

---

# CTRL-ROOM: CONTROLLABLE TEXT-TO-3D ROOM MESHES GENERATION WITH LAYOUT CONSTRAINTS

Chuan Fang<sup>1,2</sup>, Xiaotao Hu<sup>2,3</sup>, Kunming Luo<sup>1,2</sup>, Ping Tan<sup>1,2</sup>\*

<sup>1</sup> Hong Kong University of Science and Technology

<sup>2</sup> Light Illusions

<sup>3</sup> Nankai University

cfangac@connect.ust.hk, pingtan@ust.hk

## ABSTRACT

Text-driven 3D indoor scene generation could be useful for gaming, film industry, and AR/VR applications. However, existing methods cannot faithfully capture the room layout, nor do they allow flexible editing of individual objects in the room. To address these problems, we present Ctrl-Room, which is able to generate convincing 3D rooms with designer-style layouts and high-fidelity textures from just a text prompt. Moreover, Ctrl-Room enables versatile interactive editing operations such as resizing or moving individual furniture items. Our key insight is to separate the modeling of layouts and appearance. Our proposed method consists of two stages, a ‘Layout Generation Stage’ and an ‘Appearance Generation Stage’. The ‘Layout Generation Stage’ trains a text-conditional diffusion model to learn the layout distribution with our holistic scene code parameterization. Next, the ‘Appearance Generation Stage’ employs a fine-tuned ControlNet to produce a vivid panoramic image of the room guided by the 3D scene layout and text prompt. In this way, we achieve a high-quality 3D room with convincing layouts and lively textures. Benefiting from the scene code parameterization, we can easily edit the generated room model through our mask-guided editing module, without expensive editing-specific training. Extensive experiments on the Structured3D dataset demonstrate that our method outperforms existing methods in producing more reasonable, view-consistent, and editable 3D rooms from natural language prompts.


## 1 INTRODUCTION

High-quality textured 3D models are important for a broad range of applications, from interior design and games to simulators for embodied AI. Indoor scenes are of particular interest among all 3D content. Typically, 3D indoor scenes are manually designed by professional artists, which is time-consuming and expensive. While recent advancements in generative models (Poole et al., 2022; Chen et al., 2023; Lin et al., 2023; Seo et al., 2023) have simplified the creation of 3D models from textual descriptions, extending this capability to text-driven 3D indoor scene generation remains a challenge because indoor scenes exhibit strong semantic layout constraints, e.g., neighboring walls are perpendicular and the TV set often faces a sofa, which are more complicated than objects.

Existing text-driven 3D indoor scene generation approaches, such as Text2Room (Höller et al., 2023) and Text2NeRF (Zhang et al., 2023), are designed with an incremental framework. They create 3D indoor scenes by incrementally generating different viewpoints frame-by-frame and reconstructing the 3D mesh of the room from these sub-view images. However, their incremental approaches often fail to model the global layout of the room, resulting in unconvincing results that lack semantic plausibility. As shown in the left of Fig. 1 (a), the result of Tex2Room exhibits repeating objects, e.g. several cabinets in a living room, and does not follow the furniture layout patterns. We refer to this problem as the ‘*Penrose Triangle problem*’ in our paper, which has plausible 3D structures everywhere locally but lacks global consistency. Moreover, previous methods fail to enable user interactive manipulation as their resulting 3D geometry and texture are uneditable.

---

\*Corresponding author.



(a) Comparison with Text2Room Höllein et al. (2023) and MVDiffusion Tang et al. (2023c).



(b) Flexible editing by instruction or mouse clicks.

Figure 1: We present *Ctrl-Room* to achieve fine-grained textured 3D indoor room generation and editing. (a): compared with the Text2Room (Höllein et al., 2023) and MVDiffusion(Tang et al., 2023c), *Ctrl-Room* can generate rooms with more plausible 3D structures. (b): *Ctrl-Room* supports flexible editing. Users can replace furniture items or change their positions easily.

Indoor scenes might also be represented by a panorama image. Several works (Lin et al., 2019; 2021; Shum et al., 2023; Tang et al., 2023c) have been proposed to generate such a panorama from a text prompt. We might further recover the depth map of these images to build a textured 3D room model. However, these works also cannot guarantee correct room layouts. As shown on the right of Fig. 1 (a), a bedroom generated by MVDiffusion (Tang et al., 2023c) contains multiple beds, which also violates the room layout prior. Furthermore, these methods also cannot easily control the individual room objects.

To address these shortcomings, we propose a novel two-stage method to generate a high-fidelity and editable 3D room. The key insight is to separate the generation of 3D geometric layouts from that of the visual appearance, which allows us to better capture the room layout and achieve flexible editing at the same time. In the first stage, from text input, our method creates plausible room layouts with various furniture types and positions. Unlike previous scene synthesis methods (Tang et al., 2023a; Paschalidou et al., 2021) that only focus on the furniture arrangement, our approach further considers walls with doors and windows, which play an essential role in the layout. To achieve this goal, we parameterize the room by a holistic scene code, which represents a room as a set of objects. Each object is represented by a vector capturing its position, size, semantic class, and orientation. Based on our compact parameterization, we design a diffusion model to learn the 3D room layout distribution from the Structured3D dataset (Zheng et al., 2020).

Once the indoor scene layout is fixed, in the second stage, our method generates the room appearance with the guidance of the 3D room layout. We model the appearance of the room as a panoramic image, which is generated by a text-to-image latent diffusion model. Unlike previous text-to-panorama works (Tang et al., 2023c; Chen et al., 2022), our method explicitly enforces room layout constraints and guarantees plausible 3D room structures and furniture arrangement. To achieve this goal, we convert the 3D layout synthesized in the first stage into a semantic segmentation map and feed it to a fine-tuned ControlNet (Zhang & Agrawala, 2023) model to create the panorama image.

Most importantly, benefiting from the separation of layout and appearance, our method enables flexible editing on the generated 3D room. The user can replace or modify the size and position of furniture items, e.g. replacing the TV and TV stand or moving up the picture as in Fig. 1 (b), by instructions or mouse clicks. Our method can quickly update the room according to the edited room layout through our mask-guided editing module without expensive editing-specific training.

---

The updated room appearance maintains consistency with the original version while satisfying the user’s edits. To our knowledge, it’s the first work that achieves 3D indoor scene editing through a 2D diffusion model.

The main contributions of this paper are summarized as:

- We design a two-stage method for 3D room generation from text input, which separates the geometric layout generation and appearance generation. In this way, our method can better capture the room layout constraints in real-world data and produce a vivid and rich appearance at the same time.
- Our separation of geometric layout and visual appearance allows us to have flexible control and editing over the generated 3D room model. Users can adjust the size, semantic class, and position of furniture items easily.
- We introduce a novel method to generate and edit panoramic images, which achieves high-quality results with loop consistency through a pre-trained latent image diffusion model without expensive editing-specific training.

## 2 RELATED WORK

### 2.1 TEXT-BASED 3D OBJECT GENERATION

Early methods employ 3D datasets to train generation models. Chen et al. (2019) learn a feature representation from paired text and 3D data and design a GAN network to generate 3D shapes from text. However, 3D datasets are scarce which makes these methods difficult to scale. More recent methods (Nichol et al., 2022; Poole et al., 2022; Lin et al., 2023; Wang et al., 2023a; Chen et al., 2023; Wang et al., 2023b) exploit the powerful 2D text-to-image diffusion models (Rombach et al., 2022; Saharia et al., 2022) for 3D model generation. Typically, these methods generate one or multiple 2D images in an incremental fashion and optimize the 3D model accordingly. Point-E (Nichol et al., 2022) employs a text-to-image diffusion model (Rombach et al., 2022) to generate a single-view image for a point cloud diffusion model to generate a 3D point cloud. DreamFusion (Poole et al., 2022) introduces a loss based on probability density distillation and optimizes a randomly initialized 3D model through gradient descent. Magic3D (Lin et al., 2023) uses a coarse model to represent 3D content and accelerates it using a sparse 3D hash grid structure. Wang et al. (2023a) use Score Jacobian Chaining to aggregate the results of the 2D diffusion models to generate a 3D scene. Fantasia3D (Chen et al., 2023) optimizes a mesh from scratch with DMTet (Shen et al., 2021) and stable diffusion (Rombach et al., 2022). To alleviate over-saturation, over-smoothing, and low-diversity problems, ProlificDreamer (Wang et al., 2023b) models and optimizes the 3D parameters, NeRF (Mildenhall et al., 2021) or mesh, through variational score distillation. However, all these methods focus on text-based 3D object generation. They cannot be directly applied to create 3D rooms that have additional structural layout constraints.

