

FreeScene: Mixed Graph Diffusion for 3D Scene Synthesis from Free Prompts

Tongyuan Bai¹ Wangyuanfan Bai¹ Dong Chen¹ Tieru Wu^{1,3} Manyi Li² Rui Ma^{1,3*}
¹School of Artificial Intelligence, Jilin University
²School of Software, Shandong University
³Engineering Research Center of Knowledge-Driven Human-Machine Intelligence, MOE, China

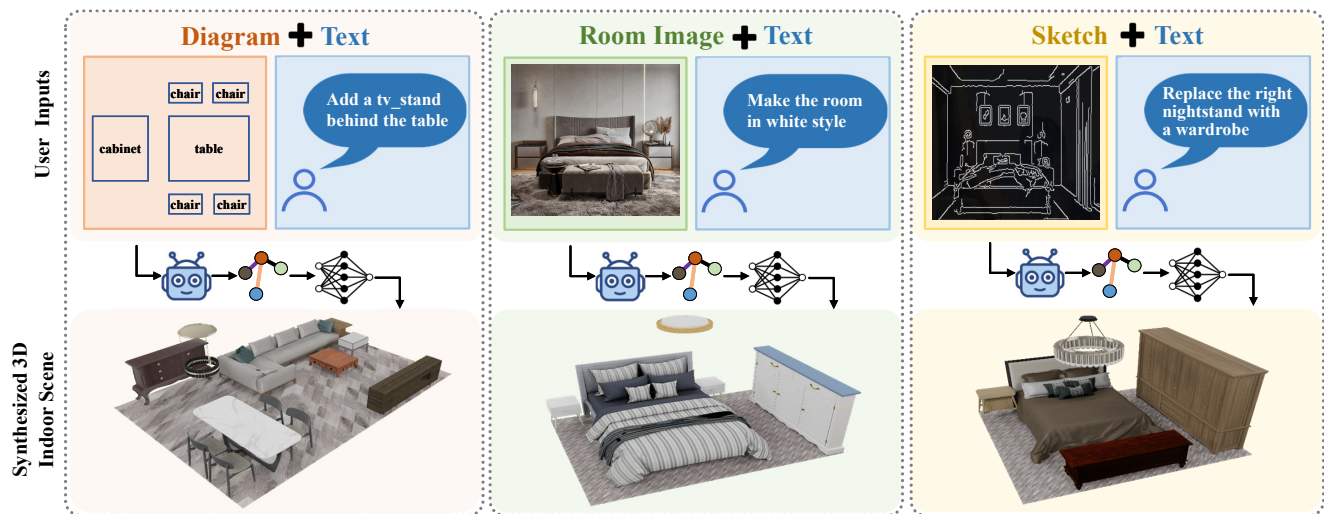


Figure 1. We introduce FreeScene, a user-friendly controllable indoor scene synthesis framework that allows convenient and effective control with free-form user inputs, including text and/or different types of images (top-view diagram in the left, realistic photograph in the middle, sketch in the right). Leveraging a multimodal agent (🤖) to extract partial graph priors (📊) from user inputs, FreeScene subsequently generates a complete and reasonable scene that conforms to the graph via a generative diffusion model (⚙️).

Abstract

Controllability plays a crucial role in the practical applications of 3D indoor scene synthesis. Existing works either allow rough language-based control, that is convenient but lacks fine-grained scene customization, or employ graph-based control, which offers better controllability but demands considerable knowledge for the cumbersome graph design process. To address these challenges, we present **FreeScene**, a user-friendly framework that enables both convenient and effective control for indoor scene synthesis. Specifically, FreeScene supports free-form user inputs including text description and/or reference images, allowing users to express versatile design intentions. The user inputs are adequately analyzed and integrated into a graph representation by a VLM-based Graph Designer. We then propose MG-DiT, a Mixed Graph Diffusion Transformer, which performs graph-aware denoising to enhance scene gener-

ation. Our MG-DiT not only excels at preserving graph structure but also offers broad applicability to various tasks, including, but not limited to, text-to-scene, graph-to-scene, and rearrangement, all within a single model. Extensive experiments demonstrate that FreeScene provides an efficient and user-friendly solution that unifies text-based and graph-based scene synthesis, outperforming state-of-the-art methods in terms of both generation quality and controllability in a range of applications.

1. Introduction

Indoor scene synthesis has demonstrated immense applications across various domains such as game design, VR/AR, and robotics, where the goal is to produce reasonable and plausible scenes that meet user requirements such as text description, sketches, images, etc. To date, there has been a significant body of related works, utilizing either graphical probabilistic models [9, 22, 31, 34, 56, 62, 63, 66] or deep

*Corresponding author

generative models [14, 26, 36, 42, 52, 53, 58, 59, 67, 68] to learn the data distribution of indoor scenes. Recently, as the diffusion model [17, 47, 48] has shown excellent generative performance compared to VAE [23] and GANs [15], it is also applied to indoor scene synthesis by iteratively denoising the scene representations, enabling the generation of plausible and high-quality scene results [18, 33, 49, 60].

However, controllability, as a critical factor for user engagement, has yet to be adequately addressed to provide both convenient and precise control over indoor scene synthesis with diffusion models. The recent DiffuScene [49] incorporates text control through a cross-attention mechanism on U-Net [43], but only obtains a rough consistency between the text conditions and generated scenes without precise control. InstructScene [28] introduces a discrete graph diffusion model to convert text conditions into scene graphs which offer fine-grained control. Although the scene graphs align better with the generated scenes, their text-to-graph requires users to provide a detailed description of the objects and their relations, while still being prone to information loss.

In this paper, we introduce FreeScene, an indoor scene synthesis framework with convenient and effective control. We allow users to express their thoughts freely with a mixture of text descriptions and images, such as top-view diagrams, realistic photos, and sketches. Our framework includes a VLM-based Graph Designer, which conducts step-wise multimodal understanding and reasoning to infer object arrangement as scene graphs, and a mixed graph diffusion transformer model coined as MG-DiT, which refines the graph to generate complete and reasonable scene configurations. Specifically, MG-DiT uses constrained sampling to preserve input conditions throughout the denoising process, enabling a variety of applications such as text-to-scene and graph-to-scene, with a single model. To the best of our knowledge, this is the first approach capable of addressing all these tasks with a unified diffusion model. Extensive experiments demonstrate that FreeScene outperforms state-of-the-art approaches in both generation quality and controllability across many applications.

Our contributions are summarized as follows:

- We introduce the FreeScene framework, which utilizes a graph designer capable of interpreting the free-form user inputs as graphs and a diffusion model to generate results for the user-friendly controllable scene synthesis.
- We propose the MG-DiT model, which employs a mixed graph diffusion process with constrained sampling to enable a variety of applications with a single model.
- We conduct comparison experiments on a series of applications to demonstrate the flexibility of our approach and the superior performance over state-of-the-art methods.

