

# DiffInDScene: Diffusion-based High-Quality 3D Indoor Scene Generation

Xiaoliang Ju<sup>1\*</sup>, Zhaoyang Huang<sup>1\*</sup>, Yijin Li<sup>2</sup>, Guofeng Zhang<sup>2</sup>, Yu Qiao<sup>3</sup>, Hongsheng Li<sup>1,4</sup>  
<sup>1</sup> CUHK MMLab <sup>2</sup> Zhejiang University <sup>3</sup> Shanghai AI Laboratory <sup>4</sup> CPII under InnoHK  
 {akira,drinkingcoder}@link.cuhk.edu.hk, hsli@ee.cuhk.edu.hk

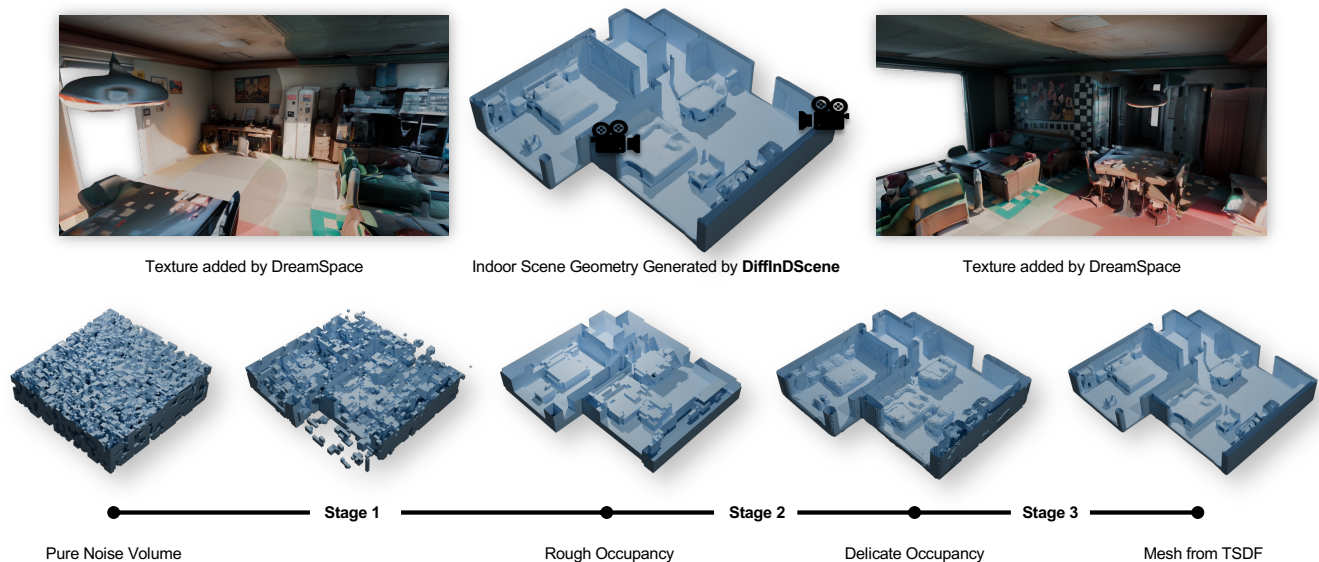


Figure 1. Coarse-to-fine indoor scene geometry generation using a sparse diffusion framework. For better visualization, the texture is produced by *DreamSpace* [46] after the scene geometry is generated by our DiffInDScene.

## Abstract

We present *DiffInDScene*, a novel framework for tackling the problem of high-quality 3D indoor scene generation, which is challenging due to the complexity and diversity of the indoor scene geometry. Although diffusion-based generative models have previously demonstrated impressive performance in image generation and object-level 3D generation, they have not yet been applied to room-level 3D generation due to their computationally intensive costs. In *DiffInDScene*, we propose a cascaded 3D diffusion pipeline that is efficient and possesses strong generative performance for Truncated Signed Distance Function (TSDF). The whole pipeline is designed to run on a sparse occupancy space in a coarse-to-fine fashion. Inspired by *KinectFusion*'s incremental alignment and fusion of local TSDF volumes, we propose a diffusion-based SDF fusion

approach that iteratively diffuses and fuses local TSDF volumes, facilitating the generation of an entire room environment. The generated results demonstrate that our work is capable to achieve high-quality room generation directly in three-dimensional space, starting from scratch. In addition to the scene generation, the final part of *DiffInDScene* can be used as a post-processing module to refine the 3D reconstruction results from multi-view stereo. According to the user study, the mesh quality generated by our *DiffInDScene* can even outperform the ground truth mesh provided by *ScanNet*.

## 1. Introduction

3D scene production is a fundamental task in 3D computer vision with many applications, such as Augmented Reality (AR), game development and embodied AI [28], where the quality of 3D scene geometry plays a paramount role. While the 3D reconstruction from multi-view stereo [1, 41] can recover scenes from real-world, the quality of the re-

\* Joint first authorship

Please visit our project page for the latest updates: <https://akirahero.github.io/diffindscene/>

sultant meshes is far from satisfactory, for the major mesh details might be lost during iterative fusion. Recently, diffusion models [11, 39] have shown their great ability in generating images and objects of high quality. Here we want to ask a question: *Can we exploit diffusion models to produce 3D scenes?* In this paper, we propose a novel framework DiffInDScene, which not only helps optimize the results of 3D reconstruction, but also generates high-quality indoor spatial geometry from scratch (see Fig. 1).