### 2.2 TEXT-BASED 3D ROOM GENERATION

**Room Layout Synthesis** Many works (Tang et al., 2023a; Paschalidou et al., 2021) have been proposed to learn the room layout distribution from 3D room models designed by professional artists. LayoutGAN (Li et al., 2019) is a generative adversarial network designed for generating class and geometric labels for room layout boxes. LayoutVAE (Jyothi et al., 2019) generates room layouts by conditioning on object label sets. Layout generation has been greatly boosted by transformer-based methods. LayoutTransformer (Gupta et al., 2021) employs self-attention to capture relationships between elements and then accomplish layout completion. ATISS (Paschalidou et al., 2021) proposes an autoregressive transformer to generate proper indoor scenes with only the room type and floor plan as the input. DiffuScene (Tang et al., 2023a) uses a fully connected scene graph to represent a scene and proposes a diffusion model to generate all objects in the scene. These methods only focus on the relative positions of different furniture items and ignore their appearance. Furthermore, unlike our methods, they do not consider walls, doors, and windows when learning layouts, which play important roles in the arrangement of real furniture.

---

**Panoramic Image Generation** Another line of works (Lin et al., 2019; 2021; Shum et al., 2023) represents an indoor scene by a panorama image without modeling 3D shapes. These methods enjoy the benefits of abundant image training data and produce vivid results. COCO-GAN (Lin et al., 2019) produces a set of patches and assemble them into a panoramic image. InfinityGAN (Lin et al., 2021) uses the information of two patches to generate the parts between them, and finally obtains a panoramic image. Shum et al. (2023) proposes a 360-aware layout generator to produce furniture arrangements and uses this layout to synthesize a panoramic image based on the input scene background. MVDiffusion (Tang et al., 2023c) simultaneously generates multi-view perspective images and proposes a correspondence-aware attention block to maintain multi-view consistency, and then transfers these images to a panorama. These methods might suffer from incorrect room layout since they do not enforce layout constraints. Furthermore, the results of these methods cannot be easily edited, e.g. resizing or moving some furniture items around, because they do not maintain an object-level representation.

**3D Room Generation** There are also 3D room generation methods. GAUDI (Bautista et al., 2022) generates immersive 3D indoor scenes rendered from a moving camera. It disentangles the 3D representation and camera poses to ensure the consistency of the scene during camera movement. Text2room (Höllein et al., 2023) incrementally synthesizes nearby images with a 2D diffusion model and recovers its depth maps to stitch these images into a 3D room model. Text2Room is the closest to our work, but it cannot handle the geometric and textural consistency among multi-posed images, resulting in the ‘*Penrose Triangle problem*’. In our method, we take both geometry and appearance into consideration and create a more geometrically plausible 3D room.

### 2.3 DIFFUSION-BASED IMAGE EDITING

Diffusion models have been employed in image editing and produced inspiring results. Prompt-to-prompt (Hertz et al., 2022) requires only text input and modifies the attention map within pre-trained diffusion models to achieve image editing capabilities. Imagic (Kawar et al., 2023) generates an edited image that aligns with the given input text through semantic optimization of text embedding within the image. But these editing methods all modify images at a high level (through text). Inspired by the idea that latent code can determine the spatial distribution of generated images (Mao et al., 2023), DragDiffusion (Shi et al., 2023) enables users to perform pixel-level spatial control on images by allowing pixel dragging. In contrast, DragonDiffusion (Mou et al., 2023) requires no training or further fine-tuning and all consistency information is from the image. Only relying on the information inside an image limits flexible editing. In pursuit of better performance and control, we additionally introduce semantic segmentation maps generated from the room layout to guide the panorama image generation to achieve flexible control.

## 3 METHOD

In order to achieve text-based 3D indoor scene generation and editing, we propose our Ctrl-Room. We first generate the room layout from the input text and then generate the room appearance according to the layout, followed by panoramic reconstruction to generate the final 3D textured mesh. This mechanism enables users to interactively edit the scene layout so as to produce customized 3D mesh results. The overall framework of our method is depicted in Fig. 2, which consists of two stages: the Layout Generation Stage and the Appearance Generation Stage. In the Layout Generation Stage, we use a holistic scene code to parametrize the indoor scene and design a diffusion model to learn its distribution. Once the holistic scene code is generated from text, we recover the room as a set of orientated bounding boxes of walls and objects. Note that users can edit these bounding boxes by dragging objects to adjust their semantic types, positions, or scales, enabling the customization of 3D scene results according to the user’s preferences. In the Appearance Generation Stage, we obtain an RGB panorama through a pre-trained latent diffusion model to represent the room texture. Specifically, we project the generated layout bounding boxes into a semantic segmentation map representing the layout. We then fine-tune a pre-trained ControlNet (Zhang & Agrawala, 2023) model to generate an RGB panorama from the input layout panorama. To ensure loop consistency, we propose a novel loop-consistent sampling during the inference process. Finally, the textured 3D mesh is obtained by estimating the depth map of the generated panorama. By separating the scene layout and appearance, our framework supports diverse editing operations on the 3D scene.

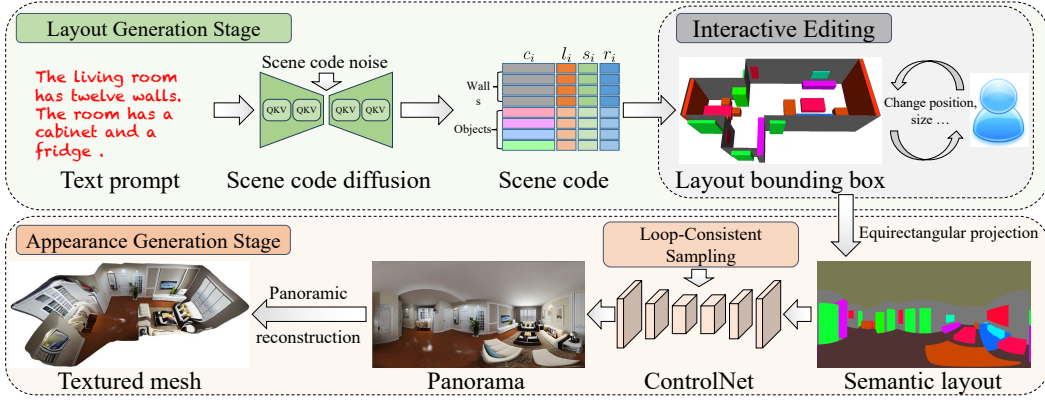


Figure 2: Overview of our method. In the Layout Generation Stage, we synthesize a scene code from the text input and convert it to a 3D bounding box representation to facilitate editing. In the Appearance Generation Stage, we project the bounding boxes into a semantic segmentation map to guide the panorama synthesis. The panorama is then reconstructed into a textured 3D mesh model.




Figure 3: (a) A 3D scene  $S$  is represented by its scene code  $x_0 = \{o_i\}_{i=1}^N$ , where each wall or furniture item  $o_i$  is a row vector storing attributes like class label  $c_i$ , location  $l_i$ , size  $s_i$ , orientation  $r_i$ . (b) During the denoising process, we rotate both the input semantic layout panorama and the denoised image for  $\gamma$  degree at each step. Here we take  $\gamma = 90^\circ$  for example.

### 3.1 LAYOUT GENERATION STAGE

**Scene Code Definition.** Different from previous methods (Paschalidou et al., 2021; Tang et al., 2023a), we consider not only furniture but also walls, doors, and windows to define the room layout. We employ a unified encoding of various objects. Specifically, given a 3D scene  $S$  with  $m$  walls and  $n$  furniture items, we represent the scene layout as a holistic scene code  $\mathbf{x}_0 = \{o_i\}_{i=1}^N$ , where  $N = m + n$ . We encode each object  $o_j$  as a node with various attributes, i.e., center location  $l_i \in \mathbb{R}^3$ , size  $s_i \in \mathbb{R}^3$ , orientation  $r_i \in \mathbb{R}$ , class label  $c_i \in \mathbb{R}^C$ . Each node is characterized by the concatenation of these attributes as  $\mathbf{o}_i = [c_i, l_i, s_i, r_i]$ . As can be seen in Fig. 3 (a), we represent a scene layout as a tensor  $\mathbf{x}_0 \in \mathbb{R}^{N \times D}$ , where  $D$  is the attribute dimension of a node. In all the data, we choose the normal direction of the largest wall as the ‘main direction’. For other objects, we take the angles between their front directions and the main direction as their rotations. We use the one-hot encoding to represent their semantic types, such as sofa or lamp. For more details on the scene encoding, please refer to the Appendix.