2. Related Work

Indoor Scene Synthesis. The goal of indoor scene synthesis [37] is to generate diverse and plausible indoor scene layouts. Some early works [9, 22, 31, 34, 62] adopt probabilistic models to parameterize the distribution of indoor scenes and sample novel scenes. However, due to the various object relationships and spatial constraints within indoor scenes, they suffer from the time-consuming optimization process and limited range of generated scenes.

Many studies have demonstrated the effectiveness of training deep neural networks on different synthetic large-scale indoor scene datasets [10, 27, 46]. The architectures of deep generative models have ranged from convolutional networks [42, 52], recursive neural networks [14, 26], graph neural networks [53, 68], to Transformers [36] and diffusion models [49, 55]. Among the recent works, ATISS [36] represents the indoor scenes as object sets and leverages the permutation equivariance of the transformer to learn the layout distribution. DiffuScene [49] trains a diffusion model to progressively denoise the layout attributes and finally generate plausible indoor scenes. To date, the attention mechanism and denoising diffusion model have been validated to be effective in capturing the complex relative placements between objects within the indoor scenes.

In parallel, some pioneer works [13, 61] study to use the pre-trained Large Language Models (LLMs) to facilitate open-vocabulary scene synthesis. They mainly focus on prompting the LLMs to generate complete scene configurations from the given simple and abstract descriptions, and directly parse the responses into scene layouts with heuristic optimizations. However, these works lack fine-grained and precise control over the generated scenes.

Controllable Scene Synthesis. Controllability is crucial for generating indoor scenes that meet user expectations. Many related works have explored scene synthesis conditioned on text descriptions [3–6, 32, 44], sketches [25, 45, 57], images [19, 21, 35], etc. They often exploit structured scene graphs to represent the specified conditions and capture object relation priors to guarantee the semantic alignment between input conditions and synthesized scenes.

Graph-based networks have shown significant potential in conditional scene synthesis. PartIT [53] utilizes a graph convolutional network to synthesize relation graphs, followed by an image-based convolution network to instantiate novel scenes, enhancing both the diversity and plausibility. Some other related works, such as SceneGraphNet [68] and CommonScenes [65], directly take scene graphs as conditions and use graph neural networks with message passing modules to complete layout attributes of the graphs. Graph-to-3D [7] and EchoScene [64] incorporate both the object shape and their placement while facilitating information exchange between them to synthesize diverse scenes. InstructScene [28] quantifies various object features with a

codebook and proposes a discrete semantic graph diffusion which significantly improves the controllability and fidelity of the synthesized indoor scenes.

Diffusion Models for Layout Synthesis. The diffusion model [17, 47, 48] has made remarkable achievement in a variety of content creation applications. As previously mentioned, some works [49, 55] directly utilize the diffusion model for indoor scene synthesis. Upon the basic design of diffusion models, PhyScene [60] uses several guidances for the diffusion model to improve the physical plausibility and interactivity. EchoScene [64] associates each object node with a denoising process and enables collaborative information exchange for better controllability and fidelity.

Indoor scene representation plays a crucial role during the iterative denoising with diffusion models. The straightforward solution is to represent scene layouts as a set of object tokens in the continuous attribute vector space [49, 60] or latent space [64]. With the development of Discrete Denoising Diffusion Probabilistic Models (D3PM) [1], several other domains that involve graph structures have started employing discrete [29, 40] or mixed diffusion [51] models to accomplish generative tasks, showcasing their potential in producing structured and high-quality outputs. Some works quantify the object attributes [20] or features [28] as discrete tokens and utilize discrete diffusion for layout synthesis. Further, MiDiffusion [18] takes object categories as discrete tokens and placement attributes as continuous values, and proposes to learn their distribution simultaneously to produce scenes with high fidelity and plausibility.

3. Preliminary

In this section, we concisely introduce the diffusion models in the continuous and discrete domains, which form the foundation for the derivation of our loss function.

Continuous Diffusion Model. The Denoising Diffusion Probabilistic Model (DDPM) [17] proposed the derivation for the forward diffusion process and reverse denoising process in continuous space. By fitting the posterior distribution of the reverse denoising with a neural network, we can obtain data samples that conform to the true distribution. DDPM initially injects Gaussian noise into the continuous variable x_0 iteratively, such that by time T , the data distribution approximates a standard normal distribution:

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

where $\beta_1, \dots, \beta_T \in (0, 1)$. The conditional distribution of the forward process at timestep t can be expressed as:

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (2)$$

where $\alpha_t := 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$. For the posterior distribution during the reverse denoising process, it is also represented as Gaussian distribution:

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\boldsymbol{\beta}}_t\mathbf{I}), \quad (3)$$

where $\tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0) := \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}\mathbf{x}_0 + \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}\mathbf{x}_t$ and the variance is defined as $\tilde{\boldsymbol{\beta}}_t := \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t}\beta_t$.

Discrete Diffusion Model. D3PM [1] defines a state transition matrix and introduces a [MASK] state to enable the forward diffusion and reverse denoising in the discrete space. For a discrete variable z with K categories, the state transition matrix is a $\mathbb{R}^{(K+2) \times (K+2)}$ matrix, where the extra two dimensions are for the empty and mask states, i.e.,

$$\mathbf{Q}_t := \begin{bmatrix} \alpha_t + \beta_t & \beta_t & \cdots & \beta_t & 0 \\ \beta_t & \alpha_t + \beta_t & \cdots & \beta_t & 0 \\ \vdots & \vdots & \ddots & \beta_t & 0 \\ \beta_t & \beta_t & \beta_t & \alpha_t + \beta_t & 0 \\ \gamma_t & \gamma_t & \gamma_t & \gamma_t & 1 \end{bmatrix}. \quad (4)$$

Therefore, at a given timestep t , a discrete variable z_t (represented as a one-hot row vector) has a probability γ_t of being masked in the subsequent time step, a probability α_t of remaining unchanged, while the remaining probability $(1 - \gamma_t - \alpha_t)$ is uniformly transited to the other $K + 1$ states (1 for empty). Then the discrete forward diffusion process can be defined as a conditional probability distribution:

$$q(z_t|z_{t-1}) = \text{Cat}(z_t; \mathbf{p} = z_{t-1}\mathbf{Q}_t), \quad (5)$$

where $\text{Cat}(z; \mathbf{p})$ is a categorical distribution over the one-hot row vector z with probabilities for each possible outcome given by the row vector \mathbf{p} . And the entire forward process can be parameterized as follows:

$$q(z_t|z_0) = \text{Cat}(z_t; \mathbf{p} = z_0\bar{\mathbf{Q}}_t), \quad (6)$$

where $\bar{\mathbf{Q}}_t = \mathbf{Q}_1\mathbf{Q}_2 \dots \mathbf{Q}_t$. Accordingly, the posterior distribution of the denoising process for discrete variables is

$$\begin{aligned} q(z_{t-1}|z_t, z_0) &= \frac{q(z_t|z_{t-1}, z_0)q(z_{t-1}|z_0)}{q(z_t|z_0)} \\ &= \text{Cat}\left(z_{t-1}; \mathbf{p} = \frac{z_t\mathbf{Q}_t^\top \odot z_0\bar{\mathbf{Q}}_{t-1}}{z_0\bar{\mathbf{Q}}_t z_t^\top}\right). \end{aligned} \quad (7)$$

4. Method

We start with problem formulation (Section 4.1) of indoor scene synthesis and then introduce our FreeScene framework. As shown in Figure 2, the framework comprises two key components: one is a carefully engineered agent named Graph Designer (Section 4.2), which analyzes the given text and/or images, and composites the extracted object and relations into a partial graph; the other is a generative model named MG-DiT (Section 4.3), which populates the partial graph to generate complete layouts of plausible and coherent indoor scenes. Finally, we describe various applications derived from MG-DiT based on constrained sampling, without the need for further training (Section 4.4).