Diffusion models are a class of generative models designed for synthesizing data by iterative denoising. The popular Denoising Diffusion Probabilistic Model (DDPM) training paradigm starts the denoising from pure Gaussian noise. Training diffusion models for room-level Truncated Signed Distance Function (TSDF) volumes is challenging because of its large size. Previous 3D diffusion models only focus on object-level 3D generations. As reported by InstantNGP [24], only 2.57% voxels are informative in common 3D scenes.

To deal with the large scale of indoor scenes, we propose a coarse-to-fine sparse diffusion pipeline consists of multiple stages. The first few stages are used to generate the occupancy volume in a coarse-to-fine manner, and the last one generates the TSDF values in the sparsely occupied voxels. For the first few occupancy generation stages, a multi-scale auto-encoder is designed to encode occupancy to latent space, providing feature guidance for occupancy generation. This approach allows us to employ latent diffusion and further compress the size of input volumes. Additionally, we propose a sparse 3D diffusion model denoising only on the sparsely occupied voxels of TSDF or the occupancy latent volumes, which saves two orders of computational and memory costs.

Although the cascaded sparse diffusion pipeline significantly reduces the required computational resources, the considerable variation in room sizes still poses a challenge when attempting to directly train on a room-level TSDF volume, at the final stage of our cascaded diffusion process. To increase the data variation within each mini-batch, we randomly crop local TSDF volumes of smaller sizes from the original large volume for training. During inference, we design a stochastic TSDF fusion algorithm that generates the entire room by iterative denoising and fusing local TSDF volumes. The fusion method allows for the generation of a complete and unified TSDF within a large scene while efficiently decomposing the scene into smaller crops, thereby conserving computing resources. Our proposed diffusion method can be also utilized to refine indoor scene meshes such as the reconstruction results from multi-view stereo, such as NeuralRecon [41]. Given the TSDF volume of reconstructed scene as occupancy condition, our diffusion model can effectively refine and optimize the TSDF volume towards the ground truth.

Our contributions can be summarized as four-fold: 1) We propose a novel framework DiffInDScene for room-level indoor scene generation with a sparse diffusion model that saves two orders of resource consumption. 2) We design a multi-scale auto-encoder to provide feature guidance for the scene occupancy generation. 3) We propose a novel algorithm that fuses diffusion-based local TSDF volumes, which enables large-scale indoor scene generation. 4) DiffInDScene exhibits a promising capability in producing high-quality room-level geometry through both generation from scratch and refinement of existing reconstructions.

## 2. Related Works

**Diffusion Models.** The diffusion model [11, 38, 39] has emerged as a promising class of generative models for learning data distributions through an iterative denoising process. They have shown impressive visual quality in diverse applications of 2D image synthesis, encompassing image inpainting [21], super-resolution [13, 34], editing [23], text-to-image synthesis [25, 33], and video generation [12, 14]. Nevertheless, the application of diffusion models in the 3D domain has received limited attention in comparison to the extensive exploration seen in the 2D domain. In the 3D domain, existing research has focused on the generation of individual objects [16, 22, 27, 48], while less attention has been paid to the synthesis of entire scenes, which possess significantly higher levels of semantic and geometric complexity, as well as the expansive spatial extent present in 3D scene synthesis.

**3D Shape Generation.** Extensive exploration has been conducted on various 3D representation methods, such as voxel [43, 45], point cloud [26, 44, 47], and implicit field [19, 37], in conjunction with different generation models such as diffusion models [37] and GANs [43]. While these approaches have shown success in object-level generation, transferring them to scene generation at a larger scale is still challenging. Firstly, the large scale leads to exponential growth in computation resources consumption in training and inference processes. Additionally, object-level generation is comparatively easier due to simpler geometries and less diversity.

**3D Scene Synthesis.** In recent decades, the field of 3D scene synthesis has experienced extensive investigation, particularly driven by the proliferation of 3D indoor scene datasets [4, 6] and advancements in 3D deep learning [20, 29, 30]. However, current methods mainly focus on synthesizing plausible 3D scene arrangements [3, 8, 10, 31]. They usually learn to synthesize the scene graph as the intermediate scene representation and retrieve objects from available dataset. In contrast to these methods, we aim to simultaneously synthesize both the scene arrangement and the detailed geometry. Text2Room [15] is the most related work to ours in recent years, which leverages pre-trained

2D text-to-image models to synthesize a sequence of images and then conduct an iterative reconstruction. While such method can produce room-scale geometry, the results are often fragmentary and distorted, limiting their practical applications in areas such as gaming or AR.

### 3. Methodology

#### 3.1. Overview

To generate room-level 3D geometry, the greatest challenge is the large scale, as it requires substantial computing resources. We employ a cascaded diffusion model to generate the whole room in a coarse-to-fine manner. The first stage is to generate the coarse structure of the whole room. The following stages further refine the rough shape to a 3D occupancy field with higher resolutions. At the final stage, the resolution increases to the highest level, and we crop the whole scene to overlapped pieces to generate the final detailed Truncated Signed Distance Function (TSDF) volume. In every stage, we use a separate sparse diffusion model to reduce the resource consumption, which exclusively denoises on sparsely distributed occupancy. In our implementation, we use 3 stages to create indoor geometry up to size of  $512 \times 512 \times 128$ .

Such cascaded solution has three advantages. First, the computation resource consumption is constrained in all stages. Second, compared with piece-wise generation or incremental generation methods, the first stage of our model is able to sketch the global structure of the scene, which helps to generate a complete layout with unified and detailed structure. Third, every stage can be trained independently, and the generation process can be stopped in advance when generating an unsatisfied layout at early stage.