**Scene Code Diffusion.** With the scene code definition, we build a diffusion model to learn its distribution. A scene layout is a point in  $\mathbb{R}^{N \times D}$ . The forward diffusion process is a discrete-time Markov chain in  $\mathbb{R}^{N \times D}$ . Given a clean scene code  $\mathbf{x}_0$ , the diffusion process gradually adds Gaussian noise to  $\mathbf{x}_0$ , until the resulting distribution is Gaussian, according to a pre-defined, linearly increased noise schedule  $\beta_1, \dots, \beta_T$ :

$$q(\mathbf{x}_t | \mathbf{x}_0) := \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t} \mathbf{x}_0, (1 - \sqrt{\alpha_t}) \mathbf{I}) \quad (1)$$

where  $\alpha_t := 1 - \beta_t$  and  $\bar{\alpha}_t := \prod_{r=1}^t \alpha_r$  define the noise level and decrease over the timestep  $t$ . A neural network is trained to reverse that process, by minimizing the denoising objective,

$$\mathcal{L} = \mathbf{E}_{\mathbf{x}_0, t, y, \epsilon} \|\epsilon - \epsilon_\theta(x_t, t, y)\|^2, \quad (2)$$

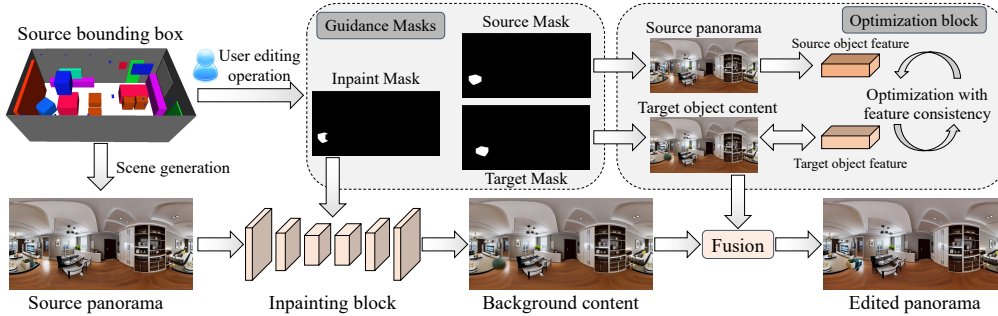


Figure 4: **Mask-guided Editing.** After editing the scene bounding box, we derive guidance masks from the changes in the semantic layout panoramas. We fill in unoccluded regions and optimize the DIFT (Tang et al., 2023b) features to keep the identity of moved objects unchanged.

where  $\epsilon_\theta$  is the noise estimator which aims to find the noise  $\epsilon$  added into the input  $x_0$ . Here,  $y$  is the text embedding of the input text prompts. The denoising network is a 1D UNet (Ronneberger et al., 2015), with multiple self-attention and cross-attention layers designed for input text prompts. The denoising network  $\epsilon_\theta$  takes the scene code  $\mathbf{x}_t$ , text prompt  $y$ , and timestep  $t$  as input, and denoises them iteratively to get a clean scene code  $\hat{\mathbf{x}}_0$ . Then we represent  $\hat{\mathbf{x}}_0$  as a set of orientated bounding boxes of various semantic types to facilitate interactive editing.

### 3.2 APPEARANCE GENERATION STAGE

Given the layout of an indoor scene, we seek to obtain a proper panorama image to represent its appearance. Instead of incrementally generating multi-view images like (Höllerin et al., 2023), we generate the entire panorama at once. We utilize ControlNet (Zhang & Agrawala, 2023) to generate a high-fidelity panorama conditioned by the input 3D scene layout. After getting the scene panorama, we recover the depth map by the method in (Shen et al., 2022) to reconstruct a textured mesh through Poisson reconstruction (Kazhdan et al., 2006) and MVS-texture (Wachter et al., 2014).

**Fine-tune ControlNet.** ControlNet is a refined Stable Diffusion (Rombach et al., 2022) model conditioned on an extra 2D input. To condition ControlNet on the scene layout, we convert the bounding box representation into a 2D semantic layout panorama through equirectangular projection. In this way, we get a pair of RGB and semantic layout panoramic images for each scene. However, the pre-trained ControlNet-Segmentation (github, 2023) is designed for perspective images, and cannot be directly applied to panoramas. Thus, we fine-tune it with our pairwise RGB-Semantic layout panoramas on the Structured3D dataset Zheng et al. (2020). As the volume of Structured3D is limited, we apply several techniques during the fine-tuning to augment the training data, including standard left-right flipping, horizontal rotation, and Pano-Stretch (Sun et al., 2019).

**Loop-consistent Sampling.** A panorama should be loop-consistent. In other words, its left and right should be seamlessly connected. Although the panoramic horizontal rotation in data augmentation may improve the model’s implicit understanding of the expected loop consistency, it lacks explicit constraints and might still produce inconsistent results. Therefore, we propose an explicit loop-consistent sampling mechanism in the denoising process of the latent diffusion model. As shown in Fig. 3(b), we rotate both the input layout panorama and the denoised image for  $\gamma$  degree in the sampling process, which applies explicit constraints for the loop consistency during denoising. A concurrent work (Wu et al., 2023) also uses a similar method for panoramic outpainting. More qualitative results in Appendix Figure.2 verify that our simple loop-consistent sampling method achieves good results without introducing additional learnable parameters.

### 3.3 MASK-GUIDED EDITING

A user can modify the generated 3D room by changing the position, semantic class, and size of object bounding boxes. Our method will update the panorama accordingly to reflect the user’s editing in 3D space, which achieves zero-shot editing without expensive re-training.

The editing should achieve two goals, i.e. altering the content according to the user’s input, and maintaining appearance consistency for scene objects. We propose a mask-guided image editing as

illustrated Fig. 4, where a chair’s position is moved. In the following, we will explain our method with this example. We denote the semantic panorama from the edited scene as  $P_{\text{edited}}$ , then we derive the guidance masks based on its difference from the original one  $P_{\text{ori}}$ . The source mask  $\mathbf{m}_{\text{src}}$  shows the position of the original chair, and the target mask  $\mathbf{m}_{\text{tar}}$  indicates the location of the moved chair, and the inpainting mask  $\mathbf{m}_{\text{inpaint}} = \{m | m \in \mathbf{m}_{\text{src}} \text{ and } m \notin \mathbf{m}_{\text{tar}}\}$  is the unoccluded region. Given these guidance masks, our method includes two steps: the inpainting step and the optimization step. We first fill in the inpaint area by feeding the inpaint mask  $\mathbf{m}_{\text{inpaint}}$  and edited semantic panorama  $P_{\text{edited}}$  to the inpainting step. Then, in our optimization step, we optimize the DIFT (Tang et al., 2023b) feature to maintain the visual consistency of relocated objects.