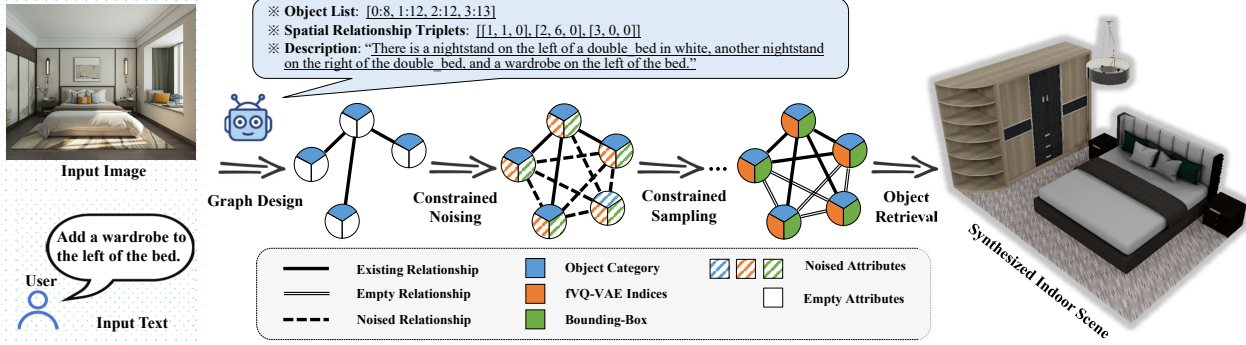


Figure 2. Overview of FreeScene. Given a reference image, text, or either of them, a carefully designed agent called the **Graph Designer** analyzes the objects and their spatial relationships, constructing a partial graph prior that captures both the object types and their spatial relationships. Next, leveraging the proposed **MG-DiT**, we apply a constrained sampling method that preserves the integrity of the graph prior throughout the noising and denoising processes, while generating the remaining parts. Ultimately, we obtain the complete room layout and perform object retrieval from the 3D-FUTURE dataset [12], retrieving each object with the most similar OpenCLIP [30] feature (derived from the generated fVQ-VAE [28] indices) within the same category.

4.1. Problem Formulation

We define $\mathcal{S} := \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_K\}_{k=1}^K$ to represent the indoor scene dataset, where each scene $\mathcal{S}_k := \{\mathcal{O}_k, \mathcal{G}_k\}$ consists of an object collection $\mathcal{O}_k = \{o_i^k\}_{i=1}^N$ with N objects and an edge collection $\mathcal{G}_k = \{e_{i \rightarrow j} \in \{1, 2, \dots, Q\} | i, j \in \{1, 2, \dots, N\}, i \neq j\}$, where $e_{i \rightarrow j}$ denotes the spatial relationship from the i_{th} objects to the j_{th} out of Q discrete labels. For objects, such as o_i^k denoting the i_{th} object in the k_{th} scene, we have $o_i^k := (c_i^k, v_i^k, s_i^k, t_i^k, r_i^k)$ containing the discrete attributes, i.e. object category $c_i^k \in \{1, 2, \dots, P\}$ in P discrete labels, feature indices $v_i^k \in \{1, 2, \dots, K_f\}^{n_f}$ with n_f tokens in K_f discrete labels, and continuous attributes, i.e. object size $s_i^k \in \mathbb{R}^3$, position $t_i^k \in \mathbb{R}^3$, and orientation $r_i^k \in \mathbb{R}^2$ (in the form of $\cos \theta$ and $\sin \theta$), respectively. Specifically, the feature indices v_i^k refer to n_f discrete tokens quantized from the features stored in the codebook $\mathcal{Z} \in \mathbb{R}^{K_f \times d_z}$, for which we leverage fVQ-VAE [28] to compress the OpenCLIP [30] features of the objects. Finally, assuming y_k is the textual description of the k_{th} scene in the dataset, our target is to optimize the parameters θ to maximize the following log-likelihood:

$$\log p_{\theta}(\mathcal{S}) = \sum_{k=1}^K \log p_{\theta}(\mathcal{S}_k | y_k) = \sum_{k=1}^K \log p_{\theta}(\mathcal{O}_k, \mathcal{G}_k | y_k). \quad (8)$$

4.2. Graph Designer

Our Graph Designer leverages pre-trained Vision-Language Models (VLMs) to extract the objects and their spatial relationships from user-provided text and/or image prompts. The VLM responses are then transformed into the graph representation, which forms the partial graph condition for the subsequent diffusion model for seamless integration and effective processing. The partial graph is defined as a list of object categories $\{c_0, c_1, \dots, c_{N^-}\}$ and an edge collection \mathcal{G}^- forming a relationship matrix, complemented with a global textual description. Here, N^- represents the number

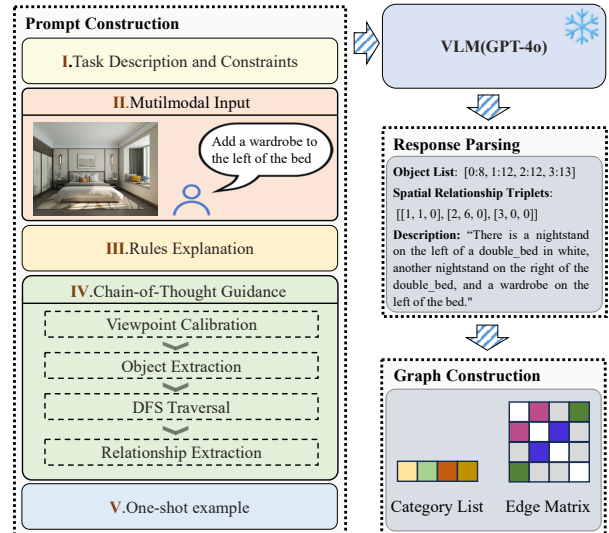


Figure 3. Graph Designer. The Graph Designer utilizes a one-shot Chain-of-Thought (CoT) method to guarantee the accuracy of the extraction results. It then parses and preprocesses the VLM responses to ensure compatibility with the MG-DiT inputs.

of objects extracted from the free-form user inputs, which acts as a subset of the final scene, and \mathcal{G}^- denotes the relationships corresponding to these extracted objects.

