

GeoSceneGraph: Geometric Scene Graph Diffusion Model for Text-guided 3D Indoor Scene Synthesis

Antonio Ruiz^{1,2}, Tao Wu¹, Andrew Melnik², Qing Cheng¹, Xuqin Wang¹,
Lu Liu¹, Yongliang Wang¹, Yanfeng Zhang¹, Helge Ritter²

¹Huawei Technologies

²Center for Cognitive Interaction Technology (CITEC), Bielefeld University

Keywords: 3D Scene Generation, Diffusion Models, Graph Neural Networks, Embodied AI, Extended Reality

Abstract

Methods that synthesize indoor 3D scenes from text prompts have wide-ranging applications in film production, interior design, video games, virtual reality, and synthetic data generation for training embodied agents. Existing approaches typically either train generative models from scratch or leverage vision-language models (VLMs). While VLMs achieve strong performance, particularly for complex or open-ended prompts, smaller task-specific models remain necessary for deployment on resource-constrained devices such as extended reality (XR) glasses or mobile phones. However, many generative approaches that train from scratch overlook the inherent graph structure of indoor scenes, which can limit scene coherence and realism. Conversely, methods that incorporate scene graphs either demand a user-provided semantic graph, which is generally inconvenient and restrictive, or rely on ground-truth relationship annotations, limiting their capacity to capture more varied object interactions. To address these challenges, we introduce GeoSceneGraph, a method that synthesizes 3D scenes from text prompts by leveraging the graph structure and geometric symmetries of 3D scenes, without relying on predefined relationship classes. Despite not using ground-truth relationships, GeoSceneGraph achieves performance comparable to methods that do. Our model is built on Equivariant Graph Neural Networks (EGNNs), but existing EGNN approaches are typically limited to low-dimensional conditioning and are not designed to handle complex modalities such as text. We propose a simple and effective strategy for conditioning EGNNs on text features, and we validate our design through ablation studies.

1. Introduction

The automatic synthesis of high-quality 3D indoor scenes is a problem that presents significant challenges in computer vision and computer graphics. Key challenges that need to be addressed include: 1) generating realistic and diverse 3D indoor scenes, where the aesthetics of objects are visually appealing, and their arrangement is meaningful, reflecting real-world functionality and spatial organization; and 2) ensuring the model can process control signals to meet user-specific requirements. This problem is highly relevant, as automating scene generation can significantly reduce the time and cost of producing computer-generated imagery (CGI) for applications such as film production, interior design, video games, and extended reality (XR), a term that encompasses augmented reality (AR), mixed reality (MR), and virtual reality (VR). Moreover, synthetic 3D scenes can provide rich ground-truth data for training models in 3D scene understanding and reconstruction, and can also serve as environments for training embodied agents.

Recently, this problem has been addressed by deep generative methods that employ variational autoencoders (VAEs) (Kingma, 2013), generative adversarial networks (GANs) (Goodfellow et al., 2020), autoregressive models or diffusion models (Sohl-Dickstein et al., 2015, Ho et al., 2020). Approaches based on GANs (Zhang et al., 2020, Yang et al., 2021b) are able to synthesize 3D scenes, but they often struggle with scene diversity, and the generation process is difficult to control. VAE-based methods (Purkait et al., 2020, Yang et al., 2021a) offer better diversity in the generated samples, but the overall quality tends to be lower. On the other hand, auto-regressive models (Paschalidou et al., 2021, Nie et al., 2023, Wang et al., 2021) fail to capture semantic relationships between objects, and they are prone

to accumulating prediction errors. Although some diffusion-based methods (Tang et al., 2024, Hu et al., 2024, Maillard et al., 2025) tackle these issues, they do not incorporate the graph structure of the 3D scene, which adversely impacts their performance. By contrast, other diffusion-based approaches (Zhai et al., 2023, Zhai et al., 2024, Lin and Mu, 2024) show that accounting for the inherent graph structure of the scene enhances both quality and controllability. Nonetheless, methods relying on the scene graph as a direct control signal (Zhai et al., 2023, Zhai et al., 2024) tend to be less user-friendly. InstructScene (Lin and Mu, 2024) overcomes these challenges by allowing text control signals and generating semantic scene graphs with a two-stage approach. Nevertheless, its semantic scene graph generation (first stage) relies on a static and predefined vocabulary of semantic relations for the graph edges, limiting its ability to handle open-vocabulary relationships. Previous methods have also overlooked the inherent geometric symmetries of 3D scenes, such as translations and rotations. However, recent works on 3D molecule generation (Hooeboom et al., 2022, Xu et al., 2023) demonstrate that leveraging these symmetries can enhance performance and improve generalization.