#### 3.2. Cascaded Diffusion for Indoor Geometry Generation

We propose a sparse cascaded diffusion model as shown in Fig. 2 (a). In our implementation, a 3-stage diffusion process is utilized to generate a complete indoor scene starting from noise. The first 2 stages are utilized to generate and refine the 3D binary scene occupancy volume, and the final is used to generate TSDF value within the occupancy. As TSDF volume only retains information near the object surface, we simply define all voxels containing valid TSDF values as the binary scene occupancy.

Assume we have multi-scale occupancy embeddings  $z^{(1)}$ ,  $z^{(2)}$  of a TSDF volume  $x$  with increasing resolutions, and their binary occupancy masks are  $M_{z^{(1)}}$ ,  $M_{z^{(2)}}$ ,  $M_x$ , satisfying

$$M_{z^{(2)}} = G_1 \left( z^{(1)}, M_{z^{(1)}} \right), \quad (1)$$

$$M_x = G_2 \left( z^{(1)}, z^{(2)}, M_{z^{(2)}} \right), \quad (2)$$

where  $G_1, G_2$  are occupancy decoders, and we will explain them in detail in Section 3.3 together with the occupancy latents  $z^{(1)}$ ,  $z^{(2)}$ . Then the 3-stage diffusion processes  $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$  can be established as follows.

- Stage 1 generates the occupancy latent code  $z^{(1)}$  of the lowest resolution. Given a fixed volume  $z_T^{(1)}$  and its occupancy mask  $M_{z_T^{(1)}}$ , the diffusion process performs denoising operation in  $T$  timesteps to obtain  $z_0^{(1)}$  as

$$z_0^{(1)} = \mathcal{D}_1(z_T^{(1)}, M_{z_T^{(1)}}), \quad (3)$$

where  $z_T^{(1)}$  is filled with Gaussian noise, and  $M_{z_T^{(1)}}$  serves as a bounding box to specify the scene size of generation. If we set the block  $M_{z_T^{(1)}}[0 : W, 0 : H, 0 : L] = 1$  and keep all other elements equal to 0, the model can generate a scene with a maximum size of  $nW \times nH \times nL$ , where  $n$  denotes compression ratio of  $z_T^{(1)}$  compared to  $x$ .

- Stage 2 generates latent code of higher resolution  $z^{(2)}$  conditioned on  $z^{(1)}$  as

$$z_0^{(2)} = \mathcal{D}_2(z_T^{(2)}, M_{z_T^{(2)}}; z_0^{(1)}), \quad (4)$$

where  $M_{z_T^{(2)}}$  is obtained utilizing Eq. (1), and  $z_T^{(2)}$  is filled with another Gaussian noise volume.

- Stage 3 generates the final TSDF volume with all those generated latent codes  $z^{(1)}$ ,  $z^{(2)}$  as input conditions

$$x_0 = \mathcal{D}_3(x_T, M_{x_T}; z_0^{(1)}, z_0^{(2)}), \quad (5)$$

where  $M_{x_T}$  is obtained by Eq. (2), and  $x_T$  is filled with Gaussian noise.

Compared with generating occupancy directly, the occupancy embedding can guide the refinement of the occupancy throughout the diffusion process, particularly at the initial stage of generation.

**Sparse Diffusion.** We follow DDPM [11, 38] to implement the sparse diffusion. Each stage of our model employs a separate sparse diffusion, with the only distinction lies in the types of input and output. In this section, we use  $v$  to represent any kind of volumes such as TSDF  $x$  and latent codes  $z^{(1)}$ ,  $z^{(2)}$ , and  $y$  to denote the diffusion condition with  $M$  as their shared sparsity mask.

DDPM [11, 38] transforms a sample volume  $v_0$ , to a white Gaussian noise  $v_T \sim \mathcal{N}(0, 1)$  in  $T$  steps. In each step  $t$ , the sample  $v_t$  is obtained by adding i.i.d. Gaussian noise with variance  $\beta_t$  and scaling the sample in the previous step  $v_{t-1}$  with  $\sqrt{1 - \beta_t}$ :

$$q(v_t | v_{t-1}) = \mathcal{N} \left( v_t; \sqrt{1 - \beta_t} v_{t-1}, \beta_t \mathbf{I} \right), \quad (6)$$

which is also called the forward direction. On the other hand, the reverse process can be depicted as:

$$p_\theta(v_{t-1} | v_t, y, M) = \mathcal{N}(v_{t-1}; \mu_\theta, \Sigma_\theta), \quad (7)$$

