Jia, R.; Gao, L.; Gong, M.; Zhao, B.; Maybank, S.; and Tao, D. 2021b. 3d-future: 3d furniture shape with texture. *International Journal of Computer Vision*, 1–25.
- Hinton, G. E.; and Salakhutdinov, R. R. 2006. Reducing the Dimensionality of Data with Neural Networks. *Science*, 313(5786): 504–507.
- Hinton, G. E.; and Zemel, R. 1993. Autoencoders, Minimum Description Length and Helmholtz Free Energy. In *Advances in Neural Information Processing Systems (NIPS)*, volume 6.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Hyvärinen, A. 2005. Estimation of Non-Normalized Statistical Models by Score Matching. *Journal of Machine Learning Research*, 6: 695–709.
- Ju, X.; Huang, Z.; Li, Y.; Zhang, G.; Qiao, Y.; and Li, H. 2024. DiffInDScene: Diffusion-Based High-Quality 3D Indoor Scene Generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 4526–4535.
- Kingma, D. P.; Mohamed, S.; Rezende, D. J.; and Welling, M. 2014. Semi-supervised Learning with Deep Generative Models. In *Advances in Neural Information Processing Systems (NIPS)*, 3581–3589.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *the 2nd International Conference on Learning Representations (ICLR)*.
- Lee, D.; Kim, C.; Kim, S.; Cho, M.; and Han, W. 2022. Autoregressive Image Generation using Residual Quantization. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 11513–11522.
- Li, Y.; Zou, Z.; Liu, Z.; Wang, D.; Liang, Y.; Yu, Z.; Liu, X.; Guo, Y.; Liang, D.; Ouyang, W.; and Cao, Y. 2025. TripoSG: High-Fidelity 3D Shape Synthesis using Large-Scale Rectified Flow Models. *CoRR*, abs/2502.06608.
- Lipman, Y.; Chen, R. T. Q.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. In *the 11th International Conference on Learning Representations (ICLR)*.
- Liu, X.; Gong, C.; and Liu, Q. 2023. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. In *the 11th International Conference on Learning Representations (ICLR)*.
- Luo, S.; and Hu, W. 2021a. Diffusion Probabilistic Models for 3D Point Cloud Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2837–2845.
- Luo, S.; and Hu, W. 2021b. Score-Based Point Cloud Denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 4583–4592.
- Maillard, L.; Sereyjol-Garros, N.; Durand, T.; and Ovsjanikov, M. 2024. DeBaRA: Denoising-Based 3D Room Arrangement Generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, 109202–109232.
- Maoutsa, D.; Reich, S.; and Oppen, M. 2020. Interacting Particle Solutions of Fokker–Planck Equations Through Gradient–Log–Density Estimation. *Entropy*, 22.
- Paschalidou, D.; Kar, A.; Shugrina, M.; Kreis, K.; Geiger, A.; and Fidler, S. 2021. ATISS: autoregressive transformers for indoor scene synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Qi, C.; Su, H.; Mo, K.; and Guibas, L. 2017a. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 77–85.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30.

- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training. *OpenAI Blog*.
- Ravi, N.; Reizenstein, J.; Novotny, D.; Gordon, T.; Lo, W.-Y.; Johnson, J.; and Gkioxari, G. 2020. Accelerating 3D Deep Learning with PyTorch3D. *arXiv preprint arXiv:2007.08501*.
- Razavi, A.; van den Oord, A.; and Vinyals, O. 2019. Generating Diverse High-Fidelity Images with VQ-VAE-2. In *Advances in Neural Information Processing Systems (NeurIPS)*, 14837–14847.
- Ritchie, D.; Wang, K.; and Lin, Y. 2019. Fast and Flexible Indoor Scene Synthesis via Deep Convolutional Generative Models. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 6182–6190.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 10674–10685.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351, 234–241.
