

# EchoScene: Indoor Scene Generation via Information Echo over Scene Graph Diffusion

Guangyao Zhai<sup>1</sup> Evin Pinar Örnek<sup>1</sup> Dave Zhenyu Chen<sup>1</sup> Ruotong Liao<sup>2,3</sup>  
Yan Di<sup>1</sup> Nassir Navab<sup>1</sup> Federico Tombari<sup>1,4</sup> Benjamin Busam<sup>1,3</sup>

<sup>1</sup>Technical University of Munich   <sup>2</sup>Ludwig Maximilian University of Munich

<sup>3</sup>Munich Center for Machine Learning   <sup>4</sup>Google

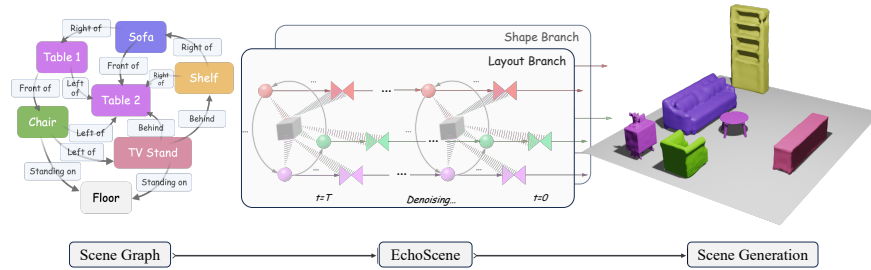
[guangyao.zhai@tum.de](mailto:guangyao.zhai@tum.de)

<https://sites.google.com/view/echoscene>

**Abstract.** We present *EchoScene*, an interactive and controllable generative model that generates 3D indoor scenes on scene graphs. EchoScene leverages a dual-branch diffusion model that dynamically adapts to scene graphs. Existing methods struggle to handle scene graphs due to varying numbers of nodes, multiple edge combinations, and manipulator-induced node-edge operations. EchoScene overcomes this by associating each node with a denoising process and enables collaborative information exchange, enhancing controllable and consistent generation aware of global constraints. This is achieved through an information echo scheme in both shape and layout branches. At every denoising step, all processes share their denoising data with an information exchange unit that combines these updates using graph convolution. The scheme ensures that the denoising processes are influenced by a holistic understanding of the scene graph, facilitating the generation of globally coherent scenes. The resulting scenes can be manipulated during inference by editing the input scene graph and sampling the noise in the diffusion model. Extensive experiments validate our approach, which maintains scene controllability and surpasses previous methods in generation fidelity. Moreover, the generated scenes are of high quality and thus directly compatible with off-the-shelf texture generation. Our code and models are open-sourced.

## 1 Introduction

Controllable Scene Generation (CSG) refers to synthesizing realistic 3D scenes according to input prompts while enabling specific entities within the scene to be user-interactive [5, 6, 49]. It has successfully been applied in robotics [38, 57], Virtual Reality / Augmented Reality [2], and autonomous driving [31, 45]. Recently, combining CSG with scene graph diffusion models has attracted significant research interest [32, 53, 58], since on the one hand, diffusion models empower more realistic and diverse 3D content generation [9, 10, 44, 48], on the other hand scene graphs capture the structure of the scene in a compact manner and facilitate intuitive user manipulation [15, 58]. Importantly, users can modify the input scene graph to dynamically change the generated scene.



**Fig. 1: EchoScene Schematic.** EchoScene uses a dual-branch diffusion model to generate 3D scenes from scene graphs. In both branches, each node is allocated a denoising process, and different processes are aware of global states through layout and shape "echoes" (waves in different colors) with an information exchange unit (grey block) along the denoising steps.

Despite its significant progress so far, CSG with scene graph diffusion still suffers from two open challenges. **First**, due to varying numbers of graph nodes and manipulator-induced node-edge operations, the input scene graphs dynamically describe global scene states, thus demanding adaptability from networks to accurately represent changing states. This status includes nodes appearing or disappearing and edges altering during scene creation. **Second**, it is crucial yet difficult when encapsulating both fine-grained node classes and diverse edge combinations into a network to be aware of global constraints. Existing solutions either simplify graphs by preserving only node sets, avoiding complex edge encoding in the diffusion while losing key structural relationships [53], or convert nodes and edges into tokens for transformer-based denoiser [32], which is effective for the layout generation from simple graphs but not scalable for larger ones due to exponential token growth. For a graph with  $P$  nodes and  $Q$  types of edges, the maximum token count grows with  $Q \cdot P!$ , making such a strategy hard to model intricate relations and computationally infeasible for large graphs. CommonScenes [58] collapses graphs to triplets (subject, predicate, object) and employs graph convolution to aggregate features for each node, which are then conditioned on individual denoising processes for shape generation. While this approach adapts well to varying graph sizes, each denoising process is isolated from the other, resulting in insufficient awareness of global shape states and, thus, undesired inter-object inconsistency. Additionally, its reliance on VAE and GAN for layout generation complicates synchronized training.

To fully utilize the controllability of scene graphs for scene generation and manipulation, the feature aggregation of multiple edges to the respective nodes is essential. In an attempt to flexibly encapsulate information from an indefinite number of nodes without losing global graph constraints, it can be helpful to allocate an individual denoising process per node and encourage information communication among all processes to achieve a common goal state. Integrating both concepts, CSG becomes promising with complex scene graphs by adhering to graph constraints and manipulating object-level diffusion. With *EchoScene*

we move a step further and introduce a scene-generative framework equipped with dual-branch diffusion models capable of producing object shapes and scene layouts simultaneously (see Fig. 2 for the pipeline). Each branch contains multiple denoising processes, with an equal number of nodes in the scene graph. The denoiser is weight-shared for all processes without introducing additional costs. To consolidate global awareness in every denoising process, we introduce an *information echo* scheme to exchange information at each denoising step in the respective branches, as illustrated in the middle of Fig. 1. Here, an information echo in the graph diffusion happens when a node first sends the denoising data and other features to an information exchange unit, which echoes back the aggregated features to the node as the condition signal for its denoiser. More clearly, for a single denoising process, the echo route is:  $\{current\ denoising\ input \rightarrow information\ exchange\ unit \rightarrow denoising\ conditioner\}$ .

In this design, EchoScene is able to generate and edit scenes from complex scene graphs with an indefinite number of nodes and edges. More importantly, the information echo scheme supports temporally global information exchange within both branches, thus continuously making the generation compliant with the scene graph description. For example, in the shape branch, shape echoes enable steady awareness of object appearance along the denoising steps for each generation process, bringing more inter-object consistency to the global style. Last but not least, the pure diffusion-based model dynamics enables synchronized training. EchoScene outperforms the previous state-of-the-art method on generation fidelity by a large margin and is more robust when handling graph manipulation. Moreover, we show that scenes generated by EchoScene are compatible with off-the-shelf texture generators, *e.g.*, SceneTex [9], facilitating potential downstream applications.