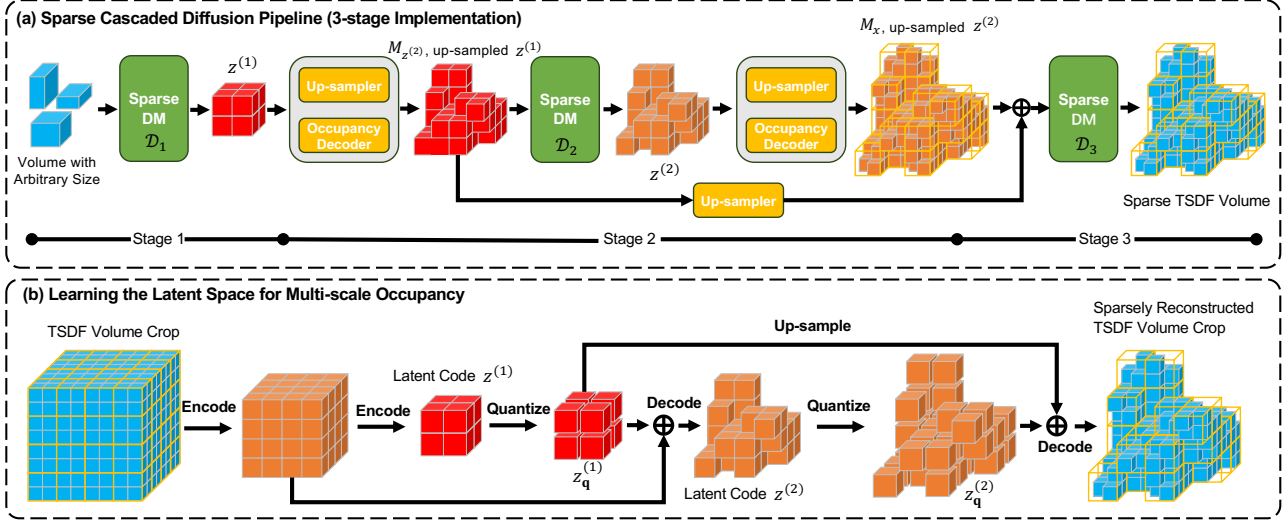


Figure 2. Sparse cascaded diffusion with a multi-scale occupancy embedding.

where  $y$  denotes the extra condition. Diffusion models are trained to reverse the forward process, predicting  $\mu_\theta$  as

$$\mu_\theta(v_t, y, M, t) = \frac{1}{\sqrt{\alpha_t}} \left( v_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(v_t, y, M, t) \right), \quad (8)$$

where  $\alpha_t := 1 - \beta_t$ ,  $\bar{\alpha}_t := \prod_{s=0}^t \alpha_s$ , and  $\epsilon_\theta$  is the neural network with the parameter set  $\theta$ , which has a UNet-like structure. With the occupancy mask  $M$ ,  $\epsilon_\theta$  denoises only sparsely occupied voxels with sparse convolutions and attentions. We implement  $\epsilon_\theta$  using the engine of TorchSparse [42]. The details of network architecture are provided in the supplementary materials.

A mean square error loss masked by  $M$  is used to supervise the noise prediction as

$$L_{\text{diff}} = E_{t, v_0, \epsilon, y, M} \left( M \odot \left[ \|\epsilon - \epsilon_\theta(v_t, y, M, t)\|_2^2 \right] \right). \quad (9)$$

### 3.3. Learning the Latent Space for Multi-scale Occupancy

To obtain hierarchical occupancy embeddings and their decoders, we design a multi-scale Patch-VQGAN as Fig. 2(b), inspired by VQ-VAE-2 [32]. The encoder takes TSDF volume  $x$  as input, and outputs the occupancy embedding  $z^{(1)}, z^{(2)}$ . These embeddings are expected to be decoded to the occupancy mask  $M_{z^{(2)}}$  and  $M_x$  as Eq. (1) and Eq. (2), where ground truth of  $M_{z^{(2)}}$  is obtained by downsampling from  $M_x$  via maxpooling. To capture more shape details, we also add a TSDF decoding head. After training, the encoder is no longer utilized, and only the occupancy decoders are employed to convert the latent code to occupancy.

Formally, we define the encoder as  $E$ , the decoders as  $G$ , and the element-wise quantization of latent volume as  $\mathbf{q}(\cdot)$ . Any ground truth TSDF volume  $x \in \mathbb{R}^{H \times W \times L}$  can be encoded progressively into  $(z^{(1)}, z^{(2)}) = E(x)$ , where  $z^{(1)}, z^{(2)} \subseteq \mathbb{R}^d$ . The quantization is formulated as

$$\mathbf{q}(z) := \left( \arg \min_{z_p \in Z} \|z_{ijk} - z_p\| \right) \in \mathbb{R}^{h \times w \times l \times d}, \quad (10)$$

where  $Z = \{z_k\}_{k=1}^K \subset \mathbb{R}^d$  denotes the discrete codebook of size  $K$ , and  $z_{ijk}$  represents the latent vector at coordinate  $i, j, k$  of the latent volume. For simplicity, we define  $z_{\mathbf{q}}^{(1)} := \mathbf{q}(z^{(1)})$ ,  $z_{\mathbf{q}}^{(2)} := \mathbf{q}(z^{(2)})$  in Fig. 2(b).

During the training process, we randomly crop cubes of  $96 \times 96 \times 96$  from the original TSDF volumes as data samples, so that diverse crops from different scenes are contained in each mini-batch. The training loss is defined in Eq. (11), where  $L_{\text{rec}}, L_{\text{vq}}, L_{\text{GAN}}$  denote the reconstruction loss, vector quantization loss and the adversarial loss from a simple multi-layer discriminator,

$$L = L_{\text{rec}} + \lambda_1 L_{\text{vq}} + \lambda_2 L_{\text{GAN}}, \quad (11)$$

where  $\lambda_1$  and  $\lambda_2$  are hyper parameters to weight losses of different types. For the encoder-decoder training, we use a  $L_1$  loss to supervise the TSDF value, and binary cross-entropy (BCE) to supervise the occupancy masks, as shown in Eq. (12).

$$L_{\text{rec}} = \text{BCE}(\hat{M}_x, M_x) + \text{BCE}(\hat{M}_{z^{(2)}}, M_{z^{(2)}}) + \|x - \hat{x}\|_1 \quad (12)$$

The  $L_{\text{vq}}, L_{\text{GAN}}$  are defined similar to [7].