- Salimans, T.; Karpathy, A.; Chen, X.; and Kingma, D. P. 2017. PixelCNN++: Improving the PixelCNN with Discretized Logistic Mixture Likelihood and Other Modifications. In *the 5th International Conference on Learning Representations (ICLR)*.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *the 9th International Conference on Learning Representations (ICLR)*.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency models. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 32211–32252.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Takida, Y.; Ikemiya, Y.; Shibuya, T.; Shimada, K.; Choi, W.; Lai, C.; Murata, N.; Uesaka, T.; Uchida, K.; Liao, W.; and Mitsufuji, Y. 2024. HQ-VAE: Hierarchical Discrete Representation Learning with Variational Bayes. *Trans. Mach. Learn. Res.*, 2024.
- Takida, Y.; Shibuya, T.; Liao, W.; Lai, C.; Ohmura, J.; Uesaka, T.; Murata, N.; Takahashi, S.; Kumakura, T.; and Mitsufuji, Y. 2022. SQ-VAE: Variational Bayes on Discrete Representation with Self-annealed Stochastic Quantization. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, volume 162 of *Proceedings of Machine Learning Research*, 20987–21012.
- Tang, J.; Nie, Y.; Markhasin, L.; Dai, A.; Thies, J.; and Nießner, M. 2024. DiffuScene: Denoising Diffusion Models for Generative Indoor Scene Synthesis. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 20507–20518.
- van den Oord, A.; Dieleman, S.; Zen, H.; Simonyan, K.; Vinyals, O.; Graves, A.; Kalchbrenner, N.; Senior, A. W.; and Kavukcuoglu, K. 2016a. WaveNet: A Generative Model for Raw Audio. In *The 9th ISCA Speech Synthesis Workshop, (SSW)*, 125.
- van den Oord, A.; Kalchbrenner, N.; Espeholt, L.; Kavukcuoglu, K.; Vinyals, O.; and Graves, A. 2016b. Conditional Image Generation with PixelCNN Decoders. In *Advances in Neural Information Processing Systems (NIPS)*, 4790–4798.
- van den Oord, A.; Vinyals, O.; and Kavukcuoglu, K. 2017. Neural Discrete Representation Learning. In *Advances in Neural Information Processing Systems (NIPS)*, 6306–6315.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Łukasz Kaiser; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS)*, 6000–6010.
- Wang, X.; Yeshwanth, C.; and Nießner, M. 2021. SceneFormer: Indoor Scene Generation with Transformers. In *International Conference on 3D Vision, 3DV 2021, London, United Kingdom, December 1-3, 2021*, 106–115.
- Wei, Q. A.; Ding, S.; Park, J. J.; Sajjani, R.; Poulenard, A.; Sridhar, S.; and Guibas, L. 2023. LEGO-Net: Learning Regular Rearrangements of Objects in Rooms. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 19037–19047.
- Wu, L.; Wang, D.; Gong, C.; Liu, X.; Xiong, Y.; Ranjan, R.; Krishnamoorthi, R.; Chandra, V.; and Liu, Q. 2023. Fast Point Cloud Generation with Straight Flows. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 9445–9454.
- Wu, Z.; Feng, M.; Wang, Y.; Xie, H.; Dong, W.; Miao, B.; and Mian, A. 2024. External Knowledge Enhanced 3D Scene Generation from Sketch. In *Computer Vision – ECCV 2024*, 286–304.
- Yang, G.; Huang, X.; Hao, Z.; Liu, M.; Belongie, S. J.; and Hariharan, B. 2019. PointFlow: 3D Point Cloud Generation With Continuous Normalizing Flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 4540–4549.
- Zhai, G.; Örnek, E. P.; Chen, D. Z.; Liao, R.; Di, Y.; Navab, N.; Tombari, F.; and Busam, B. 2025. EchoScene: Indoor Scene Generation via Information Echo Over Scene Graph Diffusion. In *Computer Vision - ECCV 2024*, 167–184.
- Zhai, G.; Örnek, E. P.; Wu, S.-C.; Di, Y.; Tombari, F.; Navab, N.; and Busam, B. 2023. CommonScenes: generating commonsense 3D indoor scenes with scene graph diffusion. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Zheng, C.; and Vedaldi, A. 2023. Online Clustered Codebook. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 22741–22750.