Vision-language models (VLMs) (Feng et al., 2023, Çelen et al., 2024, Yang et al., 2024, Sun et al., 2025, Pun et al., 2025, Chen et al., 2025, Bucher and Armeni, 2025) and agent-based methods (Yang et al., 2025, Ling et al., 2025, Ran et al., 2025, Berdoz et al., 2026, Xia et al., 2026) have demonstrated significant success in addressing this problem, particularly with open-ended prompts. However, high computational costs and inference latency limit their deployment on resource-constrained and offline devices. Efficient, smaller-scale models therefore remain relevant. Consequently, as our method belongs to the

class of task-specific approaches, we focus the remainder of this paper on those rather than VLM-based methods.

To address shortcomings of previous task-specific methods, we propose GeoSceneGraph, a novel diffusion-based approach that incorporates the inherent graph structure and geometric symmetries of 3D scenes without relying on ground-truth object relations from a static, predefined vocabulary. In addition, we introduce a text-conditioning strategy for Equivariant Diffusion Models (Hoogeboom et al., 2022) that use Equivariant Graph Neural Networks (EGNN) (Satorras et al., 2021). This strategy first fuses input feature vectors with the diffusion time-step condition through a ResNet (He et al., 2016) and self-attention Transformer (Vaswani et al., 2017) blocks, then integrates the text context into the message passing steps of the EGNN.

Our main contributions can be summarized as follows:

- We introduce a method for text-guided 3D indoor scene synthesis based on Equivariant Diffusion Models (Hoogeboom et al., 2022). Our approach leverages the scene graph structure and its geometric symmetries without relying on ground truth semantic edges. Despite not relying on ground truth semantic edges, our method still achieves competitive performance compared to approaches that use them.
- We propose a novel conditioning mechanism for Equivariant Diffusion Models, which is able to handle complex control signals such as text.
- We present various ablation studies exploring different conditioning strategies that validate our design choices.

2. Method

In this section we describe our method. Following previous works (Paschalidou et al., 2021, Tang et al., 2024, Lin and Mu, 2024), we denote a collection of 3D scenes with the set $\mathcal{S} := \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_M\}$. Then, we define a 3D scene \mathcal{S}_k as a collection of 3D objects, $\mathcal{S}_k := \{o_{k1}, o_{k2}, \dots, o_{kN_k}\}$ arranged according to a specific spatial layout relative to a world coordinate system, where the origin is placed at the center of the floor. Each object o_{ki} is associated with layout and appearance features. The layout features include $x_{ki} \in \mathbb{R}^3$, representing the position of the centroid of the object in the world coordinate system, $\theta_{ki} \in [0, 2\pi]$ (or its unit circle representation $(\cos \theta_{ki}, \sin \theta_{ki})$) indicates the yaw-axis rotation angle, restricting the object rotation to the vertical axis, and the size of its bounding box $s_{ki} \in \mathbb{R}^3$. The appearance features consist of the object class $c_{ki} \in \{1, \dots, K_c\}$, where K_c is the number of object classes in the dataset, and shape code $f_{ki} \in \mathbb{R}^d$. The shape codes are extracted from OpenCLIP ViT-bigG/14 embeddings (Liu et al., 2023), and their dimensionality is reduced with a VAE (see more details in Section 2.1). Given these features, a 3D scene \mathcal{S}_k , can be constructed by retrieving objects from a database with prebuilt 3D object models using a two-step filtering process. First, the selection is made based on the class c_{ki} , and then, on the nearest neighbor of the shape feature f_{ki} . Once the objects are chosen, the respective layout features are applied, i.e., object o_{ki} is translated to position x_{ki} , rotated by θ_{ki} about the vertical axis, and resized according to s_{ki} . We refer to Figure 1 to see the pipeline of our method.