The main contributions of this work can be summarized as follows:

1. We present *EchoScene*, a scene generation method with a dual-branch diffusion model on dynamic scene graphs, to simultaneously generate layouts and shapes with more controllability.
2. We introduce an information echo scheme inside each branch of EchoScene that allows multiple denoising processes to exchange their denoising status among each other at each time step, bringing global awareness to each individual process.
3. We show in the experiments that the proposed generative framework achieves more generation fidelity, more robustness in handling graph manipulation, and effectively handles inter-object style consistency.

## 2 Related work

**Semantic Scene Graph.** Scene graphs have become a fundamental tool for semantic scene understanding, offering a structured and symbolic representation of scenes through nodes (objects) and edges (relationships) [21]. Their versatility spans various domains, including image manipulation, caption generation,

and visual question answering, demonstrating their capacity to enrich 2D image generation [14, 20, 27]. Beyond 2D, scene graphs have been extended to 3D scene understanding [56, 62], dynamic modeling [47], and notably, controllable scene synthesis [29], showcasing their adaptability in representing complex spatial and semantic relations in multi-dimensional spaces [26, 29, 40, 47, 56, 58, 61]. EchoScene proposes an indoor 3D scene generation method conditioned fully on semantic scene graphs coupled with a scene controlling ability.

**Diffusion Models.** Applications of diffusion models [50] have witnessed rapid expansion, from generating intricate images to modeling complex distributions in various data types [19, 52]. Recent advances have focused on improving their flexibility and realism, with significant efforts to enhance conditional generation and refine the models’ understanding of input conditions [4, 8, 25, 39, 59]. Score Distillation Sampling (SDS) was proposed to enable training Neural Radiance Fields through pretrained diffusion models [44]. A line of work addressed text-to-3D via SDS [34, 48], studied the compositionality of scenes via SDS [13, 17, 33, 43], and more recently via Gaussian Splatting [60]. These methods can create realistic 3D assets from text via a small number of images; however, they are limited in terms of number of objects within a scene while capturing their relationships. Comparably, a set of works used diffusion models to learn 3D shapes explicitly [12, 35]. Finally, diffusion models have recently been shown to be useful for graph generative processes [24, 54, 55]. EchoScene draws from the recent advances in diffusion models, where both shape and layout generation are through graph diffusion processes.

**Controllability in Scene Synthesis.** The quest for controllable scene synthesis has led to diverse methodologies, ranging from text descriptions to spatial layouts and probabilistic grammars for creating 3D scenes. Methods have explored generating scenes from images [11], texts [5], and layouts [1], employing various strategies such as autoregressive models and deep priors to achieve controllable outputs [22, 37, 42]. Recent work has also seen graph-based conditioning emerge as a powerful approach for controlled scene synthesis, enabling more precise manipulation of scene elements and their relations [29, 32, 53, 58, 61]. CommonScenes [58] came up with a scene graph conditioned scene generation method, which is based on a VAE and a GAN for capturing the layout and creating shapes via a latent diffusion model. Even though these methods have high quality in layout generation, their shape generation abilities are limited, as well as their manipulation quality still has room for improvement. EchoScene addresses these challenges through a diffusion-based layout generation with a novel information echo scheme.

### 3 Preliminary

**Scene Graph.** The scene graph we use is semantic scene graph [7], denoted as  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ , which serves as a structured representation of a visual scene. In such representation,  $\mathcal{V} = \{v_i \mid i = 1, \dots, N\}$  refers to the set of object nodes, while  $\mathcal{E} = \{e_{i \rightarrow j} \mid i, j = 1, \dots, N, i \neq j\}$  represents the set of directed edges

connecting each pair of nodes  $v_i \rightarrow v_j$ . As structured in the Fig. 2.A.1, each node  $v_i$  and edge  $e_{i \rightarrow j}$  can encompass various extensible attributes, *e.g.*, object categorical information.

**Scene Graph Convolution.** The scene graph convolution we use in this work, originally from [20] and applied in [15, 36, 58], is the fundamental module processing semantic scene graphs, as known as triplet-GCN. A typical  $K$ -layer triplet-GCN module is:

$$\begin{aligned} (\alpha_{v_i}^{(k)}, \beta_{e_{i \rightarrow j}}^{(k+1)}, \alpha_{v_j}^{(k)}) &= h_1(\beta_{v_i}^{(k)}, \beta_{e_{i \rightarrow j}}^{(k)}, \beta_{v_j}^{(k)}), k = 0, \dots, K-1, \\ \beta_{v_i}^{(k+1)} &= \alpha_{v_i}^{(k)} + h_2\left(\text{Avg}(\alpha_{v_j}^{(k)} \mid v_j \in N_G(v_i))\right), \end{aligned} \quad (1)$$

where  $k$  represents an individual layer in triplet-GCN, and  $h_1, h_2$  are MLP.  $N_G(v_i)$  encompasses all nodes that are directly connected to  $v_i$ , and **Avg** denotes the process of average pooling. The initial attributes of nodes and edges, denoted by  $(\beta_{v_i}^{(0)}, \beta_{e_{i \rightarrow j}}^{(0)}, \beta_{v_j}^{(0)})$ , able to be customized for specific usages.

**Contextual Graph.** CLIP features are proven to bring strong inter-object information to the scene graph [58]. Thus, a semantic scene graph evolves to a contextual graph by using the readily available CLIP text encoder to infer the semantic information for each node and each triplet in the graph to provide semantic anchors  $(p_i, p_{i \rightarrow j})$ , and each node and edge have their own learnable vectors  $(o_i, \tau_{i \rightarrow j})$ . As shown in the Fig. 2.A.1, the node and edge feature of a contextual graph is  $v_i := \{p_i, o_i\}$  and  $e_{i \rightarrow j} := \{p_{i \rightarrow j}, \tau_{i \rightarrow j}\}$ , respectively.

**Conditional Diffusion Models.** Diffusion models learn to estimate a target distribution through a progressive Markov Process of length  $T$  where a denoiser  $\varepsilon_\theta$  is trained to gradually remove noise from diffused noisy versions  $\mathbf{d}_t$  with  $t \in \{1, \dots, T\}$  of a data sample  $\mathbf{d}$  [19, 52]. Latent Diffusion Models (LDMs) [46] efficiently establish this process in latent space where a conditioner  $\mathbf{c}$  can guide the generation process [12]. Learning the denoising is reached by minimization of the objective

$$\mathcal{L}_{LDM} = \mathbb{E}_{\mathbf{d}, \varepsilon \sim \mathcal{N}(0,1), t} [\|\varepsilon - \varepsilon_\theta(\mathbf{d}_t, t, \mathbf{c})\|_2^2]. \quad (2)$$

Once the denoiser  $\varepsilon_\theta$  is learned, the inverse process can be established using Langevin dynamics [51].