Prompt Construction. In order to ensure the output data format and efficient information extraction, we develop a one-shot Chain-of-Thought (CoT) prompt template for the VLMs in which a one-shot formatted example is given to regulate the output format following [61]. Moreover, we exploit CoT [54] to steer the VLMs to tackle the inference task logically and systematically. Our one-shot CoT prompt consists of the following components: (1) *Task description and constraints*, which outline the task objectives and specify the types of objects and relationships that might be involved. (2) *Multimodal user input* in the form of text and/or

image, where the image can be a real or synthesized image, a line drawing sketch, or a top-view diagram. (3) *Rule explanation*, which specifies certain rules governing the response generation process and establishes connections between these rules and the output through detailed descriptions. (4) *Chain-of-Thought guidance*, which provides the step-by-step hint to guide the systematic reasoning process of the VLM. (5) *One-shot example*, a demonstration with a strictly defined format. We instruct the VLMs to adhere to this format, ensuring seamless subsequent parsing.

Chain-of-Thought Guidance. To enhance the VLM to stably and precisely infer graph conditions from the free-form user inputs, we formulate the Chain-of-Thought guidance, which consists of the following steps: (1) *Viewpoint calibration*. We instruct the VLM to determine the viewpoint of input images, so that subsequent relationship judgments can accurately distinguish spatial relationships corresponding to different viewpoints. (2) *Object extraction*. We guide the VLM to identify the objects belonging to a set of predefined categories from the given image and text. (3) *DFS traversal*. Since calculating pairwise relationships between all objects often leads to highly disorganized outputs or the omission of critical relationships, we design a detailed set of instructions to guide the VLM to start from a key root node and use a depth-first search (DFS) approach to hierarchically traverse all objects, facilitating an organized relationship extraction. (4) *Relationship extraction*. Based on the traversal results from previous step, we individually evaluate each relationship between the parent and child nodes, then integrate the findings to produce the final output.

Graph Construction. The VLM output consists of an object list, a relation triplet list, and a textual description of the room. As shown in Figure 3, the object list contains the indices and categories of the objects. The relation triplet list contains the relation type and indices of the two adjacent objects, forming a relationship matrix. The outputs are parsable and can be converted to a graph representation, thanks to the one-shot example provided in the prompt. We utilize GPT-4o and apply regular expressions to parse its outputs. The full prompts are provided in Supplementary.

4.3. Mixed Graph Diffusion Transformer

We propose Mixed Graph Diffusion Transformer (MG-DiT) to facilitate a variety of applications including text-to-scene generation, zero-shot graph-to-scene generation, etc. with a single model through constrained sampling. During the inference phase, we can seamlessly input the combination of the object category list, relationship matrix, and textual description produced by the Graph Designer as the partial graph condition of MG-DiT for indoor scene synthesis.

Mixed Graph Diffusion. Drawing inspiration from MiDiffusion [18], we inject domain-specific noise into both continuous and discrete variables simultaneously during train-

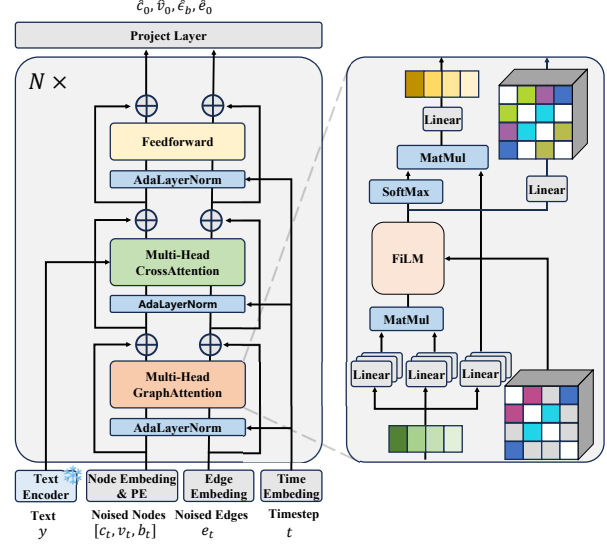


Figure 4. Mixed Graph Diffusion Transformer. MG-DiT conditions on text and timesteps to jointly denoise continuous bounding box features along with discrete graph and fVQ-VAE features. The network consists of input and output processing layers, stacked with several Mixed Graph Transformer Blocks, as illustrated.

ing. For a scene in the form of a graph, we have $\mathbf{b} = [\mathbf{s}, \mathbf{t}, \mathbf{r}]$ to represent the continuous attributes, i.e. object size, position, and orientation, and $\mathbf{z} = [\mathbf{c}, \mathbf{v}, \mathbf{e}]$ to denote the discrete attributes, i.e. object category, feature indices (fVQ-VAE indices), and the relationship category. Then the mixed forward diffusion process can be expressed as follows:

$$q(\mathbf{b}_t, \mathbf{z}_t | \mathbf{b}_{t-1}, \mathbf{z}_{t-1}) := \tilde{q}(\mathbf{b}_t | \mathbf{b}_{t-1}) \cdot \hat{q}(\mathbf{z}_t | \mathbf{z}_{t-1}). \quad (9)$$

Correspondingly, we can factor the posterior distribution of the denoising process as:

$$q(\mathbf{b}_{t-1}, \mathbf{z}_{t-1} | \mathbf{b}_t, \mathbf{z}_t, \mathbf{b}_0, \mathbf{z}_0) = \tilde{q}(\mathbf{b}_{t-1} | \mathbf{b}_t, \mathbf{b}_0) \cdot \hat{q}(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0). \quad (10)$$

We introduce MG-DiT to fit this posterior distribution:

$$p_\theta(\mathbf{b}_{t-1}, \mathbf{z}_{t-1} | \mathbf{b}_t, \mathbf{z}_t, \mathbf{y}) := \tilde{p}_\theta(\mathbf{b}_{t-1} | \mathbf{b}_t, \mathbf{z}_t, \mathbf{y}) \cdot \hat{p}_\theta(\mathbf{z}_{t-1} | \mathbf{b}_t, \mathbf{z}_t, \mathbf{y}), \quad (11)$$

where θ is the MG-DiT network parameter. The negative log-likelihood is employed as our loss function. Following DDPM [17], we adopt the variational lower bound and omit the constant term L_T :

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_b + \mathcal{L}_z \\ &= \mathbb{E}_{\mathbf{b}_0, \mathbf{z}_0, t \in [1, T], \epsilon_b, q(\mathbf{z}_t | \mathbf{z}_0)} [\\ &\quad \|\epsilon_b - \epsilon_b(\sqrt{\bar{\alpha}_t} \mathbf{b}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_b, t, \mathbf{y})\|^2 \\ &\quad + D_{KL}(\tilde{q}(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0) \| \hat{p}_\theta(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{y}))], \end{aligned} \quad (12)$$

where t is sampled from a uniform distribution $\mathcal{U}(1, T)$ and ϵ is sampled from a standard normal distribution $\mathcal{N}(0, \mathbf{I})$. **Network Architecture.** As depicted in Figure 4, in each