### 3.4. Local Fusion for Global Diffusion

The volume cropping operation in the training of the final stage diffusion raises a question: how to infer on a complete



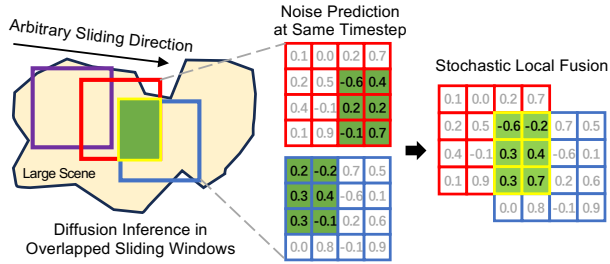


Figure 3. Stochastic TSDF Fusion.

scene using a model trained on crops? The independent generation in a crop-by-crop manner as image generators such as [21, 33] may cause inconsistent results between adjacent crops. To address this issue, a fusion algorithm is proposed to perform the joint diffusion process concurrently on the overlapping local volumes.

During inference, we split the indoor space into  $K$  overlapping 3D crops that cover the entire room, and generate the TSDF for the entire room by concurrently diffusing the  $K$  crops with stochastic fusion. We denote the global TSDF at the timestep  $t$  as  $x_t$  and the  $k$ -th crop at the timestep  $t$  as  $x_t^k$  and the global TSDF at the timestep  $t$  as  $x_t$ . At time step  $t$ , we need to obtain the global TSDF  $x_t$  by fusing local TSDFs  $x_t^k$  and then update local TSDFs from the global TSDF. After synchronizing local TSDFs with fusion, each crop steps to the next time step independently. Specifically, for a voxel grid  $\mathbf{p}$ , suppose  $\mathcal{G}(\mathbf{p})$  is a set of crops that cover  $\mathbf{p}$ , we need to obtain  $x_t(\mathbf{p})$  by fusing the crops  $\{x_t^k(\mathbf{p}) | k \in \mathcal{G}(\mathbf{p})\}$  overlapping on  $\mathbf{p}$ .

**Stochastic TSDF Fusion.** A straightforward fusion algorithm is taking the average TSDFs of the local crops:  $x_t(\mathbf{p}) = \frac{1}{|\mathcal{G}(\mathbf{p})|} \sum_{k \in \mathcal{G}(\mathbf{p})} x_t^k(\mathbf{p})$ , which is also adopted by the classical KinectFusion [17]. However, average fusion significantly reduces the variance of the sample distribution. To ensure the generation quality and global consistency, we propose stochastic fusion to keep the distribution and fuse TSDFs in the reverse process. An example of 2-windows fusion is shown in Fig. 3. Specifically, we randomly sample an index  $k$  from  $\mathcal{G}(\mathbf{p})$  in a uniform distribution to update the global TSDF  $x_t(\mathbf{p}) = x_t^k(\mathbf{p})$ ,

$$x_t(\mathbf{p}) \sim \mathcal{N}(\mu_t^k(\mathbf{p}), \Sigma_t^k(\mathbf{p})), k = \text{RandomSelect}(\mathcal{G}(\mathbf{p})). \quad (13)$$

### 3.5. Extension to Refining 3D Reconstruction

Given a rough geometry represented by the occupancy volume of a scene, the TSDF volume can be generated by directly adopting the final stage of our cascaded diffusion framework.

The most direct application is to recover/refine scenes produced by multi-view stereo methods or even LiDAR

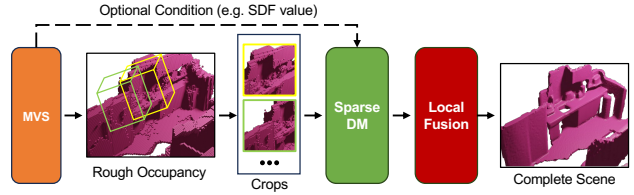


Figure 4. Refine or recover scene from MVS methods.

mappers as depicted in Fig. 4, as obtaining an occupancy of the scene is relatively straightforward using Multi-view Stereo (MVS) techniques. If the MVS method provides a TSDF result, it can be utilized as the diffusion condition to our cascaded diffusion network. The diffusion process is applied to the crops of the occupancy volume as inputs, and the results are fused using our local fusion module at each timestep, so that we can obtain a complete and refined scene reconstruction.

## 4. Experiment

DiffInDScene is a versatile tool capable of generating detailed indoor scene geometry at the room level. It is not only capable of creating scenes from scratch but also has the ability to refine or recover scenes using occupancy fields produced by multi-view stereo methods. To train DiffInDScene for scene generation, we utilized the 3D-FRONT dataset [9] provided by Alibaba, which consists of 6813 furnished houses with meshes. For training purposes, we selected 5913 houses with dimensions smaller than  $512 \times 512 \times 128$ , with a voxel size of 0.04m. For scene refinement or recovery on multi-view stereo (MVS), we trained DiffInDScene using the training split of the ScanNet dataset, with the NeuralRecon [41] as the MVS module. The ScanNet dataset contains 1613 indoor scenes, each accompanied by ground-truth camera poses and surface reconstructions. To ensure consistency, we adopted the same data split as NeuralRecon for training and evaluation purposes.