## 4 Method

**Overview.** We present *EchoScene*, a method that accomplishes scene generation through layout and shape generation from scene graphs. EchoScene evolves the contextual graph to the latent space utilizing an encoder and a manipulator based on triplet-GCN, as shown in Fig. 2.A and as introduced in Sec. 4.1. Then, as shown in Fig. 2.B, it conditions latent nodes to layout branch (Fig. 2.B.1) and shape branch (Fig. 2.B.2) separately based on an information echo scheme defined in Sec. 4.2. In the layout branch, each diffusion process interacts with

each other through layout echo to acquire global layout awareness at each denoising step, so that the final layout generation is compliant with the scene graph description (see Sec. 4.3 and Fig. 3A). Similarly, in the shape branch, each diffusion process interacts with each other through shape echo to be aware of other shape appearances at each denoising step, so that the final generated shapes in the scene are consistent (see Sec. 4.4 and Fig. 3B). Moreover, our generated scenes are compatible with off-the-shelf texture generators, *e.g.*, SceneTex [9], making more photorealistic appearances.

#### 4.1 Graph Preprocessing

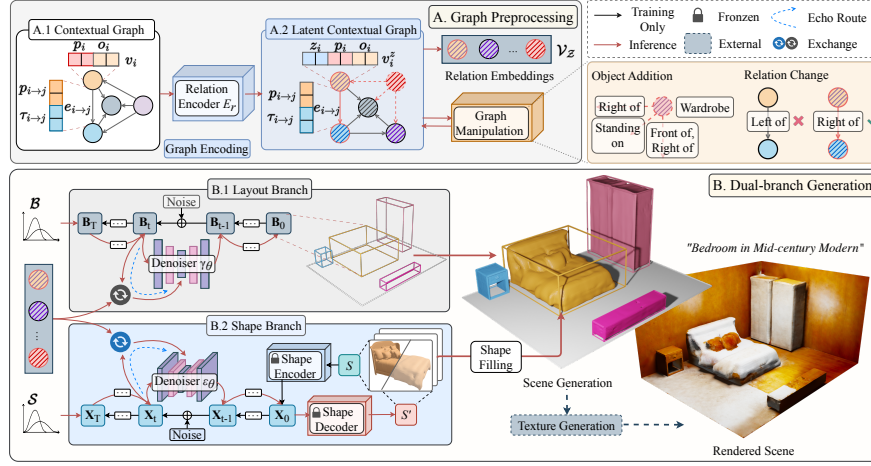
**Graph Encoding.** To make the layout and shape branches aware of the semantic and spatial information among the objects, we first encode the contextual graph to have latent relation embedding for each node by a triplet-GCN-based encoder  $E_r$ . In this case, we initialize the input embeddings of layer 0 with the features from the contextual graph, thus,  $(\beta_{v_i}^{(0)}, \beta_{e_{i \rightarrow j}}^{(0)}, \beta_{v_j}^{(0)}) = (v_i, e_{i \rightarrow j}, v_j)$  in Eq. 1. After the encoding, node features evolve to  $\mathcal{V}_Z = \{v_i^z \mid i = 1, \dots, N\}$ , where  $v_i^z$  consists of an explicit semantic representation  $v_i$  with an implicit embedding  $z_i$  encapsulating relation information with other nodes.

**Graph Manipulation.** The manipulation includes node addition and relation change, mimicking the user interaction. An example of manipulation is shown in Fig. 2. We inherit the manipulation functionality of Graph-to-3D [15] and CommonScenes [58] by setting up a graph manipulator, another triplet-GCN module, to adjust the graph in the latent space, but with an orthogonal strategy during training. We illustrate more details in the Supplementary Materials. After the manipulation, Node features in the scene graph are updated to  $\mathcal{V}_Z = \{v_i^z \mid i = 1, \dots, M, M \geq N\}$ . Subsequently, we send the updated graph to the diffusion-based layout branch and shape branch, where we set each node with an individual denoising process.

#### 4.2 Information Echo Scheme

Before introducing the layout and shape branch, we illustrate the core of the two branches—the information echo scheme in graph diffusion.

**Challenges from a Dynamic Graph.** While "Denoising Diffusion Probabilistic Models" (DDPM) [19] has significantly advanced the field of generative models by producing realistic and diverse content, their application to dynamic data structures, such as scene graphs, is limited. Scene graphs contain indefinite nodes and multiple edge combinations between nodes, describing compact and flexible relationships among objects. Some concurrent work based on DDPM [32, 53] to ours either models the object graph as a set or preserves it with a single edge between two nodes, where the representation ability of a scene graph degrades. Besides, they set a maximum number of tokens and model them in a single DDPM. As such, the local controllability of the generation is limited when users want to manipulate the specific elements. We believe a better solution is for each



**Fig. 2: Overview of EchoScene.** Our pipeline consists of graph preprocessing and two collaborative branches *Layout Branch* and *Shape Branch*. The details of two branches in one step are shown in Fig. 3. During inference, EchoScene evolves the contextual graph to the latent space, where a manipulator optionally adjust the graph by editing nodes and edges. Then, EchoScene samples a random noise from Gaussian Distribution  $\mathcal{B}$  and  $\mathcal{S}$  for both branches conditioned on the latent graph to generate shapes and layouts. Finally, the generated shapes are populated into layouts to synthesize the scenes. Moreover, an external texture generator [9] can be optionally utilized to provide a more photorealistic appearance.

node in the graph to be allocated with an individual denoising process, jointly forming a diffusion model for specific tasks, which, in our case, are shape and layout generation. Thus, the content generation is fully controllable by node and edge manipulation.

**Echo-Formed Condition.** However, even though each denoising process theoretically realizes more controllability, it also brings isolation problems. As each generation proceeds individually, there is no awareness of scene content during the denoising steps, which makes the generation inconsistent with global constraints in the graph. To achieve both highly controllable and consistent generation at the same time, we propose to couple the inverse iterative conditional diffusion mechanism on graph node features with a recurrent intermediate message exchange strategy over the graph  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ . In analogy to an echo in the real world, a node can send its information and receive it back along with information from other nodes during the steps of the inverse diffusion process.