App. \ Var.	c	e	s	$[t, o]$	v
Text-to-scene	gen	gen	gen	gen	gen
Graph-to-scene	f/p	f/p	gen	gen	gen
Re-arrangement	fixed	gen	fixed	gen	fixed
Completion	partial	partial	partial	partial	partial
Stylization	fixed	fixed	fixed	fixed	gen

Table 1. Feature fixation solutions for different applications. ‘gen’ represents normal denoising. ‘fixed’ indicates that the features of all objects in a scene remain unchanged during the denoising process, while ‘partial’ signifies that only partial existing features are fixed. ‘f/p’ means we can either fix all attributes for a full condition, or only fix partial existing attributes for a partial condition.

timestep of the denoising process, we embed and concatenate the noisy attributes (c, v, s, t, r) as node features, incorporating sinusoidal positional encodings. Simultaneously, the noisy attribute e is embedded as edge features. These nodes and edges are fed into a DiT network [38] to fit the posterior distribution, in which we integrate the cross-attention modules to handle text input. Following [8, 28, 50], we employ FiLM [39] as the interaction mechanism between nodes and edges, which can be expressed by the following function:

$$\text{FiLM}(sim, e) = \gamma(e) \left(\frac{sim - \mu(sim)}{\sigma(sim)} \right) + \beta(e), \quad (13)$$

where $\gamma(\cdot)$ and $\beta(\cdot)$ represent the mapping from the edge matrix to the new scale and bias. The network’s output consists of the predicted noise $\hat{e}_s, \hat{e}_r, \hat{e}_t$ for the bounding box, as well as the predicted values of \hat{c}_0, \hat{v}_0 , and \hat{e}_0 at timestep 0, which are used to calculate the posterior probability distribution during the denoising process.

4.4. Zero-Shot Applications

Constrained sampling has been employed in prior works [28, 36, 49] for tasks such as scene re-arrangement and completion to maintain some fixed elements. For instance, in re-arrangement tasks, the categories and sizes of objects are fixed while their positions are perturbed and denoised. In completion tasks, all attributes of existing objects are fixed while the rest of objects are generated. We extend this approach by fixing the graph defined by the graph designer or users, while generating the remaining attributes. This enables the use of a single model for both text-to-scene and graph-to-scene generation, without the need for additional training. Besides, we can apply MG-DiT to other zero-shot downstream applications. The variables for noising and denoising process encompass c, v, b and e , where $b = [s, t, o]$. At each step of the sampling process, we fix the corresponding variables to support different applications, and the settings of different applications are presented in Table 1. For simplicity, we use the same notation as in Section 4.1 but omit the subscripts and superscripts.

5. Experiments

Datasets. Following previous works, we conduct experiments on 3D-FRONT [11] dataset, an indoor scene dataset composed of 6,813 houses with 14,629 rooms. During the object retrieval phase of our indoor scene synthesis, we select the closest furniture models from 3D-FUTURE [12] based on object categories and fVQ-VAE features. We conduct experiments on three types of scenes: 4,041 bedrooms, 900 dining rooms, and 813 living rooms. For further details on data preprocessing, please refer to the Supplementary.

Baselines. To validate the effectiveness of our approach, we evaluate its performance from *three* perspectives: text-to-scene generation, zero-shot graph-to-scene generation, and other applications. We compare our method with the state-of-the-art approaches including: **1)** ATISS [36], a transformer-based autoregressive model to which we added a text-based cross-attention layer to enable text-to-scene generation; **2)** DiffuScene [49], a diffusion model for indoor scene synthesis with the U-Net architecture; **3)** InstructScene [28], a dual-model approach that first predicts the graph using a text-to-graph model and then generates the scene using a graph-to-scene model.

Metrics. We focus on two key aspects for a comprehensive evaluation of our approach. **1) For the quality of the generated scenes**, we primarily assess the following metrics: Fréchet Inception Distance (FID) ([16]), FID^{CLIP} [24] (which uses CLIP features [41] to compute FID scores), Kernel Inception Distance (KID) [2], and Scene Classification Accuracy (SCA) [36]. SCA involves training a discriminator model to distinguish between real and generated data, assessing the ability of the generated data to fool the discriminator. The closer the model’s accuracy on generated data is to 50%, the more realistic the generated results are perceived. **2) For the controllability of the models**, we adopt and extend the iRecall metric introduced in InstructScene [28]. For text-to-scene generation, iRecall measures the probability that each sentence in the input text occurs in the generated results. For graph-to-scene generation, iRecall reflects the probability that each relationship in the graph occurs in the generated scene.

5.1. Evaluation for Graph Designer

To validate the effectiveness of Graph Designer, we compared the graph extraction results of one-shot CoT prompts and one-shot prompts without CoT across a collection of images and text, including the photos and sketches of different views and room types, manually created top-view diagrams, and a set of textual scene descriptions. Both of the two prompt templates comprise a formatted example to ensure the consistency of the results. The detailed prompt templates can be found in Supplementary. Additionally, during testing, we input either images or text into the VLMs to evaluate the performance on each specific modality. We as-

Place two chairs on either side of a dining table. Position a TV stand to the left of the table, with a sofa situated behind the TV stand.

There is a coffee table together with a TV stand placed in front of a sofa, and a cabinet positioned to the left of the TV stand.

There is a nightstand on either side of a white bed, and a wardrobe is positioned in front of the bed.

Text Prompts




Figure 5. Qualitative comparisons on text-to-scene generation.

Graph Designer	\uparrow Object iRecall%	\uparrow Rel Acc%
Image	85.23/72.07	77.56/34.65
with CoT vs.	Sketch	64.97/58.25
without CoT	Diagram	91.22/88.13
	Text	98.56/95.56
		74.63/44.32
		89.45/85.60

Table 2. Quantitative comparisons on graph extraction performance between one-shot CoT prompts (with CoT) and one-shot prompts (without CoT). For each pair of data, number before and after the slash is with CoT and without CoT, respectively.

sess the extraction performance from two aspects: the ratio of correctly identified objects w.r.t. the total number of objects named **Object iRecall** and the proportion of correctly recognized relationships named **Rel Acc**.

Table 2 shows the quantitative evaluations of one-shot CoT prompts and one-shot prompts without CoT. As demonstrated, one-shot CoT prompts can significantly enhance the graph extraction performance of the VLMs across various modalities, particularly in terms of relationship accuracy. This is because one-shot prompts are prone to producing contradictory or erroneous relationships. In contrast, one-shot CoT prompts can effectively guide the VLM to hierarchically navigate the entire scene, precisely establishing accurate relationships throughout. We also present a comparison of the graph extraction performance on a specific image-text pair example between one-shot CoT prompts and one-shot prompts, the results are provided in Supplementary due to space constraints.