Given a collection of 3D scenes \mathcal{S} , we can learn a parameterized generative model $p_\phi(\mathcal{S})$ that generates both appearance

and layout features, and then use the process described above to synthesize a novel 3D scene. Furthermore, to incorporate text control signals, we can also learn a conditional generative model $p_\phi(\mathcal{S}|y)$, where y denotes the user text prompt.

For clarity in notation, we will omit the subscript ki from here onwards. Our approach is inspired by Equivariant Diffusion Models (EDM) (Hoogeboom et al., 2022, Xu et al., 2023). These models utilize Equivariant Graph Neural Networks (EGNNs) (Satorras et al., 2021) as the base architecture to perform the denoising process in molecular generation. We model a geometric scene graph $G = (V, E)$, corresponding to scene \mathcal{S} , similar to a molecular graph, where the nodes $v \in V$ are endowed with 3D geometric coordinates x and node features h , i.e. $v = \langle x, h \rangle$. In our scenario, $x = \{x_1, x_2, \dots, x_{N_i}\}$ represents the 3D coordinates of the centroids of the objects in the 3D scene, while the rest of node features are denoted by $h = \{h_1, h_2, \dots, h_{N_i}\}$. Each feature vector h_i is composed of three components: $h_i = [c_i, f_i, b_i]$, where c_i is the class, f_i is the shape code, and $b_i = [s_i, \cos \theta_i, \sin \theta_i]$ represents the scaled and rotated bounding box of object o_i . Here, $[\cdot, \cdot]$ denotes concatenation of the inputs. For the edges E , we assume full connectivity between the nodes. The goal is to train an EDM p_ϕ conditioned on text prompts to learn the denoising process of the node features and 3D coordinates from the geometric scene graphs $G = (V, E)$. This will allow us to generate layout and appearance features of the objects to construct a novel 3D scene. We extend the original framework of EGNNs (Satorras et al., 2021), adapting it to handle higher-dimensional feature vectors and more complex conditioning, such as text control signals. We review the concepts of equivariance and invariance in Section 2.2, and provide a detailed overview of EGNNs in Section 2.3.

2.1 Text-aligned Shape Autoencoder

Similar to InstructScene (Lin and Mu, 2024), we obtain the object shape codes using OpenCLIP ViT-bigG/14 (Liu et al., 2023), which is a model that aligns CLIP (Radford et al., 2021) embeddings with 3D objects. The advantage of using OpenCLIP embeddings is that they ease the learning for the model to capture more nuanced and complex object descriptions, a capability that is more challenging to learn when the shape codes are extracted solely from point clouds, as in DiffuScene (Tang et al., 2024). However, instead of employing a vector-quantized variational autoencoder (VQ-VAE) (Van Den Oord et al., 2017) as is done in InstructScene, we use a standard VAE (Kingma, 2013) to reduce the high dimensionality of the OpenCLIP embeddings. Our method works well with embeddings in a continuous latent space, which allows us to use a lower dimensionality of 64 instead of 256 (as in InstructScene). The original dimensionality of OpenCLIP embeddings is 1280.

2.2 Equivariance and Invariance

Let $\rho(\mathfrak{g})$ be the linear representation related to the group element $\mathfrak{g} \in \mathfrak{G}$, where \mathfrak{G} is a group. A function f is \mathfrak{G} -equivariant to the action of a group \mathfrak{G} if: $f(\rho(\mathfrak{g})x) = \rho(\mathfrak{g})f(x)$, for all $\mathfrak{g} \in \mathfrak{G}$, i.e., the group action has the same effect if it is applied first to the input x or after to the output $f(x)$ (Bronstein et al., 2021). In the case of invariance, applying the group action on the input x has no effect on the output $f(x)$. In this work, we make use of the special Euclidean group $SE(3)$, that is generated by orthogonal rotation matrices $R \in SO(3)$ and translation vectors $t \in \mathbb{R}^3$. Let $x = (x_1, x_2, \dots, x_N) \in \mathbb{R}^{N \times 3}$ denote a point