**Inpainting Step.** Denoting the original image as  $\mathbf{x}_0^{\text{ori}}$ , we replace pixels outside the inpainting mask  $\mathbf{m}_{\text{inpaint}}$  with  $\mathbf{x}_t^{\text{ori}}$  during the diffusion process. This simple strategy keeps the outside region unchanged. At each reverse diffusion step, we compute:

$$\mathbf{x}_t^{\text{ori}} \sim \mathcal{N}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0^{\text{ori}}, (1 - \bar{\alpha}_t \mathbf{I})), \quad (3)$$

$$\mathbf{x}_t^{\text{new}} \sim \mathcal{N}(\mu_\theta(x_t, t, y, P_{\text{edited}}), \Sigma_\theta(x_t, t, y, P_{\text{edited}})), \quad (4)$$

$$\hat{\mathbf{x}}_{t-1}^{\text{new}} = \mathbf{m}_{\text{inpaint}} \odot \mathbf{x}_t^{\text{new}} + (1 - \mathbf{m}_{\text{inpaint}}) \odot \mathbf{x}_t^{\text{ori}}, \quad (5)$$

where  $\mathbf{x}_t^{\text{ori}}$  is obtained through propagating  $\mathbf{x}_0^{\text{ori}}$  in diffusion process, and  $\mathbf{x}_t^{\text{new}}$  is sampled from the fine-tuned ControlNet model, which takes the edited semantic layout panorama  $P_{\text{edited}}$  and text prompt  $y$  as input. As the propagated  $\mathbf{x}_t^{\text{ori}}$  is unaware of the new content  $\mathbf{x}_t^{\text{new}}$ , this may result in distracting boundaries of the inpainted area. To better blend the new content  $\mathbf{x}_t^{\text{new}}$  and its surrounding background  $\mathbf{x}_t^{\text{ori}}$  in the inpainted area, we update the computation of  $\hat{\mathbf{x}}_{t-1}^{\text{new}}$  to,

$$\hat{\mathbf{x}}_{t-1}^{\text{new}} = \mathbf{m}_{\text{inpaint}} \odot \mathbf{x}_t^{\text{new}} + (1 - \mathbf{m}_{\text{inpaint}}) \odot (\mathbf{x}_t^{\text{ori}} \cdot \lambda_{\text{ori}} + \mathbf{x}_{t+1}^{\text{new}} \cdot \lambda_{\text{new}}), \quad (6)$$

where  $\lambda_{\text{ori}}$  and  $\lambda_{\text{new}}$  are hyper-parameters to adjust the weight for fusing the inpainted area and unchanged area. The final result of inpainting is  $\hat{\mathbf{x}}_0^{\text{new}}$ .

**Optimization Step.** When the user moves the position of a furniture item, we need to keep its appearance unchanged before and after the movement. The recent work, DIFT (Tang et al., 2023b), finds the learned features from the diffusion network allow for strong semantic correspondence. Thus, we maintain the consistency between the original and moved furniture by requiring their latent features to be consistent. In particular, we extract latent features  $F_t^l$  of the layer  $l$  in the denoising U-Net network, at timestep  $t$ . Then we construct a loss function using the latent features from source area  $\mathbf{m}_{\text{src}}$  in source panorama  $\mathbf{x}_0^{\text{ori}}$  and target area  $\mathbf{m}_{\text{tar}}$  in inpainted panorama  $\hat{\mathbf{x}}_0^{\text{new}}$ .

For conciseness, we denote the target image  $\hat{\mathbf{x}}_0^{\text{edit}}$  initialized by  $\hat{\mathbf{x}}_0^{\text{new}}$ . We first propagate the original image  $\mathbf{x}_0^{\text{ori}}$  and  $\hat{\mathbf{x}}_0^{\text{edit}}$  to get  $\mathbf{x}_t^{\text{ori}}$  and  $\hat{\mathbf{x}}_t^{\text{edit}}$  at timestep  $t$  by diffusion process, respectively. At each iteration, we use the same ControlNet model to denoise both  $\mathbf{x}_t^{\text{ori}}$  and  $\hat{\mathbf{x}}_t^{\text{edit}}$  and extract the latent features of them, denoted as  $F_t^{\text{ori}}$  and  $F_t^{\text{edit}}$ , respectively. Based on the strong correspondence between the features, the source mask area  $\mathbf{m}_{\text{src}}$  and the target area  $\mathbf{m}_{\text{tar}}$  in  $F_t^{\text{ori}}$  and  $F_t^{\text{edit}}$  need to have high similarity. Here, we utilize the cosine embedding loss to measure the similarity, and define the optimization loss function as follows:

$$\mathcal{L}_{\text{obj}} = -\cos(\text{sg}(F_t^{\text{ori}} \odot \mathbf{m}_{\text{src}}), F_t^{\text{edit}} \odot \mathbf{m}_{\text{tar}}). \quad (7)$$

Here,  $\text{sg}$  is the stop gradient operator, the gradient will not be back-propagated for the term  $\text{sg}(F_t^{\text{ori}} \odot \mathbf{m}_{\text{src}})$ . Then we minimize the loss iteratively. At each iteration,  $\hat{\mathbf{x}}_t^{\text{edit}}$  is updated by taking one gradient descent step with a learning rate  $\eta$  to minimize the loss  $\mathcal{L}_{\text{obj}}$  as,

$$\hat{\mathbf{x}}_t^{k+1} = \hat{\mathbf{x}}_t^k - \eta \cdot \frac{\partial \mathcal{L}_{\text{obj}}}{\partial \hat{\mathbf{x}}_t^k}. \quad (8)$$

After  $M$  steps optimization, we apply the standard denoising process to get the final result  $\hat{\mathbf{x}}_0^{\text{edit}}$ .

## 4 EXPERIMENTS

We validate our method with experiments and compare it with previous methods on indoor scene generation in accuracy, speed, and result quality. We further show various scene editing results to demonstrate the flexible control of our method.

Table 1: Quantitative Comparison of panorama and mesh generation.

Method	Panorama Metrics				2D Rendering Metrics		3D Mesh User Study	
	FID ↓	CS ↑	IS ↑	Time/s ↓	CS ↑	IS ↑	PQ ↑	3DS ↑
Text2Light (Chen et al., 2022)	56.22	21.45	<b>4.198</b>	81.56	-	-	2.732	2.747
MVDiffusion (Tang et al., 2023c)	34.76	<b>23.93</b>	3.21	208.5	-	-	3.27	3.437
Text2Room (Höllein et al., 2023)	-	-	-	≥ 9,000	25.90	2.90	2.487	2.588
Ours	<b>21.02</b>	21.58	3.429	<b>61.1</b>	<b>25.97</b>	<b>3.14</b>	<b>3.89</b>	<b>3.746</b>



Figure 5: Qualitative comparison with previous works. For each method, we show a textured 3D mesh in the first row and a few rendered perspective images in the second row.

#### 4.1 EXPERIMENT SETUP


**Dataset:** We experiment on the 3D indoor scene dataset Structured3D (Zheng et al., 2020), which consists of 3,500 houses with 21,773 rooms designed by professional artists. Photo-realistic rendered images, such as RGB panorama, semantic panorama, depth map, and normal map are provided in each room. Each room is also labeled with the room layout, represented by the intersection lines of walls, roofs, and floors. 3D bounding boxes of furniture items are provided but noisy. We parse these boxes for all the ‘living rooms’ and ‘bedrooms’. Then, we follow (Wang et al., 2021) to generate text prompts describing the scene layout. Please refer to Appendix Sec.1.3 for more details about data preprocessing. The filtered dataset for training and evaluation includes 4,961 bedrooms and 3,039 living rooms. For each room type, we use 80% of rooms for training and the remaining for testing.

**Metrics:** We choose Frechet Inception Distance (FID) (Heusel et al., 2017), CLIP Score (CS) (Radford et al., 2021), and Inception Score (IS) (Salimans et al., 2016) to measure the image quality of generated panoramas. We also compare the time cost to synthesize an RGB panorama of size  $512 \times 1024$ . To compare the quality of 3D room models, we follow Text2Room (Höllein et al., 2023) to render images of the 3D room model and measure the CLIP Score (CS) and Inception Score (IS). We further conduct a user study and ask 61 users to score Perceptual Quality (PQ) and 3D Structure Completeness (3DS) of the final room mesh on scores ranging from 1 to 5.

#### 4.2 COMPARISON WITH PREVIOUS METHODS

**Quantitative Comparison** To evaluate our generated panoramic images, we follow MVDiffusion (Tang et al., 2023c) to crop perspective images from the 1,181 generated panoramas on the test split and evaluate the FID, CS, IS on the cropped multi-view images. These quantitative results are summarized on the left of Table 1, where our method achieves the best score in FID, which indicates that our method can better capture the room appearance because of its faithful recovery of the room layout. However, our score on CS is slightly lower than MVDiffusion, which seems to be insensitive to the number of objects as we illustrated in the supplementary file, and cannot reflect the quality of room layouts. The IS score depends on the semantic diversity of the cropped images as captured by an image classifier. It turns out Text2Light has the best IS score, since its generated indoor scenes often contain unexpected objects. More panoramic visualization can be found in Appendix Sec.1.5.