### 4.1. Indoor 3D Scene Generation

In this section, we will present the results of our indoor scene generation, both quantitatively and qualitatively. To the best of our knowledge, Text2Room stands as the state-of-the-art solution to generate 3D geometry of room-level indoor scenes. Since there are few other works that can directly generate scene structures and output the corresponding mesh models, here we also compare with "Text2Room + Poisson" to enrich the experimental comparison, which means a Poisson reconstruction [18] is added as a refinement of Text2Room.

**Metrics.** We employ mesh quality as a measure to reflect the overall quality of the generated 3D scene geometry. Regarding mesh quality, we noticed that noisy meshes and

Table 1. Geometry quality comparison, including mesh quality (Aspect Ratio, Circularity, and Shape Regularity) and user study on the completeness and perceptual quality.

	Text2Room [15]	Text2Room + Poisson [18]	Ours
Aspe. mean $\uparrow$	0.416	0.443	<b>0.473</b>
Aspe. var $\downarrow$	0.022	0.029	<b>0.009</b>
Circ. mean $\uparrow$	0.674	0.709	<b>0.781</b>
Circ. var $\downarrow$	0.052	0.057	<b>0.022</b>
Shap. mean $\uparrow$	0.716	0.730	<b>0.816</b>
Shap. var $\downarrow$	0.045	0.060	<b>0.023</b>
Completeness $\uparrow$	2.532	3.228	<b>4.856</b>
Perceptual $\uparrow$	2.472	2.812	<b>4.836</b>

high-quality meshes exhibit distinct distributions of triangle shapes. Noisy meshes and problematic regions are characterized by triangles with low aspect ratio, circularity, and shape regularity, as introduced in [2, 5]. Therefore, these 3 factors are selected as the objective metric of mesh quality.

In addition, we perform a user study similar to the one conducted in Text2Room [15] as the subjective metric on the mesh quality, including the completeness and perceptual quality. The scores range from 1 to 5, with higher scores indicating a better alignment with the evaluation metrics.

**Quantitative Results.** For the objective metrics, We randomly select 11 scenes from different methods, and calculate the average value of all triangular mesh faces. For the user study, we randomly choose 5 scenes generated by each method and gather ratings from a group of 40 users with basic knowledge in 3D modeling. We summarize the evaluation results as Table 1 shows. Our method demonstrates superior performance on those metrics compared to the other approaches. While employing Poisson reconstruction as a refinement step enhances the results of Text2Room, it does not lead to a significant improvement.

**Qualitative Results.** The generated scene samples from different methods are listed in Fig. 5. We compare both untextured meshes and textured scene renderings in our evaluation. While Text2Room is capable of reconstructing the scene from images, it often results in serious distortions and fragmentation. As indicated by our quantitative evaluation, applying Poisson reconstruction helps in filling the holes, but it has limited impact on improving the overall structure of the scene. Our method can produce larger rooms with clearer structures and more complex layout. Furthermore, by incorporating DreamSpace as a post-processing step, we can achieve high-quality scene renderings. This combination allows for enhanced visual output and improved overall scene representation.

Compared with the middle-size house showed in Fig.5, two larger and more complex samples are shown in Fig. 6. Due to limited space, we demonstrate additional samples

Table 2. Resource consumption comparison.

Batch Size	TFLOPs	Parameters(M)	GPU Memory(GB)		
	1	1	1	2	4
Sparse	0.008	161.5	11.8	15.3	22.8
Dense	3.290	161.5	22.8	-	-

showcasing larger views and **conditional Sketch-to-Scene generation** in the supplementary materials.

## 4.2. Ablation Study

**Sparse or Dense.** To compare the dense diffusion and the sparse diffusion, we implement two networks using sparse and dense convolution respectively, with exactly the same structures. Several randomly cropped TSDF volumes ( $96 \times 96 \times 96$ ) from ScanNet dataset are fed into these two models. The resource consumption of the two strategies are shown in Table 2. This experiment is conducted on the platform equipped with RTX3090 GPU with 24GB memory, with the diffusion model running in training status. The sparse diffusion requires fewer resources and runs faster with fewer parameters, due to the characteristic of the occupancy distribution. In our test data, the largest occupancy rate of the TSDF crops are less than 20%.

**Diffusion with Fusion.** To compare the different fusion methods mentioned in Section 3.4, we perform the final stage of our cascaded diffusion model on the given occupancy from a scene of ScanNet dataset. The results are shown in Fig. 7. Individual diffusion presents significant inconsistency between adjacent crops. Average fusion generates meshes of lower quality because the sample distribution during diffusion is disturbed. Our stochastic diffusion generates high-quality meshes with global consistency.

## 4.3. Refinement and Recovery on MVS

**Metrics.** Our objective is to obtain high-quality 3D scenes from MVS, aiming for results comparable to the ground truth in various aspects, including details, completeness, tightness, and sharpness. As a generative model, the proposed method operates within a relaxed occupancy, which may cause drift from the ground truth. As a result, conventional correspondence-based metrics for 3D reconstruction may be not suitable, while we include such evaluation [36] in the supplementary materials. Instead, we employ metrics such as normal error distribution and mesh quality.

Normal error is used to evaluate the similarity of surface orientation between the reconstructed and ground truth meshes. Smaller normal errors indicate better alignment to the ground truth. We filter out outliers with errors larger than  $90^\circ$  and analyze the percentage of inlier normals below a threshold ( $< T^\circ$  ratio) and the mean normal error of those inliers. We use the same mesh quality metrics as Section 4.1.




Figure 5. Comparison of indoor scene generation. The texture renderings of our results are produced by DreamSpace [46]. The ceilings are removed from the original meshes for a better visualization.