The cornerstone of echo is the introduction of an information exchange unit  $U$  that enables dynamic and interactive diffusion processes among the elements of a dynamic graph. The incorporation of  $U$  on the conditioner during the inverse diffusion process allows each element within the group to share and receive information. At every denoising step, each diffusion process assembles the current denoising data  $\mathbf{d}_t^i$  with other available node attributes, where we choose latent

feature  $v_i^z$ , to construct a new node set  $\mathcal{V}_{\mathcal{D}_t} := \{f(\mathbf{d}_t^i, v_i^z, \pi(t)) \mid i = 1, \dots, M\}$ , where  $f(\cdot)$  depicts a feature assembling function, and we choose simple concatenation here.  $\pi(\cdot)$  illustrates the temporal process embedding function. Then, the process sends  $\mathcal{V}_{\mathcal{D}_t}$  to  $U$ , which comprehensively understands group dynamics according to graph edges  $\mathcal{E}$ , by subscribing and aggregating information from all processes. In our task, we set  $U$  based on triplet-GCN. The feature aggregation function  $U(\cdot)$  is based on Eq. 1, with input  $(\beta_{v_i}^{(0)}, \beta_{e_{i \rightarrow j}}^{(0)}, \beta_{v_j}^{(0)}) = (v_i^{\mathbf{d}_t}, e_{i \rightarrow j}, v_j^{\mathbf{d}_t})$ , where  $v_i^{\mathbf{d}_t}$  and  $v_j^{\mathbf{d}_t}$  are elements in  $\mathcal{V}_{\mathcal{D}_t}$ . The denoiser of each process receives the aggregated features for conditioning. In this form, one sending and one receiving step constitute an ‘information echo.’ Note that the Langevin dynamics here are different from the ones in a normal diffusion model, where we introduce the relationship of a group of denoising data into the condition  $\mathcal{C}_{\mathcal{D}_t}$ , making it temporarily changed and steadily introducing constraints to generate next states. With forward process variances  $\beta_t$  and  $\sigma_t = \sqrt{\beta_t}$  as well as the shorthand notations  $\alpha_t := 1 - \beta_t$  and  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$ , we define one conditional step in the denoising procedure as:

$$\begin{aligned} \mathbf{d}_{t-1}^i &= \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{d}_t^i - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \varepsilon_\theta(\mathbf{d}_t^i, \pi(t), \mathcal{C}_{\mathcal{D}_t}) \right) + \sigma_t \mathbf{z}, \\ \mathcal{C}_{\mathcal{D}_t} &= U(\mathcal{G}_{\mathcal{D}_t}), \mathcal{G}_{\mathcal{D}_t} = \{\mathcal{V}_{\mathcal{D}_t}, \mathcal{E}\}, \end{aligned} \quad (3)$$

where  $\mathbf{z} \sim \mathcal{N}(0, 1)$  if  $t > 1$ , else  $\mathbf{z} = 0$ . The information echo is repeated at every timestep, collectively forming the information over the scene graph diffusion.

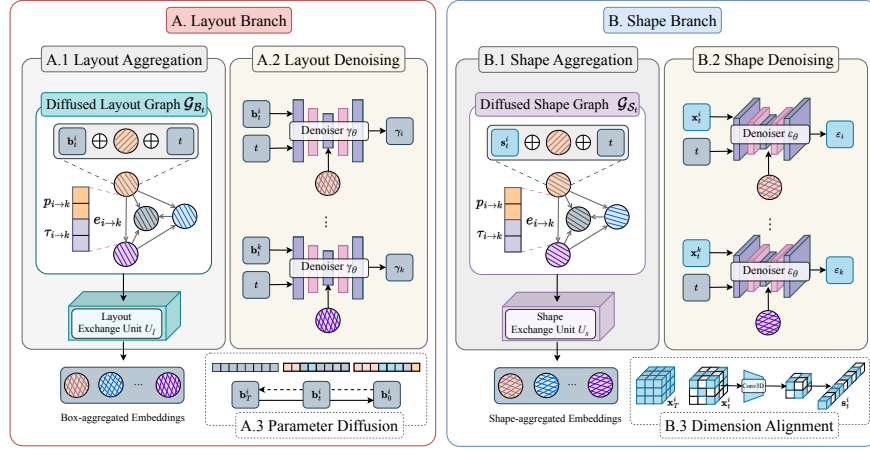
### 4.3 Layout Branch

As described in Sec. 4.2, we model the layout generation by setting each node with its own denoising process and encouraging them to interact with their diffused layout with each other at every denoising step.

**Layout Parametrization.** The scene layout is represented by object bounding boxes. Initially, each bounding box  $\mathbf{b}_0^i$  has 7 parameters, e.g., location  $(x, y, z)$ , size  $(l, h, w)$ , and a yaw angle  $\theta$ . Before training, we normalize location and size by their maximum and minimum values on each axis. For angles, we calculate its sine and cosine values. Thereby, the final representation contains 8 parameters:  $\mathbf{b}_0^i = \{x, y, z, l, h, w, \sin \theta, \cos \theta\}$ , as shown in Fig. 3A.3.

**Layout Echoes.** Since the bounding box generation needs to be compliant with the spatial constraints described in the scene graph, state observation from other nodes is needed to determine the bounding box of a specific node. In this branch, we encourage all nodes to communicate their current diffused bounding boxes  $\mathbf{B}_t := \{\mathbf{b}_t^i \mid i = 1, \dots, M\}$  in the graph by a layout exchange unit  $U \mapsto U_l$  to achieve global awareness at each time step. As shown in Fig. 3A.1, in the implementation, we substitute denoising data  $\mathbf{d}_t^i$  to the diffused bounding box  $\mathbf{b}_t^i$ , resulting in  $\mathcal{V}_{\mathcal{D}_t} \mapsto \mathcal{V}_{\mathcal{B}_t}$  and  $\mathcal{G}_{\mathcal{D}_t} \mapsto \mathcal{G}_{\mathcal{B}_t}$ . Thus, layout echoes happen at each time step by using Eq. 3, which are essential for the functionality of this branch.

**Training objective.** We follow a normal DDPM training routine, in which we set 1000 time steps for all diffusion processes with weight-shared  $\gamma_\theta$ . In each



**Fig. 3: One Step of Dual-Branch Information Echo.** For each time step, we encourage the layout (left) and shape (right) branches to exchange information within each branch for all objects in the same scene.

time step, the objective is to minimize the noise prediction errors supervised by  $\gamma$  sampled from Gaussian Distribution:

$$\mathcal{L}_{layout} = \mathbb{E}_{\mathbf{B}, \gamma \sim \mathcal{N}(0,1), t} [\|\gamma - \gamma_\theta(\mathbf{B}_t, \pi(t), U_l(\mathcal{G}_{B_t}))\|_2^2], \mathcal{G}_{B_t} = \{\mathcal{V}_{B_t}, \mathcal{E}\}. \quad (4)$$

#### 4.4 Shape Branch

As shown in Fig. 2B.2, we pretrain a VQ-VAE as a shape en-decoder and treat the latent codes  $\mathbf{X}_0$  in the bottleneck of the VQ-VAE as the ground truth of the LDM, as deployed in CommonScenes [58]. Similar to our layout branch, we introduce the information echo scheme to shape generation, enhancing the generation consistency.

**Isolation Problems.** Unlike layout generation, shape generation can be conditioned solely with the semantic information [12, 30] or relation information [58], which means the shape branch is already functional when allocating each node to a denoising process driven by semantics or latent relation embeddings. However, each process still lacks consistency control from the shape appearance aspect. In other words, each generation process only focuses on its own, while ignoring others' shape generation state throughout the whole denoising. Thus, the generations are isolated from each other, resulting in suboptimal inter-object consistency of the global style.