(a) Resize the TV (b) Replace the chair by a new one  
 Figure 6: Room editing examples. (a) resize the TV, (b) replace the chair with a new one.

In terms of running time, our method takes the shortest time. Averagely speaking, our method takes only 61 seconds to generate a panorama, and another 20 seconds to generate the textured 3D mesh. In comparison, MVDiffusion takes 208 seconds, about 3.5 times longer, to synthesize a panorama. Text2Room needs at least 150 minutes to finish a textured 3D room generation.

We then compare the 3D room models in terms of their rendered images. Because of the expensive running time of Text2Room (Höllerle et al., 2023), we only test on 12 examples for this comparison. In this comparison, we further skip Text2Light and MVDiffusion since we have compared them on panoramas. As the room layout is better captured with a large FOV, we render 60 perspective images of each scene with a  $140^\circ$  FOV and evaluate their CS and IS scores respectively. Please refer to Appendix Sec.1.4.1 for more details. The results of this comparison are shown in the middle of Table 1. Our method obtains better scores on both metrics than Text2Room.

We further evaluate the quality of the textured 3D mesh model by user studies. For those panorama generation methods, we utilize the depth estimation work (Shen et al., 2022) to reconstruct a textured 3D room mesh. More details can be found in the Appendix. The results of the user study are shown on the right of Table 1. Users prefer our method over others, for its clear room layout structure and furniture arrangement.

**Qualitative Comparison** Fig. 5 shows some results generated by different methods. The first row shows a textured 3D room model, and the second row shows some cropped perspective images from the panorama. As we can see, Text2Light (Chen et al., 2022) cannot generate a reasonable 3D indoor scene. It even fails to ensure the loop consistency of the generated panorama, which leads to distorted geometry and room model. Both MVDiffusion (Tang et al., 2023c) and Text2Room Höllerle et al. (2023) can generate vivid local images as demonstrated by the perspective renderings in the second row. But they fail to capture the more global scale room layout. Similar effects can be seen from the Fig. 1 (a). These two methods often repeat a dominating object, e.g. a bed in the bedroom or fireplace in the living room, multiple times at different places and violate the room layout constraint. In comparison, our method does not suffer from these problems and generates high-quality results. More examples are provided in the supplementary file.

### 4.3 INTERACTIVE SCENE EDITING

We demonstrate the scene editing capability of our method in Fig. 6. In this case, we resize the TV and replace the chair in the generated results. Fig. 1 (b) shows two additional examples of replacing the TV and TV stand and moving the picture upwards. Our method can keep the visual appearance of the moved/resized objects unchanged after editing. More examples can be found in the supplementary file.

## 5 CONCLUSION AND DISCUSSION

We present **Ctrl-Room**, a flexible method to achieve editable and structurally plausible 3D indoor scene generation. It consists of two stages, the layout generation stage and the appearance generation stage. In the layout generation stage, we design a scene code to parameterize the scene layout and learn a text-conditioned diffusion model for text-driven layout generation. In the appearance generation stage, we fine-tune a ControlNet model to generate a vivid panorama image of the room

---

with the guidance of the layout. Finally, a high-quality 3D room with a structurally plausible layout and realistic textures can be generated. We conduct extensive experiments to demonstrate that **Ctrl-Room** outperforms existing methods for 3D indoor scene generation both qualitatively and quantitatively, and supports interactive 3D scene editing.

## REFERENCES

- Miguel Angel Bautista, Pengsheng Guo, Samira Abnar, Walter Talbott, Alexander Toshev, Zhuoyuan Chen, Laurent Dinh, Shuangfei Zhai, Hanlin Goh, Daniel Ulbricht, et al. Gaudi: A neural architect for immersive 3d scene generation. *Proc. NeurIPS*, 35:25102–25116, 2022.
- Kevin Chen, Christopher B Choy, Manolis Savva, Angel X Chang, Thomas Funkhouser, and Silvio Savarese. Text2shape: Generating shapes from natural language by learning joint embeddings. In *Proc. ACCV*, pp. 100–116. Springer, 2019.
- Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023.
- Zhaoxi Chen, Guangcong Wang, and Ziwei Liu. Text2light: Zero-shot text-driven hdr panorama generation. *ACM Trans. Graphics*, 41(6):1–16, 2022.
- github. Controlnetgithubmodel. <https://github.com/lllyasviel/ControlNet-v1-1-nightly#controlnet-11-segmentation>, 2023.
- Kamal Gupta, Justin Lazarow, Alessandro Achille, Larry S Davis, Vijay Mahadevan, and Abhinav Shrivastava. Layouttransformer: Layout generation and completion with self-attention. In *Proc. ICCV*, pp. 1004–1014, 2021.
- Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Proc. NeurIPS*, 30, 2017.
- 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.
- Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, and Greg Mori. Layoutvae: Stochastic scene layout generation from a label set. In *Proc. ICCV*, pp. 9895–9904, 2019.
- Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proc. CVPR*, pp. 6007–6017, 2023.
- Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, pp. 0, 2006.
- Jianan Li, Jimei Yang, Aaron Hertzmann, Jianming Zhang, and Tingfa Xu. Layoutgan: Generating graphic layouts with wireframe discriminators. *arXiv preprint arXiv:1901.06767*, 2019.
- Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Proc. CVPR*, pp. 300–309, 2023.
- Chieh Hubert Lin, Chia-Che Chang, Yu-Sheng Chen, Da-Cheng Juan, Wei Wei, and Hwann-Tzong Chen. Coco-gan: Generation by parts via conditional coordinating. In *Proc. ICCV*, pp. 4512–4521, 2019.
- Chieh Hubert Lin, Hsin-Ying Lee, Yen-Chi Cheng, Sergey Tulyakov, and Ming-Hsuan Yang. Infinitygan: Towards infinite-pixel image synthesis. *arXiv preprint arXiv:2104.03963*, 2021.

- 
- Jiafeng Mao, Xueting Wang, and Kiyoharu Aizawa. Guided image synthesis via initial image editing in diffusion model. *arXiv preprint arXiv:2305.03382*, 2023.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- Chong Mou, Xintao Wang, Jiechong Song, Ying Shan, and Jian Zhang. Dragondiffusion: Enabling drag-style manipulation on diffusion models. *arXiv preprint arXiv:2307.02421*, 2023.
- Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022.
- Despoina Paschalidou, Amlan Kar, Maria Shugrina, Karsten Kreis, Andreas Geiger, and Sanja Fidler. Atiss: Autoregressive transformers for indoor scene synthesis. *Proc. NeurIPS*, 34:12013–12026, 2021.
- Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proc. ICML*, pp. 8748–8763. PMLR, 2021.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proc. CVPR*, pp. 10684–10695, 2022.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pp. 234–241. Springer, 2015.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Proc. NeurIPS*, 35:36479–36494, 2022.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.
- Junyoung Seo, Wooseok Jang, Min-Seop Kwak, Jaehoon Ko, Hyeonsu Kim, Junho Kim, Jin-Hwa Kim, Jiyoung Lee, and Seungryong Kim. Let 2d diffusion model know 3d-consistency for robust text-to-3d generation. *arXiv preprint arXiv:2303.07937*, 2023.
- Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Proc. NeurIPS*, 34:6087–6101, 2021.
- Zhijie Shen, Chunyu Lin, Kang Liao, Lang Nie, Zishuo Zheng, and Yao Zhao. Panoformer: Panorama transformer for indoor 360 depth estimation. In *Proc. ECCV*, pp. 195–211. Springer, 2022.
- Yujun Shi, Chuhui Xue, Jiachun Pan, Wenqing Zhang, Vincent YF Tan, and Song Bai. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. *arXiv preprint arXiv:2306.14435*, 2023.
- Ka Chun Shum, Hong-Wing Pang, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Conditional 360-degree image synthesis for immersive indoor scene decoration. *arXiv preprint arXiv:2307.09621*, 2023.
- Cheng Sun, Chi-Wei Hsiao, Min Sun, and Hwann-Tzong Chen. Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation. In *Proc. CVPR*, pp. 1047–1056, 2019.