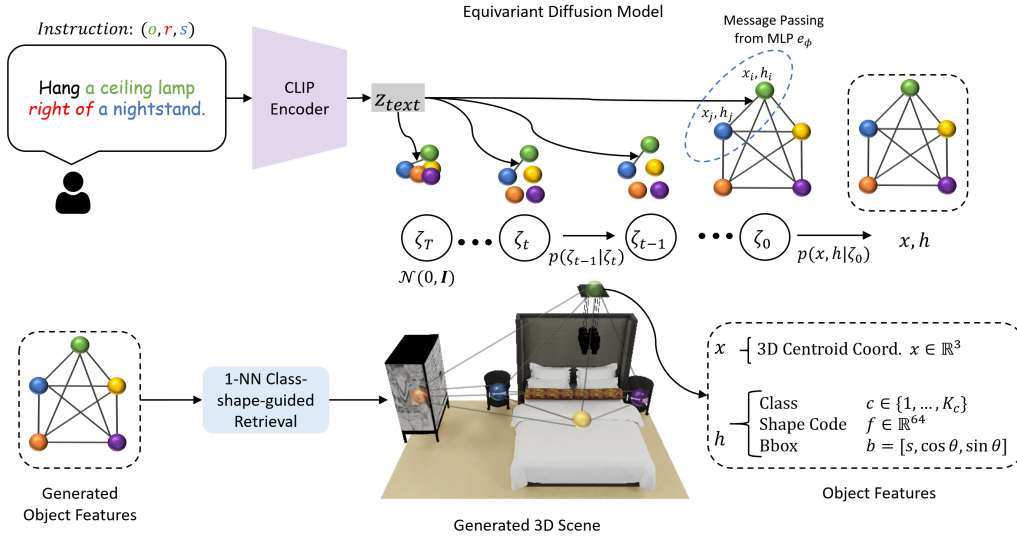


Figure 1. Overall pipeline of our method. First, we encode the text prompt with a CLIP encoder, and use this text embedding to condition the diffusion process through all message passing steps in the denoising EGNN. Once object features are sampled with the denoising process, we generate the scene by retrieving objects via 1-NN search and positioning them with 3D coordinates and bounding box parameters.




Figure 2. EGNN Architecture. Our EGNN architecture denoises the noisy node features $[x + \epsilon_x, h + \epsilon_h]$ through a three-phase pipeline composed of MLP encoders, EGCL layers and MLP decoders. Our novelty is highlighted by the dotted red boxes.

cloud with N points, and let $h = (h_1, h_2, \dots, h_N) \in \mathbb{R}^{N \times n_f}$ represent the concatenated features of these points, where n_f is the dimension of each feature vector h_i . Then, we want a vector-valued function $(z_x, z_h) = f(x, h)$ that is $SE(3)$ -invariant on h and $SE(3)$ -equivariant on x , i.e.:

$$(Rz_x + t, z_h) = f(Rx + t, h). \quad (1)$$

2.3 Equivariant Graph Neural Networks (EGNNs)

$E(n)$ -Equivariant Graph Neural Networks (Satorras et al., 2021) satisfy the constraint from Eq. 1. They have been successfully used for molecular generation (Hooeboom et al., 2022, Xu et al., 2023). In these works, all pairwise interactions between atoms are taken into account, and therefore all edges $e_{ij} \in E$, $i \neq j$, are included in the EGNN. An EGNN is composed of multiple equivariant graph convolutional layers (EGCLs), where the updates $[x^{l+1}, h^{l+1}] = \text{EGCL}[x^l, h^l]$ follow the formulation proposed in the original paper (Satorras et al., 2021):

$$m_{ij}^l = e_\phi(h_i^l, h_j^l, d_{ij}^2, a_{ij}), \quad h_i^{l+1} = h_\phi\left(h_i^l, \sum_{j \neq i} m_{ij}^l e_{ij}^l\right),$$

$$x_i^{l+1} = x_i^l + \sum_{j \neq i} \frac{x_j^l - x_i^l}{d_{ij} + 1} x_\phi(h_i^l, h_j^l, d_{ij}^2, a_{ij}), \quad (2)$$

where l indexes the EGCL layer, $d_{ij} = \|x_i^l - x_j^l\|_2$ is the Euclidean distance between the coordinates x_i^l and x_j^l at layer l ,

a_{ij} are optional edge attributes, and $\tilde{e}_{ij}^l = m_\phi(m_{ij}^l)$. The functions e_ϕ , h_ϕ , m_ϕ and x_ϕ are parameterized with learnable multi-layer perceptrons (MLPs). An EGNN is composed of a set of L EGCL layers, where the input of the first layer is $[x^0, h^0]$. We use this framework and extend it to our problem of 3D scene generation (see Section 2.4).