**Shape Echoes.** We solve the problem by introducing shape observation of other processes for each process, which is achieved by shape echoes. As shown in Fig. 3B.3, before shape echoes happen, we first feed diffused shape codes

$\mathbf{X}_t := \{\mathbf{x}_t^i, \mid i = 1, \dots, M\}$  to 3D convolutional layers and flatten to  $\mathbf{S}_t := \{\mathbf{s}_t^i, \mid i = 1, \dots, M\}$ , aligning the dimension with the node’s attribute. The meaning of the  $\mathbf{X}_t$  and  $\mathbf{S}_t$  becomes more obvious when the time step is closer to 0. Similar to the layout branch, we substitute denoising data  $\mathbf{d}_t^i$  to the diffused shape code  $\mathbf{s}_t^i$ , resulting in  $\mathcal{V}_{\mathcal{D}_t} \mapsto \mathcal{V}_{\mathcal{S}_t}$  and  $\mathcal{G}_{\mathcal{D}_t} \mapsto \mathcal{G}_{\mathcal{S}_t}$ . Thus, with a shape exchange unit  $U \mapsto U_s$ , shape echoes happen at each time step by using Eq. 3.

**Training objective.** For each latent diffusion process, the weight-shared denoiser  $\varepsilon_\theta$  takes time step  $t$ , shape codes  $\mathbf{X}_t$ , and shape-aggregated feature as input to predict the noise, which is supervised by  $\varepsilon$  sampled from Gaussian Distribution. The objective of the training is to minimize the noise prediction errors:

$$\mathcal{L}_{shape} = \mathbb{E}_{\mathbf{X}, \varepsilon \sim \mathcal{N}(0,1), t} [\|\varepsilon - \varepsilon_\theta(\mathbf{X}_t, \pi(t), U_s(\mathcal{G}_{\mathcal{S}_t}))\|_2^2], \mathcal{G}_{\mathcal{S}_t} = \{\mathcal{V}_{\mathcal{S}_t}, \mathcal{E}\} \quad (5)$$

#### 4.5 Dual-Branch Joint Training

Since both branches are based on diffusion models, the framework supports end-to-end synchronized training optimized by two losses weighted by  $\lambda_1, \lambda_2$ :

$$\mathcal{L} = \lambda_1 \mathcal{L}_{layout} + \lambda_2 \mathcal{L}_{shape}. \quad (6)$$

## 5 Experiments

**Dataset.** We conduct our experiments on SG-FRONT dataset [58], which provides scene-graph annotations for the high-quality 3D-FRONT [16] with household environments. SG-FRONT contains 15 relationship types and 45K object instances from three types of scenes.

**Metrics.** We follow the metrics used in [32, 57, 58], to measure the *fidelity and diversity* of generated scenes, where we adopt Fréchet Inception Distance (FID) [18], FID<sub>CLIP</sub> [28] and Kernel Inception Distance (KID) [3] metrics [42]. To measure the *scene graph consistency*, we follow the scene graph constraints [15], which measure the accuracy of a set of relations on a generated layout. To measure the *shape consistency and diversity* following [58], we test dining rooms, a typical scenario in which tables and chairs should be in suits.

**Baselines.** For retrieval-based methods *i.e.*, a layout generation network *3D-SLN* [36], a progressive method to add objects one-by-one designed in [15], *Graph-to-Box* from [15], *DiffuScene* [53] with text conditioning, and *InstructScene* [32] with graph conditioning. For generative methods, *Graph-to-3D* [15], modeling shape and layouts both in VAE. CommonScenes [58], which model scene layouts through VAE and generate each object via LDM conditioned on relation embeddings. *EchoLayout/CommonLayout+SDFusion* that we design stacking the layout branch and a text-to-shape generation model SDFusion [12] sequentially. More details are in Supplementary Material.

**Implementation Details.** The training, evaluation, and visualization are carried out on a single NVIDIA A40 GPU with 40 GB of memory. Training is optimized with AdamW with an initial learning rate of 1e-4. The values of  $\lambda_1, \lambda_2$  are set to 1.0. More details are provided in Supplementary Material.

Method	Shape Representation	Bedroom			Living room			Dining room		
		FID	FID <sub>CLIP</sub>	KID	FID	FID <sub>CLIP</sub>	KID	FID	FID <sub>CLIP</sub>	KID
3D-SLN [36]	Retrieval	57.90	5.45	3.85	77.82	7.02	3.65	69.13	7.99	6.23
Progressive [15]		58.01	5.67	7.36	79.84	7.41	4.24	71.35	8.28	6.21
Graph-to-Box [15]		54.61	5.26	2.93	78.53	6.88	3.32	67.80	7.75	6.30
CommonLayout [58]		52.69	5.22	2.82	76.52	6.58	2.08	65.10	7.55	6.11
DiffuScene [53]		52.02	5.01	2.52	81.61	7.52	1.23	65.90	7.39	0.09
InstructScene [32]		45.40	3.87	1.06	75.83	6.98	4.15	61.56	6.49	4.90
<b>EchoLayout (Ours)</b>		46.53	4.24	0.33	75.54	6.35	1.60	59.66	6.24	2.63
Graph-to-3D [15]	DeepSDF [41]	63.72	6.01	17.02	82.96	7.80	11.07	72.51	7.25	12.74
CommonLayout+SDFusion [12]		68.08	5.61	18.64	85.38	7.23	10.04	64.02	6.92	5.08
EchoLayout+SDFusion [12]		57.68	4.96	10.54	83.66	6.83	9.62	65.55	7.02	4.99
CommonScenes [58]		57.68	4.86	6.59	80.99	7.05	6.39	65.71	7.04	5.47
<b>EchoScene (Ours)</b>		48.85	4.26	1.77	75.95	6.73	0.60	62.85	6.28	1.72