5.2. Text-to-Scene Generation

We compare with three existing methods to demonstrate the effectiveness of our method. Additionally, we provide

Text-to-scene	FID	FID ^{CLIP}	KID _{$\times 1e^{-3}$}	SCA%	iRecall%	
Bedroom	ATISS	122.37	8.23	0.74	58.77	44.06
	DiffuScene	129.34	9.66	0.81	59.13	53.32
	InstructScene	114.86	6.52	0.68	56.37	72.71
	Ours	111.21	6.43	0.35	54.94	73.69
	Ours+GD	108	6.07	0.21	53.16	81.40
Livingroom	ATISS	120.10	7.43	16.44	57.02	36.31
	DiffuScene	135.93	10.71	29.07	58.98	39.58
	InstructScene	111.52	5.91	8.65	55.32	57.21
	Ours	110.55	5.83	7.95	55.24	58.16
	Ours+GD	108.22	5.23	3.87	54.05	71.81
Diningroom	ATISS	134.75	9.82	24.25	59.83	32.56
	DiffuScene	142.37	11.76	28.36	60.88	37.17
	InstructScene	129.13	8.24	15.27	58.93	61.47
	Ours	127.28	8.01	14.83	56.82	63.39
	Ours+GD	125.32	7.14	10.68	55.59	75.01

Table 3. Quantitative comparisons on text-to-scene generation. Lower FID, FID^{CLIP}, KID and higher iRecall indicate better synthesis quality. For SCA, a score closer to 50% is better.

two configurations with our approach: the first directly employs MG-DiT for text-to-scene generation (named **Ours**); the second uses Graph Designer to preprocess the textual scene description, followed by the MG-DiT conditioned on the extracted partial graphs to generate the scenes (named **Ours+GD**). To ensure a fair comparison, we keep the floor texture and background consistent among all the results.

Table 3 presents a quantitative comparison of our method against other approaches across various metrics. As shown in the table, our method exhibits excellent text-to-scene generation performance and surpasses other approaches in various metrics. This reveals that taking graph information into account significantly improves the feature extraction ability of our method during the generation process, making

Graph-to-scene		FID	FID ^{CLIP}	KID _{×1e-3}	SCA%	iRecall%
Bedroom	InstructScene	101.86	5.66	0.13	53.68	88.84
	Ours	98.31	5.58	0.12	52.34	89.37
Livingroom	InstructScene	108.45	5.43	0.33	54.51	83.14
	Ours	107.62	5.25	0.29	53.77	83.35
Diningroom	InstructScene	123.82	8.21	14.18	56.23	84.16
	Ours	122.57	7.49	13.11	55.70	84.47

Table 4. Quantitative comparisons on graph-to-scene generation

it outperform the autoregressive model ATISS and the diffusion model DiffuScene by a considerable margin. The 1D convolution of DiffuScene makes it challenging to capture global features, while ATISS struggles to integrate global text descriptions, leading to reduced controllability. Compared to InstructScene, since our MG-DiT model employs a combination of discrete graph and continuous bounding box variables, it has to predict the graph from discrete noise, which contains [MASK] tokens or other noise states, while concurrently generating the bounding boxes. This forces the model to more effectively learn global scene features and grasp the relationships between objects during graph prediction, thereby improving the final generation quality and controllability. By directly extracting graph from text using Graph Designer as a preprocessing step, our model can achieve superior generation quality and controllability.

In Figure 5, we present a qualitative comparison with several baselines. It is evident that our MG-DiT algorithm produces higher-quality layouts compared to InstructScene. Moreover, preprocessing the text using the Graph Designer enhances controllability, yielding outputs that align more closely with the input text. Additional results are provided in the Supplementary Materials due to space constraints.

5.3. Graph-to-Scene Generation

Regarding graph-to-scene generation, we assess both the quality of the output scenes and the consistency between the generated scenes and the provided graphs. Since ATISS and DiffuScene do not involve graph priors, we only compared the graph-conditioned results of MG-DiT with InstructScene. For InstructScene, we employ constrained sampling in the text-to-graph model to fix the graph (object categories and edges) and generate fVQ-VAE indices, then use the graph-to-scene model to produce the final results. Table 4 shows the quantitative results, indicating that our hybrid graph diffusion architecture achieves improved performance in both generation quality and controllability. More qualitative results are presented in Supplementary.

5.4. More Applications

Table 5 presents comparisons with existing methods across various applications, including re-arrangement, completion, stylization and unconditioned generation without any text or graph condition. To evaluate scene stylization, we use

Applications	Re-arrangement		Completion		Stylization	Uncond.	
	FID	iRecall%	FID	iRecall%	$\Delta_{\times 1e-3}$	FID	
Bedroom	ATISS	-	-	95.51	43.77	-	126.80
	DiffuScene	119.46	55.37	100.72	52.15	-	138.24
	InstructScene	106.80	76.24	82.94	67.47	6.55	124.77
	Ours	105.71	77.49	81.83	69.25	6.79	123.42
Livingroom	ATISS	-	-	99.25	28.65	-	124.55
	DiffuScene	124.65	32.33	103.17	25.08	-	139.63
	InstructScene	107.78	53.40	92.38	48.53	0.29	117.71
	Ours	106.33	54.82	91.01	50.57	0.31	116.87
Diningroom	ATISS	-	-	117.20	35.70	-	139.26
	DiffuScene	136.37	34.62	121.53	33.97	-	148.54
	InstructScene	125.96	62.45	107.44	59.21	1.83	136.40
	Ours	124.69	63.02	105.71	60.67	1.87	132.22

Table 5. Quantitative comparisons on various applications.

metric Δ in [28] which measures the cosine similarity between the CLIP features of the target style text and the top-down view images of the generated scenes. Since ATISS and DiffuScene lack features indicating color and texture, it is challenging to conduct stylization experiments. Additionally, as an autoregressive model, ATISS is unsuitable for re-arrangement tasks. Therefore, we have excluded these experimental results. Our experimental findings clearly demonstrate that our method outperforms other baseline approaches across re-arrangement, completion, and unconditional generation tasks and achieves comparable results to InstructScene in the stylization task. More qualitative results are presented in Supplementary.

6. Conclusion

We introduce FreeScene, a novel framework designed to promote both user interactivity and the quality of indoor scene synthesis through the integration of a VLM-based Graph Designer and the MG-DiT model. Experimental results demonstrate our FreeScene can not only enhance the performance of various tasks, but also provides a more user-friendly control mechanism for scene synthesis.