2.4 Equivariant Diffusion Models for 3D Scene Graphs

In this section, we introduce the 3D Scene Equivariant Graph Neural Network (3DS-EGNN), an adaptation of the EGNN framework by Satorras et al. (Satorras et al., 2021) specifically designed for 3D scene generation. Our 3DS-EGNN differs from the original EGNN in several key aspects. First, we employ ResNet and Transformer blocks to inject time-step embeddings, which proves more effective for high-dimensional latent vectors (see Section 3.4). Second, we incorporate text embeddings directly into each message-passing and coordinate-update step; in contrast, the original framework simply appended these to the node features. Finally, we model geometric interactions using the scaled vector difference between coordinates rather than standard Euclidean distance. The remainder of this section details these components.

Our approach takes as input both the 3D coordinates x of objects and their associated features h , denoted as x^0 and h^0 , respectively. The framework consists of three main phases (see Figure 2): encoding, EGCL layers, and decoding.

For the encoding phase, following prior works (Paschalidou et al., 2021, Tang et al., 2024, Lin and Mu, 2024), each feature type in h , specifically, the class c , shape code f , and bounding box b , is processed using a multi-layer perceptron (MLP). This produces latent representations: $z_c = E_\phi^c(c)$, $z_f = E_\phi^f(f)$, $z_b = E_\phi^b(b)$. These latent variables are then combined to form a joint representation: $z = E_\phi^h([z_c, z_f, z_b])$, which captures all the input features from h . To incorporate text prompt conditioning for the 3D scene generation, we employ a CLIP text encoder (Radford et al., 2021) to extract text embeddings: $z_{\text{text}} = E^{CLIP}(y)$. Additionally, we encode the diffusion process timestep t using

positional encodings (Vaswani et al., 2017), followed by an encoder: $z_{\text{time}} = E_{\phi}^t(\text{PE}(t))$. All these embeddings, namely z , z_{text} , and z_{time} , are then utilized within the EGCL layers.

For the second phase, we use L EGCL layers for the 3DS-EGNN. At each layer, indexed by l , we combine z^l and z_{time} using a ResNet (He et al., 2016) block RN_{ϕ}^l , followed by a Transformer (Vaswani et al., 2017) block T_{ϕ}^l . Then, we compute the scaled vector difference between the node coordinates as $\delta_{ij}^l = \frac{x_j^l - x_i^l}{\|x_j^l - x_i^l\|_2^{l+1}}$, and use the text embedding z_{text} as edge attributes. Afterwards, we update the latent variables z^l and the node coordinates x^l using MLPs $e_{\phi}^l, m_{\phi}^l, z_{\phi}^l, x_{\phi}^l$. After processing x^0 and z^0 through the L EGCL layers, we obtain x^L and z^L .

Finally, in the third phase, we decode z^L back into the original input features h^L using three distinct decoders: $\hat{c} = D_{\phi}^c(z^L)$, $\hat{f} = D_{\phi}^f(z^L)$, and $\hat{b} = D_{\phi}^b(z^L)$. The coordinates x^L remain unchanged. The complete 3DS-EGNN framework is presented in Algorithm 1.