- 
- Jiapeng Tang, Yinyu Nie, Lev Markhasin, Angela Dai, Justus Thies, and Matthias Nießner. Diffuscene: Scene graph denoising diffusion probabilistic model for generative indoor scene synthesis. *arXiv preprint arXiv:2303.14207*, 2023a.
- Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent correspondence from image diffusion. *arXiv preprint arXiv:2306.03881*, 2023b.
- Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. *arXiv preprint arXiv:2307.01097*, 2023c.
- Michael Waechter, Nils Moehrle, and Michael Goesele. Let there be color! large-scale texturing of 3d reconstructions. In *Proc. ECCV*, pp. 836–850. Springer, 2014.
- Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proc. CVPR*, pp. 12619–12629, 2023a.
- Xinpeng Wang, Chandan Yeshwanth, and Matthias Nießner. Sceneformer: Indoor scene generation with transformers. In *2021 International Conference on 3D Vision (3DV)*, pp. 106–115. IEEE, 2021.
- Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolific-dreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv preprint arXiv:2305.16213*, 2023b.
- Tianhao Wu, Chuanxia Zheng, and Tat-Jen Cham. Ipo-ldm: Depth-aided 360-degree indoor rgb panorama outpainting via latent diffusion model. *arXiv preprint arXiv:2307.03177*, 2023.
- Jingbo Zhang, Xiaoyu Li, Ziyu Wan, Can Wang, and Jing Liao. Text2nerf: Text-driven 3d scene generation with neural radiance fields. *arXiv preprint arXiv:2305.11588*, 2023.
- Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023.
- Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3d: A large photo-realistic dataset for structured 3d modeling. In *Proc. ECCV*, pp. 519–535. Springer, 2020.

---

# CTRL-ROOM: CONTROLLABLE TEXT-TO-3D ROOM MESHES GENERATION WITH LAYOUT CONSTRAINTS (SUPPLEMENTARY MATERIALS)

**Anonymous authors**

Paper under double-blind review

## 1 APPENDIX

In the appendix, we first present more details about our scene code diffusion model in Sec. 1.1, then we provide our dataset pre-processing, text prompt generation, and implementation details in Sec. 1.2 and Sec. 1.3 respectively. Additional experiment results are illustrated in Sec. 1.4 and user studies in Sec. 1.6. Finally, we clarify the limitations of our method.

### 1.1 SCENE CODE DENOISING NETWORK

As shown in Fig. 1, our scene code denoising network is built upon IDDPM (?). The whole architecture is similar to IDDPM, while we replace the upsample and downsample blocks with 1D-convolution network in the U-Net, and insert attention blocks after each residual block to capture both the global context among objects and the semantic context from the input text prompt. The input encoding head processes different encoding of the node attributes, e.g., semantic class labels, box centroid, and box orientation. After adding noise, the input encoding is fed into the U-Net to obtain a denoised scene code. During the forward phase, as in IDDPM, we iteratively perform the denoising process and generate a scene code from a partial scene textual description.

### 1.2 DATASET

**Structured3D dataset preprocessing** Structured3D consists of 3,500 houses with 21,773 rooms, where each room is designed by professional designers with rich 3D structure annotations, including the room planes, lines, junctions, and orientated bounding box of most furniture, and photo-realistic 2D renderings of the room. In our work, we use the 3D orientated bounding boxes of furniture, 2D RGB panorama, and 3D lines and planes of each room. While the original dataset lacks semantic class labels for each furniture bounding box. The dataset preprocessing aims to produce clean ground truth data for our layout generation module and appearance generation module.

- **Orientated Object Bounding Box Annotation.** As the original dataset lacks semantic label for each orientated object bounding box, we first unproject the RGB panorama and depth map into a point cloud of the room, then manually annotate the object semantic class and add more accurate object bounding boxes based on the noisy annotation of the original version. As shown in Fig. 2, we annotated the orientated object bounding boxes through labelCloud (?) on 4961 bedrooms and 3039 livingrooms. To ensure high-quality annotations, the object bounding boxes of the same scene are annotated by two annotators and verified by a third annotator.
- **Scene Node Encoding.** We define our holistic scene code based on a unified encoding of walls and object bounding box. Each object  $o_j$  is treated as a node with various attributes, i.e., center location  $l_i \in \mathbb{R}^3$ , size  $s_i \in \mathbb{R}^3$ , orientation  $r_i \in \mathbb{R}$ , class label  $c_i \in \mathbb{R}^C$ . The orientated bounding box is off-the-shelf, we extract the inner walls based on the line junctions and corners of the 3D room. Then we put the orientated object bounding boxes and walls into a compact scene code. Concretely, we define an additional 'empty' object and pad it into scenes to have a fixed number of object across scenes. Each object rotation angle is parametrized by a 2-d vector of cosine and sine values. Finally, each node is characterized by the concatenation of these attributes as  $\mathbf{o}_i = [c_i, l_i, s_i, \cos r_i, \sin r_i]$ .

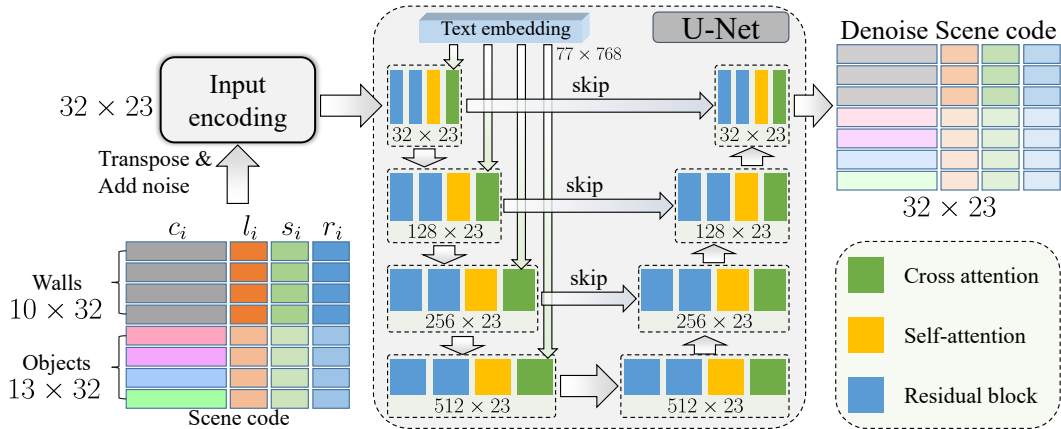


Figure 1: The detailed structure of the scene code denoising network. We here take the bedroom for example to demonstrate the scene code denoiser’s dataflow. The scene code tensor  $\mathbf{x}_0 \in \mathbb{R}^{N \times D}$ , where  $N = 23, D = 32$ .



Figure 2: Example of object bounding box annotation.

- **data filtering.** We start by filtering out those problematic scenes such as zooms with wall number less than 4 or larger than 24. We also remove those scenes with too few or too many objects. The number of walls of valid bedrooms is between 4 and 10, and that of objects is between 3 and 13. As for living rooms, the minimum and maximum numbers of walls are set to 4 and 24, and that of objects are set to 3 and 21 respectively. Thus, the number of scene nodes is  $N = 23$  in bedrooms and  $N = 45$  in living rooms.

**Text Prompt Generation** We follow the SceneFormer ? to generate text prompts describing partial scene configurations. Each text prompt contains two to four sentences. The first sentence describes how many walls are in the room, then the second sentence describes two or three existing furniture in the room. The following sentences mainly describe the spatial relations among the furniture, please refer to SceneFormer ? and DiffuScene ? for more detailed explanation of relation-describing sentences. In this way, we can get some relation-describing sentences to depict the partial scene. Finally, we randomly sampled zero to two relation-describing sentences to form the text prompt for 3D room generation.

### 1.3 IMPLEMENTATION DETAILS

#### Training and inference details.

- In the layout generation stage, We train the scene code diffusion model on our processed ‘bedroom’ and ‘livingroom’ data of Structured3D (?) for 200,000 steps. The frozen text encoder we adopted is the same as Stable Diffusion (?). The training is performed using the AdamW optimizer with a batch size of 128 and a learning rate of  $1e-4$ , utilizing 2 A6000 GPUs. During the inference process, we utilize the DDIM (?) sampler with a step size of 250 to perform scene code denoising.
- In the appearance generation stage, we fine-tune the segmentation-conditional ControlNet model based on the pairwise semantic and RGB panorama of Structured3D. The fine-tuning process is implemented on two A6000 GPUs for 150 epochs(about 3 days). In the inference phase, we generate high-fidelity and loop-consistent RGB panorama through DDIM sampler with 100 steps, rotating both semantic layout panorama and the denoised image for  $\gamma = 90^\circ$  at each step.
- As for the mask-guided editing module, we utilize the fine-tuned Control-Seg model to inpaint the background content and optimize the latents of the edited panorama. In inpainting step, the weights used to fuse the unpainted area and unchanged area are set  $\lambda_{ori} = 0.8, \lambda_{new} = 0.2$ . In the optimization step, the maximum iteration is  $M = 50$ , the learning rate  $\eta$  for optimization is initialized to 0.1 and then gradually decreases to 0.01.