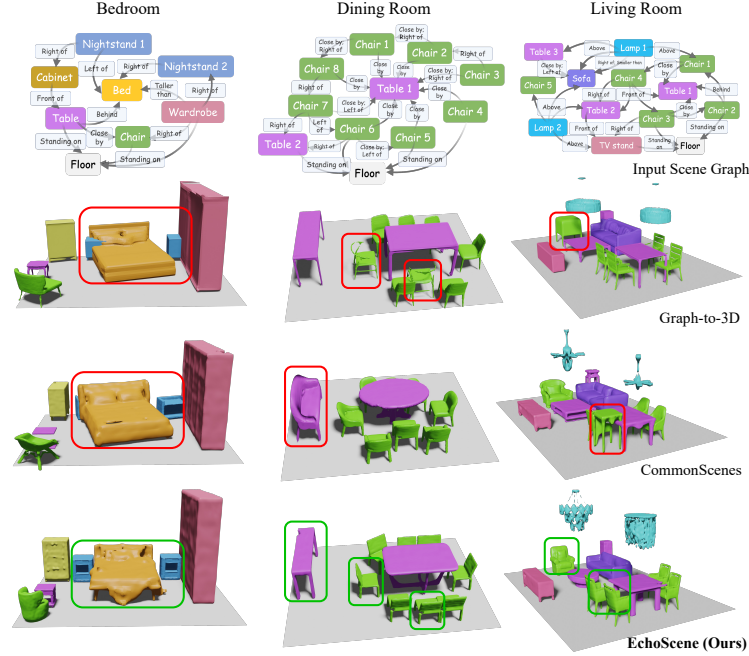
**Table 1: Scene generation realism** as measured by FID, FID<sub>CLIP</sub> and KID ( $\times 0.001$ ) scores at  $256^2$  pixels between the top-down rendering of generated and real scenes, following [53, 58] (lower is better). We color the **best** and the **second**. Two main rows are separated with respect to the reliance on an external shape database. *InstructScene\** represents its 3D layout decoder solely. Both *EchoLayout* and *CommonLayout* refer to the single layout branch.

### 5.1 Scene Fidelity

As a generation task, the fidelity of scene synthesis is essential to measure.

**Quantitative results.** We perform two types of scene synthesis, which are shape retrieval-based and generation-based synthesis, and provide FID / FID<sub>CLIP</sub> / KID in Tab. 1. For retrieval, our layout branch *EchoLayout* performs better than the layout branch of the previous SoTA CommonLayout [58] across all metrics, and concurrent work DiffuScene [53] on most of the metrics. In terms of shape generation, our 3D scene generation version *EchoScene* shows clear advantages against CommonScenes [58], for example, improving FID by 15%, FID<sub>CLIP</sub> by 12% and KID by 73% in the bedroom generation. Notably, although retrieval-based methods overall score better due to the mesh alignment with the test database, EchoScene remains competitive, showcasing generative effectiveness. Moreover, our text-to-shape baseline *EchoLayout+SDFusion* outperforms the equivalent CommonLayout+SDFusion. Despite under the issue of shape inconsistency, it still achieves similar results as CommonScenes. This also underscores the effectiveness of our layout branch, driven by layout echoes, in generating more coherent scene layouts. For shape consistency, we test the Chamfer Distance (CD) between two generated objects that are supposed to be identical (*i.e.* shape consistency [58]). We report common objects in three room types in Tab. 3 showing that our shape echoes alleviate the isolation problem in CommonScenes and thus bring better inter-object consistency.

**Qualitative results.** Graph-to-3D [15] and CommonScenes [58] have the same settings as EchoScene, which belong to (semi-/full-)generative models conditioned on scene graphs for the whole scene generation. We provide a comparison of these methods for each different room type in Fig. 4. For example, in the bedroom, beds and nightstands of Graph-to-3D and CommonScenes are in twisted poses, indicating VAE-based layout generator does not learn angles well, while



**Fig. 4: Comparisons with other generative methods.** Input scene graphs have more edges between two nodes than the ones visualized here. Red rectangles highlight the inconsistent generation. (Zoom for details)

our diffusion-based architecture can provide neat results. In the dining room, Graph-to-3D generates chairs in different styles, which breaks the inter-object consistency. Although CommonScenes can generate chairs in a suit, the generated table on the left is problematic, having a chair-like shape, which makes an unsatisfactory global appearance as well. Ours achieves higher shape consistency.

## 5.2 Graph Constraints

In this part, we evaluate the layout generation performance with respect to scene graph constraints. We compare against the closest methods: 3D-SLN [36], Graph-to-3D [15], InstructScene [32], and CommonScenes [58].

We optionally manipulate scene graphs in the latent by either adding a node and relevant edges to the graph or changing relations between two nodes, as illustrated in Fig. 2. As shown in the top and middle row of Tab. 2, for both node addition and relation change, EchoLayout and EchoScene perform better on most metrics. 3D-SLN, CommonScenes, and Graph-to-3D, for instance, loose the constraints **smaller/bigger than** and **close by** after graph manipulation, while ours largely keeps it. We believe that the improvement is caused by our scene graph diffusion strategy, where a graph-consistent sample is generated within a realistic appearance distribution.

Method	Shape Representation	Mode	left/ right	front/ behind	smaller/ larger	taller/ shorter	close by symmetrical
3D-SLN [36]	Retrieval	Change	0.89	0.90	0.55	0.58	0.10
Progressive [15]			0.89	0.89	0.52	0.55	0.08
Graph-to-Box [15]			0.91	0.91	0.86	0.91	0.66
CommonLayout [58]			0.91	0.92	0.86	0.92	0.70
<b>EchoLayout (Ours)</b>			0.94	0.93	0.92	0.92	0.72
Graph-to-3D [15]	DeepSDF [41]	Change	0.91	0.92	0.86	0.89	0.69
CommonScenes [58]			0.91	0.92	0.86	0.91	0.69
<b>EchoScene (Ours)</b>			0.94	0.96	0.92	0.93	0.74
3D-SLN [36]	Retrieval	Addition	0.92	0.92	0.56	0.58	0.05
Progressive [15]			0.92	0.91	0.53	0.54	0.02
Graph-to-Box [15]			0.94	0.93	0.90	0.94	0.67
CommonLayout [58]			0.95	0.95	0.90	0.94	0.73
<b>EchoLayout (Ours)</b>			0.98	0.99	0.97	0.96	0.74
Graph-to-3D [15]	DeepSDF [41]	Addition	0.94	0.95	0.91	0.93	0.63
CommonScenes [58]			0.95	0.95	0.91	0.95	0.70
<b>EchoScene (Ours)</b>			0.98	0.99	0.96	0.96	0.76
3D-SLN [36]	Retrieval	None	0.97	0.99	0.95	0.91	0.72
Progressive [15]			0.97	0.99	0.95	0.82	0.69
Graph-to-Box [15]			0.98	0.99	0.96	0.95	0.72
CommonLayout [58]			0.98	0.99	0.97	0.95	0.74
InstructScene* [32]			0.99	0.99	0.85	0.86	0.52
<b>EchoLayout (Ours)</b>			1.00	0.99	0.95	0.96	0.74
Graph-to-3D [15]	DeepSDF [41]	None	0.98	0.99	0.97	0.95	0.74
CommonScenes [58]			0.98	1.00	0.97	0.95	0.77
<b>EchoScene (Ours)</b>			0.98	0.99	0.96	0.96	0.74

**Table 2: Scene graph constraints** (higher is better). Top: Relationship change mode. Middle: Node addition mode. Bottom: No manipulation (*i.e.*, generation only).