Limitations & Future Work. FreeScene also has several limitations. First, the Graph Designer is limited by the capabilities of VLMs, which may struggle to capture all the graph priors, especially in scenes with a large number of objects. This limitation could be mitigated with future advancements in VLM technology. Second, the MG-DiT cannot precisely control the exact position and orientation of objects, which may result in furniture arrangements that are impractical for human use. Adding additional constraints, like affordability, to diffusion models may offer a promising solution to this issue. Last but not least, the scale and quality of available indoor scene datasets are still insufficient, which may lead to overfitting and suboptimal performance. A potential direction for future work to address this data limitation is to combine the general knowledge embedded in large models with real samples from databases, leveraging techniques such as Retrieval-Augmented Generation (RAG) for interactive room design.

Acknowledgements

This research was supported in part by the National Natural Science Foundation of China (No. 62202199, 62302269), the Excellent Young Scientists Fund Program (Overseas) of Shandong Province (No. 2023HWYQ-034) and the Fundamental Research Funds for the Central Universities.

References

- [1] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021. 3
- [2] Miłkołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv:1801.01401*, 2018. 6
- [3] Angel Chang, Manolis Savva, and Christopher D Manning. Interactive learning of spatial knowledge for text to 3d scene generation. In *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces*, pages 14–21, 2014. 2
- [4] 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.
- [5] Angel Chang, Will Monroe, Manolis Savva, Christopher Potts, and Christopher D Manning. Text to 3d scene generation with rich lexical grounding. *arXiv:1505.06289*, 2015.
- [6] Angel X Chang, Mihail Eric, Manolis Savva, and Christopher D Manning. Sceneseer: 3d scene design with natural language. *arXiv:1703.00050*, 2017. 2
- [7] Helisa Dhama, Fabian Manhardt, Nassir Navab, and Federico Tombari. Graph-to-3d: End-to-end generation and manipulation of 3d scenes using scene graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16352–16361, 2021. 2
- [8] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv:2012.09699*, 2020. 6
- [9] Matthew Fisher, Daniel Ritchie, Manolis Savva, Thomas Funkhouser, and Pat Hanrahan. Example-based synthesis of 3d object arrangements. *ACM Transactions on Graphics (TOG)*, 31(6):1–11, 2012. 1, 2
- [10] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Binqiang 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. 2
- [11] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Binqiang 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. 6
- [12] Huan Fu, Rongfei Jia, Lin Gao, Mingming Gong, Binqiang Zhao, Steve Maybank, and Dacheng Tao. 3d-future: 3d furniture shape with texture. *International Journal of Computer Vision*, 129:3313–3337, 2021. 4, 6
- [13] Rao Fu, Zehao Wen, Zichen Liu, and Srinath Sridhar. Any-home: Open-vocabulary generation of structured and textured 3d homes. In *European Conference on Computer Vision*, pages 52–70. Springer, 2025. 2
- [14] Lin Gao, Jia-Mu Sun, Kaichun Mo, Yu-Kun Lai, Leonidas J Guibas, and Jie Yang. Scenehgn: Hierarchical graph networks for 3d indoor scene generation with fine-grained geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7):8902–8919, 2023. 2
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2014. 2
- [16] 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. *Advances in Neural Information Processing Systems*, 30, 2017. 6
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 3, 5
- [18] Siyi Hu, Diego Martin Arroyo, Stephanie Debats, Fabian Manhardt, Luca Carlone, and Federico Tombari. Mixed diffusion for 3d indoor scene synthesis. *arXiv:2405.21066*, 2024. 2, 3, 5
- [19] Siyuan Huang, Siyuan Qi, Yinxue Xiao, Yixin Zhu, Ying Nian Wu, and Song-Chun Zhu. Cooperative holistic scene understanding: Unifying 3d object, layout and camera pose estimation. In *NeurIPS*, 2018. 2
- [20] Naoto Inoue, Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Layoutdm: Discrete diffusion model for controllable layout generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10167–10176, 2023. 3
- [21] Hamid Izadinia, Qi Shan, and Steven M Seitz. Im2cad. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5134–5143, 2017. 2
- [22] Z Sadehipour Kermani, Zicheng Liao, Ping Tan, and Hao Zhang. Learning 3d scene synthesis from annotated rgb-d images. In *Computer Graphics Forum*, pages 197–206. Wiley Online Library, 2016. 1, 2
- [23] Diederik P Kingma. Auto-encoding variational bayes. *arXiv:1312.6114*, 2013. 2
- [24] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of imagenet classes in fr’echet inception distance. *arXiv:2203.06026*, 2022. 6
- [25] Jeehyung Lee and Thomas A Funkhouser. Sketch-based search and composition of 3d models. In *SBIM*, pages 97–104, 2008. 2
- [26] Manyi Li, Akshay Gadi Patil, Kai Xu, Siddhartha Chaudhuri, Owais Khan, Ariel Shamir, Changhe Tu, Baoquan Chen, Daniel Cohen-Or, and Hao Zhang. Grains: Generative recursive autoencoders for indoor scenes. *ACM Transactions on Graphics (TOG)*, 38(2):1–16, 2019. 2