### 1.3.1 BASELINE IMPLEMENTATIONS

We provide implementation details for baseline methods in the following:

- MVDiffusion (?): To get a high-resolution photo realistic panorama, MVDiffusion employs 8 branches of SD (?) model and correspondence-aware attention mechanism to generate multi-view images simultaneously. We first fine-tune the pre-trained model of MVDiffusion on Structured3D for 10 epochs(about 3 days). Since each generated subview image of MVDiffusion is at  $512 \times 512$  resolution, the final panorama is pretty large. We resize the generated panorama of MVDiffusion from  $4096 \times 2048$  to  $1024 \times 512$ . Then the 8 sub-view perspective images are extracted from the post-processed panorama using the same camera settings (FOV=90°,rotation=45°). The same operation is adopted on our generated panoramic images. Finally, we combine the panorama from MVDiffusion with our panoramic reconstruction module to create a 3D mesh.
- Text2Light (?): Text2Light creates HDR panoramic images from text using a multi-stage auto-regressive generative model. We choose Text2Light as one of the baseline for our panorama generation and 3D room mesh generation. We first generate RGB panoramas from the input text using Text2Light, then lift it into 3D mesh using the same panoramic reconstruction module as our method. When evaluate the panoramic image quality, we adopt the same processing as MVDiffusion to get multi-view perspective images of Text2Light.
- Text2Room (?): Text2Room is the current state-of-the-art and off-the-shelf method for 3D room mesh generation. It utilizes 20 camera spots of a pre-defined trajectory to expand new areas as much as possible by generating 10 images at each spot. Here We use its final fused poison mesh for 3D mesh comparison. For a fair comparison of 2D renderings evaluation, we only use the renderings at the origin of the final mesh.

## 1.4 PANORAMA GENERATION COMPARISON

Fig. 3 presents additional results for panorama generation. Given a simple partial-scene text prompt, our approach obtains better RGB panorama than that of Text2Light (?) and MVDiffusion (?), which demonstrates the effectiveness of our well-designed framework. While Text2Light suffers from the inconsistent loop and unexpected content of the generated panorama, MVDiffusion fails to recover a reasonable room layout from the text prompt.

## 1.5 ADDITIONAL QUALITATIVE RESULTS

We show additional qualitative comparison results between our method and baselines in Fig. 5. We demonstrate more scene editing results of our method in Fig. 4.

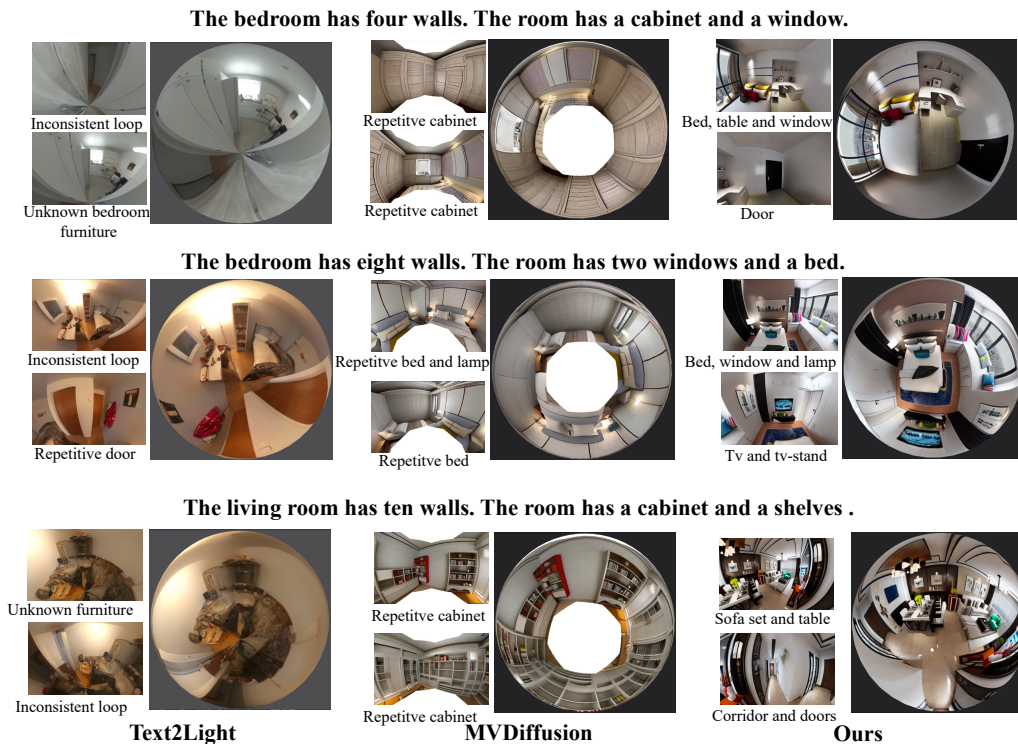


Figure 3: Qualitative comparison for panorama generation. Generated panorama is visualized in a panoramic image viewer to facilitate the user to check the global content of panorama. The left side of each column is two zoom-in views, and the right side is the fisheye view. Text2Light (?) exists serious inconsistent problem on the border of the generated panorama, it also shows a lot of unexpected stuff in the image. MVDiffusion (?) fails to synthesize reasonable content for the target room type. In contrast, our method obtains layout plausible and vivid panorama from the given text prompt of partial scene.

## 1.6 USER STUDY

Follow Text2Room ?, we conduct a user study and ask  $n = 61$  ordinary users to score the Perceptual Quality(PQ) and 3D Structure Completeness(3DS) of the generated room on a scale of 1 – 5. Different from Text2room which only demonstrates the perspective renderings of the 3D room, we directly show users the generated mesh to get a global evaluation of the whole generated 3D room. We show an example of the presented interface of the user study in Fig. 6. In total, we presented 40 top-down views from 10 scenes and report averaged results for each method. Users favor our approach, which emphasizes the superiority of our more plausible geometry, along with the vivid texture.

## 1.7 LIMITATION

Although we have shown impressive 3D room generation results, there are still some limitations in our method. Firstly, we only support single-room generation, thus we cannot produce large-scale indoor scenes with multiple rooms. Secondly, the generated 3D room still contains incomplete structures in invisible area. Users might observe obviously stretched texture because of the occlusion and poor performance of the panoramic depth estimator. It will further affect the visual quality of scene editing results. For example, in Fig. 4(b), we find the table after movement is distorted and the tabletop becomes inclined, the problem is mainly caused by the inaccurate depth estimation and the next mesh reconstruction and mesh texturing process. A promising direction is to learn a text-driven diffusion model to produce one or more RGB-D panorama images under the scene layout constraints. We leave these mentioned limitations as our future efforts.



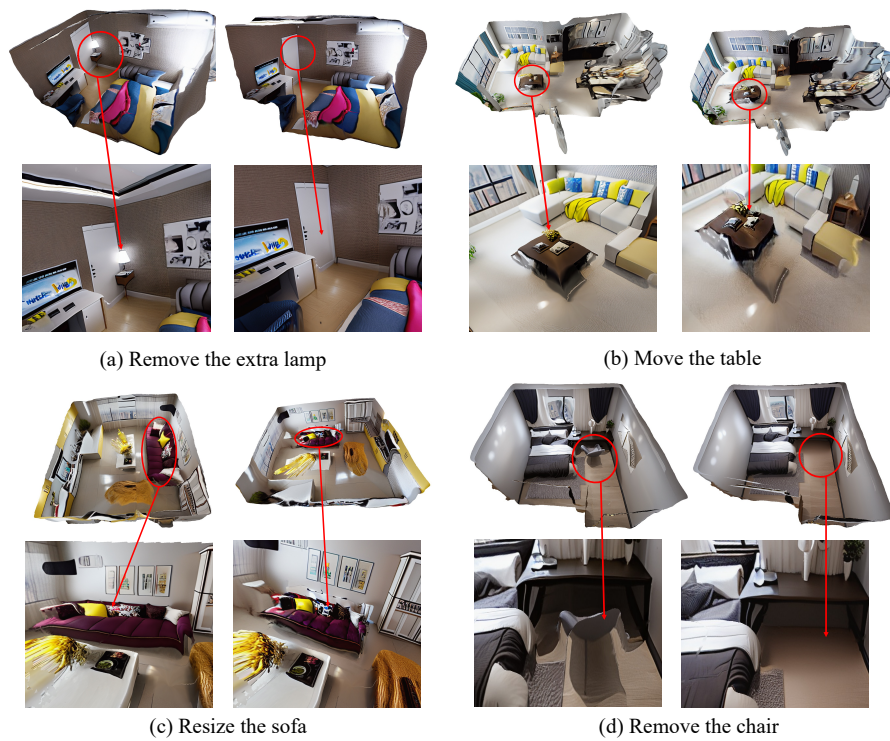


Figure 4: Additional scene editing results. In each sub-figure, the left part is the original 3D room, the right part shows the final mesh after users' interactive editing.