The decrease in **symmetrical** category compared with CommonScenes is likely caused by the same effect. Inspection of the data reveals that **symmetrical** is an extremely rare label in the supplementary of CommonScenes [58], with an occurrence of only 0.9%. Thus, the learned part of the distribution is underrepresented and less accurate. The residual echo correction diffuses the latent code towards the learned distribution, resulting in a slight variation. In the bottom row, we report the results without manipulation. All baselines show good performance, yet our EchoLayout and EchoScene are still competitive on their own track. Notably, even though InstructScene shows impressive performance on the global generation fidelity shown in Tab. 1, it struggles to maintain the multiple graph constraints simultaneously, as we described in Sec. 1, while our information echo scheme can sufficiently aggregate relations and condition each process to achieve better results. Distribution-dependant effects from data equally manifest in the results where we can see some drop for **symmetrical** again. Interestingly, the effect is less pronounced for a lower parametric final task such as layout generation. This can be explained by the fact that the variability of bounding boxes is smaller compared to shape variations.

### 5.3 Ablation study

We ablate three parts and report the performance in Tab. 4. Firstly, we ablate  $\pi(t)$  to check the influence on the model without explicit time information. We



**Fig. 5: Off-the-shelf texture creation.** A bedroom (top) and living room (bottom) generated by EchoScene and textured in different styles by SceneTex [9].

Method	Bedroom		Living room		Dining room	
	Ward.	N.stand	Chair	Table	Chair	Table
CommonScenes [58]	0.61	2.69	6.64	11.75	1.96	9.04
<b>EchoScene</b>	<b>0.14</b>	<b>1.68</b>	<b>0.99</b>	<b>3.02</b>	<b>1.75</b>	<b>1.26</b>

Ablation	FID	FID <sub>CLIP</sub>	KID	mSG
Ours w/o $\pi(t)$	40.55	<b>3.14</b>	1.69	0.87
Ours w/o shape echoes	46.88	3.81	4.17	<b>0.88</b>
Ours with concat	48.32	3.82	6.87	0.87
<b>Ours</b>	<b>39.74</b>	<b>3.14</b>	<b>1.24</b>	<b>0.88</b>

**Table 3: Inter-object Consistency.** The consistent object shapes within a scene are indicated by low CD values ( $\times 0.001$ ).

**Table 4: Ablations under three circumstances.** mSG means average graph constraints.

observe a marginal decrease, indicating even without  $\pi(t)$ , the model still learns temporal information through every denoising step. Secondly, we exclude the shape echoes to degrade the shape branch. We clearly see a drop in the fidelity evaluation, where ‘mSG’ focusing on the layouts is not influenced. Thirdly, we substitute cross-attention for concatenation to check if the condition way influences the information echo scheme. We also observe a drop in all metrics.

## 5.4 Application

The limitation of EchoScene is that it can generate semantic yet only textureless scenes, which prohibits it from downstream tasks requiring photorealistic textures. Yet, the limitation can be alleviated by directly equipping EchoScene with a subsequent off-the-shelf texture generator. Here, we demonstrate that textured scenes can be generated via SceneTex [9] in Fig. 5.

## 6 Conclusion

We have introduced EchoScene, a scene-generative method based on dual-branch diffusion models, which are layout branch and shape branch. To make two branches functional, we have proposed an information scheme encouraging information exchange within both branches at each denoising step, generating final scenes that are compliant with the scene graph description. Experiments show that EchoScene can achieve higher generation fidelity and more inter-object consistency of the global scene style.

## Acknowledgements

We are grateful to Google University Relationship GCP Credit Program for supporting this work by providing computational resources.

## References

1. Bahmani, S., Park, J.J., Paschalidou, D., Yan, X., Wetzstein, G., Guibas, L., Tagliasacchi, A.: CC3D: Layout-conditioned generation of compositional 3D scenes. In: ICCV (2023)
2. Bautista, M.A., Guo, P., Abnar, S., Talbott, W., Toshev, A., Chen, Z., Dinh, L., Zhai, S., Goh, H., Ulbricht, D., et al.: Gaudi: A neural architect for immersive 3D scene generation. In: NeurIPS (2022)
3. Bińkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying MMD GANs. In: ICLR (2018)
4. Brooks, T., Holynski, A., Efros, A.A.: InstructPix2Pix: Learning to follow image editing instructions. In: CVPR (2023)
5. Chang, A., Monroe, W., Savva, M., Potts, C., Manning, C.D.: Text to 3D scene generation with rich lexical grounding. In: ACL (2015)
6. Chang, A., Savva, M., Manning, C.D.: Learning spatial knowledge for text to 3D scene generation. In: EMNLP (2014)
7. Chang, X., Ren, P., Xu, P., Li, Z., Chen, X., Hauptmann, A.: A comprehensive survey of scene graphs: Generation and application. T-PAMI (2021)
8. Chefer, H., Alaluf, Y., Vinker, Y., Wolf, L., Cohen-Or, D.: Attend-and-Excite: Attention-based semantic guidance for text-to-image diffusion models. In: SIG-GRAPH (2023)
9. Chen, D.Z., Li, H., Lee, H.Y., Tulyakov, S., Nießner, M.: SceneTex: High-quality texture synthesis for indoor scenes via diffusion priors. In: CVPR (2024)
10. Chen, D.Z., Siddiqui, Y., Lee, H.Y., Tulyakov, S., Nießner, M.: Text2Tex: Text-driven texture synthesis via diffusion models. In: ICCV (2023)
11. Chen, Z., Wang, G., Liu, Z.: SceneDreamer: Unbounded 3D scene generation from 2d image collections. T-PAMI (2023)
12. Cheng, Y.C., Lee, H.Y., Tulyakov, S., Schwing, A., Gui, L.: SDFusion: Multimodal 3D shape completion, reconstruction, and generation. In: CVPR (2023)
13. Cohen-Bar, D., Richardson, E., Metzger, G., Giryas, R., Cohen-Or, D.: Set-the-Scene: Global-local training for generating controllable nerf scenes. In: ICCV Workshop (2023)
14. Dhano, H., Farshad, A., Laina, I., Navab, N., Hager, G.D., Tombari, F., Rupperecht, C.: Semantic image manipulation using scene graphs. In: CVPR (2020)
15. Dhano, H., Manhardt, F., Navab, N., Tombari, F.: Graph-to-3D: End-to-end generation and manipulation of 3D scenes using scene graphs. In: ICCV (2021)
16. Fu, H., Cai, B., Gao, L., Zhang, L.X., Wang, J., Li, C., Zeng, Q., Sun, C., Jia, R., Zhao, B., et al.: 3D-FRONT: 3D furnished rooms with layouts and semantics. In: ICCV (2021)
17. Gao, G., Liu, W., Chen, A., Geiger, A., Schölkopf, B.: GraphDreamer: Compositional 3D scene synthesis from scene graphs. In: CVPR (2024)
18. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017)

19. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020)
20. Johnson, J., Gupta, A., Fei-Fei, L.: Image generation from scene graphs. In: CVPR (2018)
21. Johnson, J., Krishna, R., Stark, M., Li, L.J., Shamma, D.A., Bernstein, M.S., Fei-Fei, L.: Image retrieval using scene graphs. In: CVPR (2015)
22. Jyothi, A.A., Durand, T., He, J., Sigal, L., Mori, G.: LayoutVAE: Stochastic scene layout generation from a label set. In: ICCV (2019)
23. Kenton, J.D.M.W.C., Toutanova, L.K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: NAACL-HLT (2019)
24. Kong, L., Cui, J., Sun, H., Zhuang, Y., Prakash, B.A., Zhang, C.: Autoregressive diffusion model for graph generation. In: ICML (2023)
25. Kong, X., Liu, S., Lyu, X., Taher, M., Qi, X., Davison, A.J.: Eschnet: A generative model for scalable view synthesis. In: CVPR (2024)
26. Kong, X., Yang, X., Zhai, G., Zhao, X., Zeng, X., Wang, M., Liu, Y., Li, W., Wen, F.: Semantic graph based place recognition for 3D point clouds. In: IROS (2020)
27. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual Genome: Connecting language and vision using crowdsourced dense image annotations. IJCV (2017)
28. Kynkäänniemi, T., Karras, T., Aittala, M., Aila, T., Lehtinen, J.: The role of imagenet classes in fréchet inception distance. In: ICLR (2023)
29. Li, M., Patil, A.G., Xu, K., Chaudhuri, S., Khan, O., Shamir, A., Tu, C., Chen, B., Cohen-Or, D., Zhang, H.: GRAINS: Generative recursive autoencoders for indoor scenes. TOG (2019)
30. Li, M., Duan, Y., Zhou, J., Lu, J.: Diffusion-SDF: Text-to-shape via voxelized diffusion. In: CVPR (2023)
31. Li, X., Zhang, Y., Ye, X.: DrivingDiffusion: Layout-guided multi-view driving scene video generation with latent diffusion model (2023)
32. Lin, C., MU, Y.: InstructScene: Instruction-driven 3D indoor scene synthesis with semantic graph prior. In: ICLR (2024)
33. Lin, Y., Bai, H., Li, S., Lu, H., Lin, X., Xiong, H., Wang, L.: CompoNeRF: Text-guided multi-object compositional nerf with editable 3D scene layout (2023)
34. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3D object. In: ICCV (2023)
35. Liu, Z., Feng, Y., Black, M.J., Nowrouzezahrai, D., Paull, L., Liu, W.: MeshDiffusion: Score-based generative 3D mesh modeling. In: ICLR (2023)
36. Luo, A., Zhang, Z., Wu, J., Tenenbaum, J.B.: End-to-end optimization of scene layout. In: CVPR (2020)
37. Ma, R., Patil, A.G., Fisher, M., Li, M., Pirk, S., Hua, B.S., Yeung, S.K., Tong, X., Guibas, L., Zhang, H.: Language-driven synthesis of 3D scenes from scene databases. TOG (2018)
38. Mandlekar, A., Nasiriany, S., Wen, B., Akinola, I., Narang, Y., Fan, L., Zhu, Y., Fox, D.: MimicGen: A data generation system for scalable robot learning using human demonstrations. In: CoRL (2023)
39. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: SDEdit: Guided image synthesis and editing with stochastic differential equations. In: ICLR (2022)
40. Özsoy, E., Örnek, E.P., Eck, U., Tombari, F., Navab, N.: Multimodal semantic scene graphs for holistic modeling of surgical procedures. In: Arxiv (2021)
41. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: CVPR (2019)

42. Paschalidou, D., Kar, A., Shugrina, M., Kreis, K., Geiger, A., Fidler, S.: ATISS: Autoregressive transformers for indoor scene synthesis. In: NeurIPS (2021)
43. Po, R., Wetzstein, G.: Compositional 3D scene generation using locally conditioned diffusion. In: 3DV (2023)
44. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: DreamFusion: Text-to-3D using 2d diffusion. In: ICLR (2023)
45. Pronovost, E., Ganesina, M.R., Hendy, N., Wang, Z., Morales, A., Wang, K., Roy, N.: Scenario Diffusion: Controllable driving scenario generation with diffusion. In: NeurIPS (2023)
46. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
47. Rosinol, A., Gupta, A., Abate, M., Shi, J., Carlone, L.: 3D Dynamic Scene Graphs: Actionable spatial perception with places, objects, and humans. In: RSS (2020)
48. Shi, Y., Wang, P., Ye, J., Mai, L., Li, K., Yang, X.: MVDream: Multi-view diffusion for 3D generation. In: ICLR (2024)
49. Simsar, E., Tonioni, A., Örnek, E.P., Tombari, F.: LatentSwap3D: Semantic edits on 3D image gans. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops. pp. 2899–2909 (October 2023)
50. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: ICML (2015)
51. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* **32** (2019)
52. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: ICLR (2021)
53. Tang, J., Nie, Y., Markhasin, L., Dai, A., Thies, J., Nießner, M.: DiffuScene: Scene graph denoising diffusion probabilistic model for generative indoor scene synthesis. In: CVPR (2024)
54. Verma, T., De, A., Agrawal, Y., Vinay, V., Chakrabarti, S.: VarScene: A deep generative model for realistic scene graph synthesis. In: ICML (2022)
55. Vignac, C., Krawczuk, I., Siraudin, A., Wang, B., Cevher, V., Frossard, P.: DiGress: Discrete denoising diffusion for graph generation. In: International Conference on Learning Representations (ICLR) (2023)
56. Wald, J., Dhano, H., Navab, N., Tombari, F.: Learning 3D semantic scene graphs from 3D indoor reconstructions. In: CVPR (2020)
57. Zhai, G., Cai, X., Huang, D., Di, Y., Manhardt, F., Tombari, F., Navab, N., Busam, B.: SG-Bot: Object rearrangement via coarse-to-fine robotic imagination on scene graphs. In: ICRA (2024)
58. Zhai, G., Örnek, E.P., Wu, S.C., Di, Y., Tombari, F., Navab, N., Busam, B.: CommonScenes: Generating commonsense 3D indoor scenes with scene graphs. In: NeurIPS (2023)
59. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: ICCV (2023)
60. Zhou, X., Ran, X., Xiong, Y., He, J., Lin, Z., Wang, Y., Sun, D., Yang, M.H.: GALA3D: Towards text-to-3D complex scene generation via layout-guided generative gaussian splatting (2024)
61. Zhou, Y., While, Z., Kalogerakis, E.: SceneGraphNet: Neural message passing for 3D indoor scene augmentation. In: ICCV (2019)
62. Özsoy, E., Örnek, E.P., Eck, U., Czempiel, T., Tombari, F., Navab, N.: 4D-OR: Semantic scene graphs for or domain modeling. In: MICCAI (2022)