Figure 6. Samples with more complex structures generated by DiffInDScene. Note that the floating objects in this figure are ceiling lamps, which are reserved after we remove the ceilings for visualization.

**Results.** The samples of scene reconstruction are shown in Fig. 8. As Table 3 shows, our method significantly outperforms the other three methods on all thresholds of normal

error. As for the mesh quality, we compute the mean and variance of the three scores described above as shown in Table 4. Our results present significantly better performance than the other three methods, i.e., higher scores with smaller score variance. Moreover, we also include ground truth mesh for comparison. Interestingly, our results also outperform the ground truth except for the aspect ratio, which indicates the high quality of our reconstructed meshes.

**User Study.** The user study is conducted to rank scene quality, providing subjective evaluations to complement the objective metrics. We randomly select 10 scenes from the test split of the ScanNet dataset and generate reconstructed meshes using four methods. Users are asked to rank the methods based on details, completeness, plane quality, and edge quality. We collect ranking results from 51 users. As shown in Table 5, our reconstructed mesh significantly outperforms NeuralRecon, “NeuralRecon+Laplacian Denois-



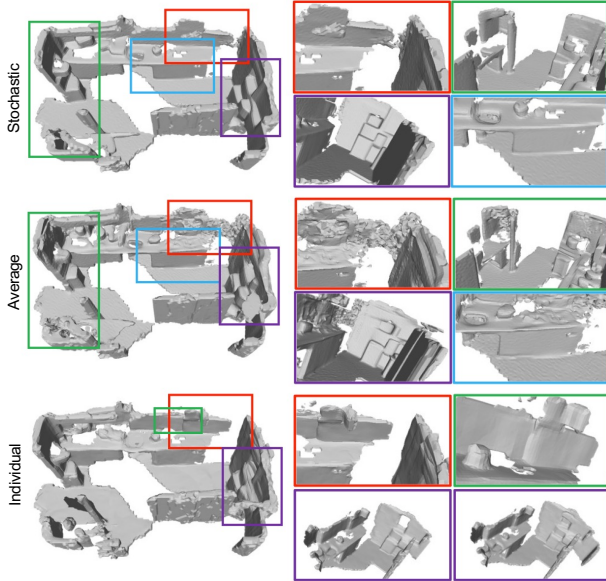


Figure 7. Comparison of different fusion methods in the final stage of our cascaded diffusion.

Table 3. Normal error comparison. “NR”, “Lap”, and “SR” denote NeuralRecon, Laplacian and SimpleRecon, respectively.

	NR [41]	NR + Lap [40]	SR [35]	NR + Ours
< 90° mean↓	34.65	37.12	35.44	<b>30.4</b>
< 90° ratio↑	100%	100%	100%	<b>100%</b>
< 45° mean↓	10.27	12.34	12.51	<b>8.09</b>
< 45° ratio↑	59.97%	57.97%	60.20%	<b>65.05%</b>
< 30° mean↓	6.45	8.17	8.31	<b>5.05</b>
< 30° ratio↑	52.88%	49.89%	51.67%	<b>59.27%</b>


Table 4. Mesh quality comparison. We compare the mean and variance of three scores: Aspect Ratio, Circularity, and Shape Regularity. NR(occ) means only the occupancy is used, without the conditional TSDF depicted by the dash line in Fig. 4.

	NR [41]	NR + Lap [40]	SR [35]	NR(occ) + Ours	NR + Ours	GT
Aspe. mean↑	0.459	0.437	0.436	0.469	0.457	<b>0.477</b>
Aspe. var↓	0.022	0.023	0.024	0.020	<b>0.016</b>	0.022
Circ. mean↑	0.740	0.712	0.708	0.742	<b>0.763</b>	0.758
Circ. var↓	0.041	0.052	0.054	0.041	<b>0.030</b>	0.034
Shap. mean↑	0.772	0.746	0.739	0.793	<b>0.797</b>	0.793
Shap. var↓	0.041	0.052	0.055	0.041	<b>0.030</b>	0.031


ing” (denoted as “NR + Lap”), and even surpasses the quality of the ground truth meshes.

## 5. Conclusion & Acknowledgement

We have presented DiffInDScene as a novel framework for diffusion-based high-quality indoor scene generation. DiffInDScene mainly consists of three modules: 1) a sparse diffusion network that efficiently denoises 3D volumes on



(a) Scene samples refined on NeuralRecon (NR).



(b) Comparison of scene quality from different methods.

Figure 8. Sample scene reconstructions on ScanNet dataset. The meshes are colored according to curvatures in sub-figure(b), where green regions denote lower curvatures.

Table 5. User study on the refinement of scene reconstruction.

	NR [41]	NR + Lap [40]	NR + Ours	GT
Details↑	12.26	6.80	17.41	<b>25.50</b>
Completeness↑	11.15	8.84	<b>21.30</b>	20.76
Tight Plane↑	5.48	12.10	<b>25.85</b>	18.17
Sharp Edge↑	8.22	10.20	<b>22.58</b>	21.19
Overall (Sum)↑	37.10	37.98	<b>87.14</b>	85.62

occupied voxels, 2) a multi-scale Patch-VQGAN for occupancy decoding, 3) a cascaded diffusion framework to generate room-level scene from scratch, and 4) a stochastic fusion algorithm for diffusion-based local TSDFs, which enables large-scale indoor scene generation.

This project is funded in part by National Key R&D Program of China Project 2022ZD0161100, by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)’s InnoHK, by General Research Fund of Hong Kong RGC Project 14204021. Hongsheng Li is a PI of CPII under the InnoHK. We would like to thank the authors of DreamSpace [46] for supporting us to produce vivid visualization. We also appreciate the participation of users in our user study.