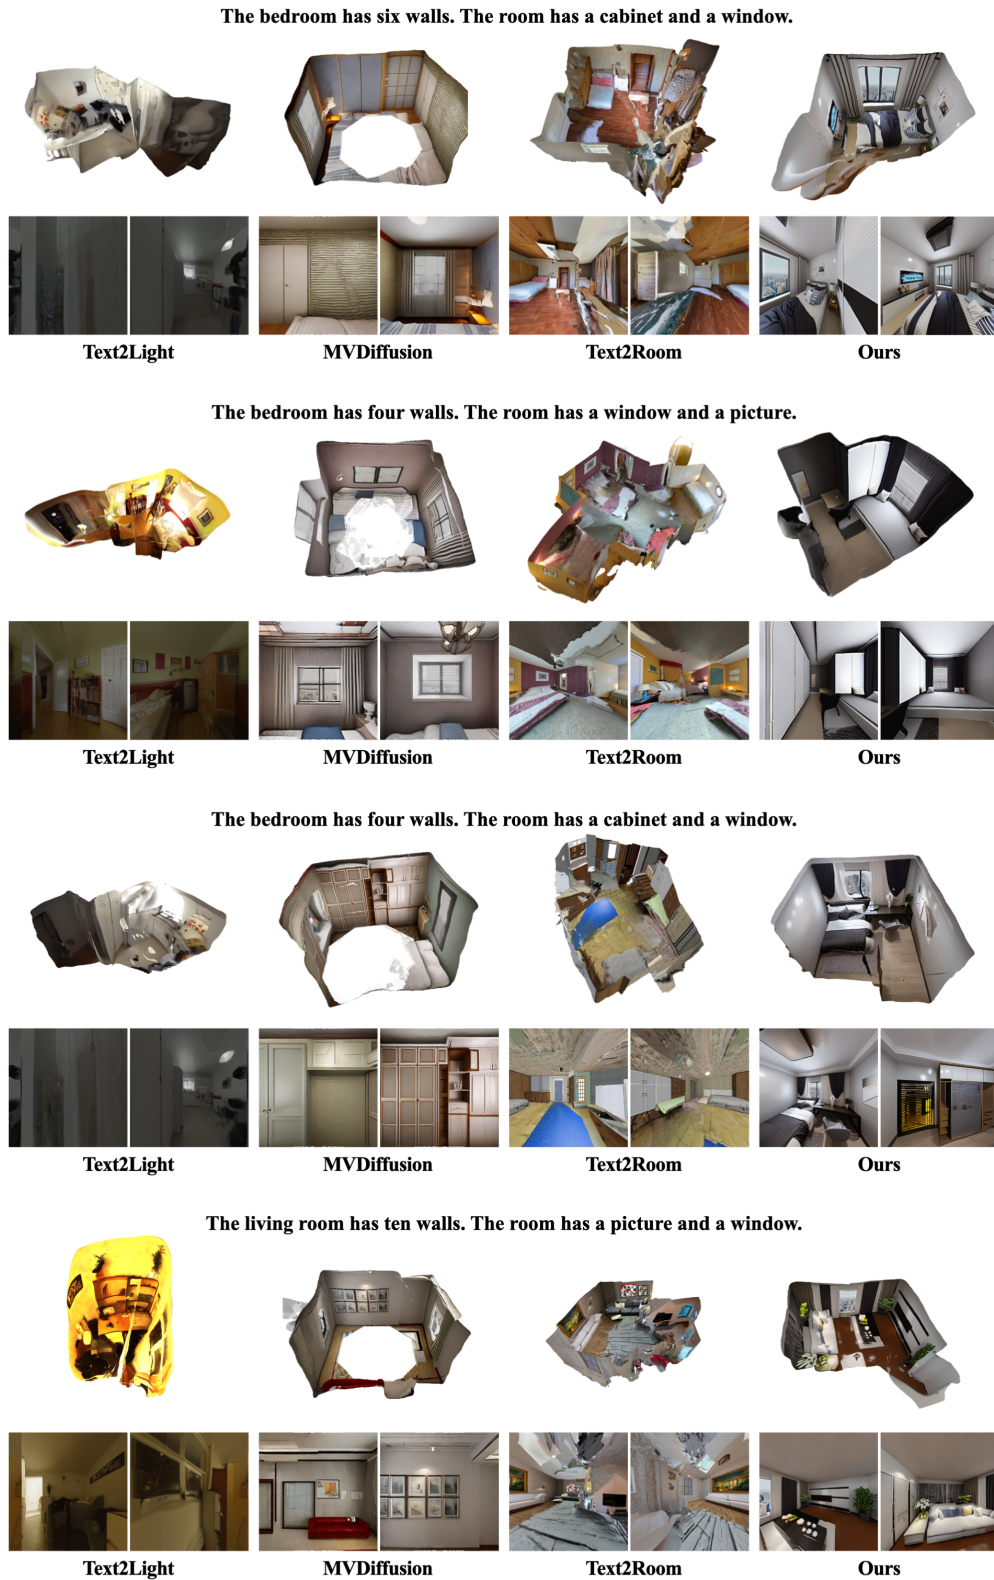


Figure 5: Additional qualitative comparison with previous works.

请给以下4个例子的 a.房屋质量, b.房屋三维结构完整性打分1-5;分数越高表示越好。  
 Given the 4 examples, Please rate the following cases by the two metrics: a. House Perceptual Quality, b. House 3D Structural Completeness. The scale of score is range from 1 to 5; The higher score means better PQ and 3DS.


房屋质量评价准则是当前房屋的纹理图案是否满足用户对这类房间的正常印象, 由用户主观判断;  
 The evaluation criterion of Perceptual Quality is whether the texture pattern of the current room meets the common sense impression of the user on the type of room, which is judged subjectively by the user.

房屋三维结构完整性评价准则是当前房屋三维结构是否符合常理, 是否能清晰判断室内家具物体, 是否存在破损、结构拉伸, 由用户主观判断;  
 The evaluation criterion of the 3D Structure Completeness of the room is whether the three-dimensional structure of the current room is in line with common sense, whether it can clearly judge the indoor furniture objects, whether there are holes and structural stretching, which is judged by the user's subjective judgment.

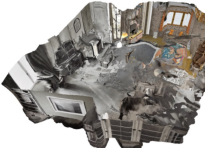
客厅有十二面墙。房间里有一个柜子和一扇窗户。柜子上方有一盏灯。  
 The living room has twelve walls. There is a cabinet and a window in the room. There is a lamp above the cabinet.

	1	2	3	4	5
(A)图-房屋质量-PQ	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
(B)图-房屋质量-PQ	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
(C)图-房屋质量-PQ	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
(D)图-房屋质量-PQ	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
(A)图-三维结构完整性-3DS	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
(B)图-三维结构完整性-3DS	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
(C)图-三维结构完整性-3DS	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
(D)图-三维结构完整性-3DS	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

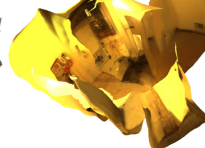
**A**



**B**



**C**



**D**




Figure 6: User study interface. We provide users with multiple top-down images from different methods and ask users to rate the given 3D meshes on a scale from 1 to 5, according to the criteria of Perceptual Quality and 3D Structure Completeness.