- [27] Wenbin Li, Sajad Saeedi, John McCormac, Ronald Clark, Dimos Tzoumanikas, Qing Ye, Yuzhong Huang, Rui Tang, and Stefan Leutenegger. InteriorNet: Mega-scale multi-sensor photo-realistic indoor scenes dataset. *arXiv:1809.00716*, 2018. [2](#)
- [28] Chenguo Lin and Yadong Mu. InstructScene: Instruction-driven 3d indoor scene synthesis with semantic graph prior. *arXiv:2402.04717*, 2024. [2](#), [3](#), [4](#), [6](#), [8](#)
- [29] Gang Liu, Jiaxin Xu, Tengfei Luo, and Meng Jiang. Graph diffusion transformers for multi-conditional molecular generation. *arXiv:2401.13858*, 2024. [3](#)
- [30] Minghua Liu, Ruoxi Shi, Kaiming Kuang, Yin hao Zhu, Xu-anlin Li, Shizhong Han, Hong Cai, Fatih Porikli, and Hao Su. Openshape: Scaling up 3d shape representation towards open-world understanding. *Advances in neural information processing systems*, 36, 2024. [4](#)
- [31] Tianqiang Liu, Siddhartha Chaudhuri, Vladimir G Kim, Qixing Huang, Niloy J Mitra, and Thomas Funkhouser. Creating consistent scene graphs using a probabilistic grammar. *ACM Transactions on Graphics (TOG)*, 33(6):1–12, 2014. [1](#), [2](#)
- [32] Rui Ma, Akshay Gadi Patil, Matthew Fisher, Manyi Li, Sören Pirk, Binh-Son Hua, Sai-Kit Yeung, Xin Tong, Leonidas Guibas, and Hao Zhang. Language-driven synthesis of 3d scenes from scene databases. *ACM Transactions on Graphics (TOG)*, 37(6):1–16, 2018. [2](#)
- [33] Léopold Maillard, Nicolas Sereyjol-Garros, Tom Durand, and Maks Ovsjanikov. Debara: Denoising-based 3d room arrangement generation. *arXiv:2409.18336*, 2024. [2](#)
- [34] Paul Merrell, Eric Schkufza, Zeyang Li, Maneesh Agrawala, and Vladlen Koltun. Interactive furniture layout using interior design guidelines. *ACM transactions on graphics (TOG)*, 30(4):1–10, 2011. [1](#), [2](#)
- [35] Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, and Jian Jun Zhang. Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [36] Despoina Paschalidou, Amlan Kar, Maria Shugrina, Karsten Kreis, Andreas Geiger, and Sanja Fidler. Atiss: Autoregressive transformers for indoor scene synthesis. *Advances in Neural Information Processing Systems*, 34:12013–12026, 2021. [2](#), [6](#)
- [37] Akshay Gadi Patil, Supriya Gadi Patil, Manyi Li, Matthew Fisher, Manolis Savva, and Haotong Zhang. Advances in data-driven analysis and synthesis of 3d indoor scenes. *Comput. Graph. Forum*, 43(1), 2024. [2](#)
- [38] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. [6](#)
- [39] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. [6](#)
- [40] Yiming Qin, Clement Vignac, and Pascal Frossard. Sparse training of discrete diffusion models for graph generation. *arXiv:2311.02142*, 2023. [3](#)
- [41] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [6](#)
- [42] Daniel Ritchie, Kai Wang, and Yu-an Lin. Fast and flexible indoor scene synthesis via deep convolutional generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6182–6190, 2019. [2](#)
- [43] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention*, pages 234–241, 2015. [2](#)
- [44] Lee M Seversky and Lijun Yin. Real-time automatic 3d scene generation from natural language voice and text descriptions. In *Proceedings of the 14th ACM international conference on Multimedia*, pages 61–64, 2006. [2](#)
- [45] HyoJong Shin and Takeo Igarashi. Magic canvas: interactive design of a 3-d scene prototype from freehand sketches. In *Proceedings of graphics interface 2007*, pages 63–70, 2007. [2](#)
- [46] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1746–1754, 2017. [2](#)
- [47] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. [2](#), [3](#)
- [48] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv:2011.13456*, 2020. [2](#), [3](#)
- [49] Jiapeng Tang, Yinyu Nie, Lev Markhasin, Angela Dai, Justus Thies, and Matthias Nießner. DiffScene: Denoising diffusion models for generative indoor scene synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20507–20518, 2024. [2](#), [3](#), [6](#)
- [50] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. *arXiv:2209.14734*, 2022. [6](#)
- [51] Clement Vignac, Nagham Osman, Laura Toni, and Pascal Frossard. Midi: Mixed graph and 3d denoising diffusion for molecule generation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 560–576. Springer, 2023. [3](#)
- [52] Kai Wang, Manolis Savva, Angel X Chang, and Daniel Ritchie. Deep convolutional priors for indoor scene synthesis. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018. [2](#)
- [53] Kai Wang, Yu-An Lin, Ben Weissmann, Manolis Savva, Angel X Chang, and Daniel Ritchie. Planit: Planning and instantiating indoor scenes with relation graph and spatial prior networks. *ACM Transactions on Graphics (TOG)*, 38(4):1–15, 2019. [2](#)

- [54] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022. 4
- [55] Qi Wei, Sijie Ding, Jeong Joon Park, Rahul Sajani, Adrien Poulenard, Srinath Sridhar, and Leonidas J. Guibas. Legonet: Learning regular rearrangements of objects in rooms. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19037–19047, 2023. 2, 3
- [56] Ken Xu, James Stewart, and Eugene Fiume. Constraint-based automatic placement for scene composition. In *Graphics Interface*, pages 25–34. Citeseer, 2002. 1
- [57] Kun Xu, Kang Chen, Hongbo Fu, Wei-Lun Sun, and Shi-Min Hu. Sketch2scene: Sketch-based co-retrieval and co-placement of 3d models. *ACM Transactions on Graphics (TOG)*, 32(4):1–15, 2013. 2
- [58] Haitao Yang, Zaiwei Zhang, Siming Yan, Haibin Huang, Chongyang Ma, Yi Zheng, Chandrajit Bajaj, and Qixing Huang. Scene synthesis via uncertainty-driven attribute synchronization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5630–5640, 2021. 2
- [59] Ming-Jia Yang, Yu-Xiao Guo, Bin Zhou, and Xin Tong. Indoor scene generation from a collection of semantic-segmented depth images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15203–15212, 2021. 2
- [60] Yandan Yang, Baoxiong Jia, Peiyuan Zhi, and Siyuan Huang. Physcene: Physically interactable 3d scene synthesis for embodied ai. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16262–16272, 2024. 2, 3
- [61] Yue Yang, Fan-Yun Sun, Luca Weihs, Eli VanderBilt, Alvaro Herrasti, Winson Han, Jiajun Wu, Nick Haber, Ranjay Krishna, Lingjie Liu, et al. Holodeck: Language guided generation of 3d embodied ai environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16227–16237, 2024. 2, 4
- [62] Lap Fai Yu, Sai Kit Yeung, Chi Keung Tang, Demetri Terzopoulos, Tony F Chan, and Stanley J Osher. Make it home: automatic optimization of furniture arrangement. *ACM Transactions on Graphics (TOG)-Proceedings of ACM SIGGRAPH 2011*, v. 30,(4), July 2011, article no. 86, 30(4), 2011. 1, 2
- [63] Lap-Fai Yu, Sai-Kit Yeung, and Demetri Terzopoulos. The clutterpalette: An interactive tool for detailing indoor scenes. *IEEE transactions on visualization and computer graphics*, 22(2):1138–1148, 2015. 1
- [64] Guangyao Zhai, Evin Pinar Örnek, Dave Zhenyu Chen, Ruo-tong Liao, Yan Di, Nassir Navab, Federico Tombari, and Benjamin Busam. Echoscene: Indoor scene generation via information echo over scene graph diffusion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. 2, 3
- [65] Guangyao Zhai, Evin Pinar Örnek, Shun-Cheng Wu, Yan Di, Federico Tombari, Nassir Navab, and Benjamin Busam. Commonsences: Generating commonsense 3d indoor scenes with scene graphs. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [66] Song-Hai Zhang, Shao-Kui Zhang, Wei-Yu Xie, Cheng-Yang Luo, Yong-Liang Yang, and Hongbo Fu. Fast 3d indoor scene synthesis by learning spatial relation priors of objects. *IEEE Transactions on Visualization and Computer Graphics*, 28(9):3082–3092, 2021. 1
- [67] Zaiwei Zhang, Zhenpei Yang, Chongyang Ma, Linjie Luo, Alexander Huth, Etienne Vouga, and Qixing Huang. Deep generative modeling for scene synthesis via hybrid representations. *ACM Transactions on Graphics (TOG)*, 39(2):1–21, 2020. 2
- [68] Yang Zhou, Zachary While, and Evangelos Kalogerakis. Scenegraphnet: Neural message passing for 3d indoor scene augmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7384–7392, 2019. 2