## References

- [1] Aljaz Bozic, Pablo Palafox, Justus Thies, Angela Dai, and Matthias Nießner. Transformerfusion: Monocular rgb scene reconstruction using transformers. *Advances in Neural Information Processing Systems*, 34:1403–1414, 2021. 1
- [2] Jan Brandts, Sergey Korotov, and Michal Krížek. On the equivalence of regularity criteria for triangular and tetrahedral finite element partitions. *Computers & Mathematics with Applications*, 55(10):2227–2233, 2008. Advanced Numerical Algorithms for Large-Scale Computations. 6
- [3] Angel Chang, Manolis Savva, and Christopher D Manning. Learning spatial knowledge for text to 3d scene generation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 2028–2038, 2014. 2
- [4] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. 2
- [5] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, Guido Ranzuglia, et al. Meshlab: an open-source mesh processing tool. In *Eurographics Italian chapter conference*, pages 129–136. Salerno, Italy, 2008. 6
- [6] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 2
- [7] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 4
- [8] Matthew Fisher, Manolis Savva, Yangyan Li, Pat Hanrahan, and Matthias Nießner. Activity-centric scene synthesis for functional 3d scene modeling. *ACM Transactions on Graphics (TOG)*, 34(6):1–13, 2015. 2
- [9] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Bin-qiang Zhao, et al. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10933–10942, 2021. 5
- [10] Qiang Fu, Xiaowu Chen, Xiaotian Wang, Sijia Wen, Bin Zhou, and Hongbo Fu. Adaptive synthesis of indoor scenes via activity-associated object relation graphs. *ACM Transactions on Graphics (TOG)*, 36(6):1–13, 2017. 2
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2, 3
- [12] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 2
- [13] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23(47):1–33, 2022. 2
- [14] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022. 2
- [15] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. *arXiv preprint arXiv:2303.11989*, 2023. 2, 6
- [16] Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural wavelet-domain diffusion for 3d shape generation. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 2
- [17] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, et al. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 559–568, 2011. 5
- [18] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, page 0, 2006. 5, 6
- [19] Muheng Li, Yueqi Duan, Jie Zhou, and Jiwen Lu. Diffusion-sdf: Text-to-shape via voxelized diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12642–12651, 2023. 2
- [20] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel cnn for efficient 3d deep learning. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [21] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. 2, 5
- [22] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2837–2845, 2021. 2
- [23] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021. 2
- [24] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 2
- [25] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 2
- [26] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generat-

- ing 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. 2
- [27] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 2
- [28] Xavier Puig, Eric Undersander, Andrew Szot, Mikael Dal-laire Cote, Tsung-Yen Yang, Ruslan Partsey, Ruta Desai, Alexander William Clegg, Michal Hlavac, So Yeon Min, et al. Habitat 3.0: A co-habitat for humans, avatars and robots. *arXiv preprint arXiv:2310.13724*, 2023. 1
- [29] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 2
- [30] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 2
- [31] Siyuan Qi, Yixin Zhu, Siyuan Huang, Chenfanfu Jiang, and Song-Chun Zhu. Human-centric indoor scene synthesis using stochastic grammar. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5899–5908, 2018. 2
- [32] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019. 4
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 2, 5
- [34] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2
- [35] Mohamed Sayed, John Gibson, Jamie Watson, Victor Prisacariu, Michael Firman, and Clément Godard. Simplerecon: 3d reconstruction without 3d convolutions. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pages 1–19. Springer, 2022. 8
- [36] Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06)*, pages 519–528. IEEE, 2006. 6
- [37] Jaehyeok Shim, Changwoo Kang, and Kyungdon Joo. Diffusion-based signed distance fields for 3d shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20887–20897, 2023. 2
- [38] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 2, 3
- [39] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 2
- [40] Olga Sorkine, Daniel Cohen-Or, Yaron Lipman, Marc Alexa, Christian Rössl, and H-P Seidel. Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 175–184, 2004. 8
- [41] Jiaming Sun, Yiming Xie, Linghao Chen, Xiaowei Zhou, and Hujun Bao. Neuralrecon: Real-time coherent 3d reconstruction from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15598–15607, 2021. 1, 2, 5, 8
- [42] Haotian Tang, Zhijian Liu, Xiuyu Li, Yujun Lin, and Song Han. Torchsparse: Efficient point cloud inference engine. *Proceedings of Machine Learning and Systems*, 4:302–315, 2022. 4
- [43] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *Advances in neural information processing systems*, 29, 2016. 2
- [44] Zijie Wu, Yaonan Wang, Mingtao Feng, He Xie, and Ajmal Mian. Sketch and text guided diffusion model for colored point cloud generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8929–8939, 2023. 2
- [45] Jianwen Xie, Zilong Zheng, Ruiqi Gao, Wenguan Wang, Song-Chun Zhu, and Ying Nian Wu. Learning descriptor networks for 3d shape synthesis and analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8629–8638, 2018. 2
- [46] Bangbang Yang, Wenqi Dong, Lin Ma, Wenbo Hu, Xiao Liu, Zhaopeng Cui, and Yuewen Ma. Dreamspace: Dreaming your room space with text-driven panoramic texture propagation. *arXiv preprint arXiv:2310.13119*, 2023. 1, 7, 8
- [47] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4541–4550, 2019. 2
- [48] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5826–5835, 2021. 2