Algorithm 1 3DS-EGNN Framework

Input: Node features $[x^0, h^0]$, time step t , text prompt y
 Encoders $E_{\phi}^h, E_{\phi}^c, E_{\phi}^f, E_{\phi}^b, E^{CLIP}, E_{\phi}^t$
 Decoders $D_{\phi}^c, D_{\phi}^f, D_{\phi}^b$
 EGCL layers $\{e_{\phi}^l, m_{\phi}^l, z_{\phi}^l, x_{\phi}^l, RN_{\phi}^l, T_{\phi}^l\}_{l=0}^{L-1}$
 Encode input features $h^0 = [c, f, b]$, time t and text y :
 $z^0 = E_{\phi}^h([E_{\phi}^c(c), E_{\phi}^f(f), E_{\phi}^b(b)])$
 $z_{\text{text}} = E^{CLIP}(y), z_{\text{time}} = E_{\phi}^t(\text{PE}(t))$
for l in $0, \dots, L-1$ **do**
 Compute $z^l = T_{\phi}^l(RN_{\phi}^l(z^l, z_{\text{time}}))$
 Compute $m_{ij}^l = e_{\phi}^l(z_i^l, z_j^l, \delta_{ij}^l, z_{\text{text}})$
 Update z_i^{l+1} and x_i^{l+1} :
 $z_i^{l+1} = z_{\phi}^l(z_i^l, \sum_{j \neq i} m_{ij}^l, m_{\phi}^l(m_{ij}^l))$
 $x_i^{l+1} = x_i^l + \sum_{j \neq i} \delta_{ij}^l x_{\phi}^l(z_i^l, z_j^l, \delta_{ij}^l, z_{\text{text}})$
end for
 Decode z^L back to $h^L = [\hat{c}, \hat{f}, \hat{b}]$:
 $\hat{c} = D_{\phi}^c(z^L), \hat{f} = D_{\phi}^f(z^L), \hat{b} = D_{\phi}^b(z^L)$
return $[x^L, h^L]$

The training and sampling procedures of the EDM remain largely as in the original EDM paper used for molecule generation (Hooeboom et al., 2022), with the key difference that we employ the 3DS-EGNN outlined above instead of the original EGNN framework. We present the respective algorithms in Algorithm 2 and Algorithm 3. We replace the variable z_t from the original EDM with ζ_t to avoid confusion with the latent variable resulting from the encoding phase in the 3DS-EGNN.

Algorithm 2 SG-EDM Training

Input: Data point $x, h = [c, f, b]$, text prompt y
 3DS-EGNN f_{ϕ}
 Sample $t \sim \mathcal{U}(1, \dots, T), \epsilon \sim \mathcal{N}(0, \mathbf{I})$
 Subtract the center of gravity in ϵ_x in $\epsilon = [\epsilon_x, \epsilon_h]$
 Compute $\zeta_t = \sqrt{\alpha_t}[x_0, h_0] + \sqrt{1 - \alpha_t}\epsilon$
 Minimize $\|\epsilon - f_{\phi}(\zeta_t, t, y)\|_2^2$

Algorithm 3 SG-EDM Sampling

Input: Text prompt y , trained 3DS-EGNN f_{ϕ}
 Sample $\zeta_T \sim \mathcal{N}(0, \mathbf{I})$
for t in $T, T-1, \dots, 1$ **do**
 Sample $\epsilon \in \mathcal{N}(0, \mathbf{I})$
 Subtract the center of gravity in ϵ_x in $\epsilon = [\epsilon_x, \epsilon_h]$
 Compute $\zeta_{t-1} = \frac{1}{\sqrt{\alpha_t}}(\zeta_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}}f_{\phi}(\zeta_t, t, y)) + \sigma_t\epsilon$
end for
return $[x, h] \sim p(x, h | \zeta_0)$

3. Experiments

3.1 Experimental Settings

Datasets. All experiments are carried out on the Scene Instruction 3D-FRONT dataset¹, an augmented version of 3D-FRONT (Fu et al., 2021). 3D-FRONT is a professionally designed synthetic dataset, comprising 6,813 houses with 14,629 rooms, where each room is furnished with high-quality 3D furniture objects extracted from the 3D-FUTURE dataset (Fu et al., 2021). However, the original dataset lacks textual descriptions of room layouts and object appearances. To address this, InstructScene (Lin and Mu, 2024) augments it by extracting perspective-dependent spatial relations using predefined geometric rules (for further details, see the supplementary material of InstructScene (Lin and Mu, 2024)). Then, 3D object captions are generated using BLIP (Li et al., 2022) and further refined with ChatGPT (Ouyang et al., 2022). The final scene instructions are created by randomly selecting relation triplets in the form (*object, relation, subject*), where both the *object* and *subject* are object names or object descriptions, and the *relation* is chosen from a predefined set of spatial relationships. For further details on dataset creation, we refer to Appendix A of InstructScene (Lin and Mu, 2024). To ensure consistency with previous methods, we adopt the same filtering process and train/test split proposed by ATISS (Paschalidou et al., 2021). The filtering process removes duplicates and problematic scenes, resulting in 4041 bedrooms, 900 dining rooms and 813 living rooms.

Baselines. We compare our method against three state-of-the-art task-specific methods for 3D indoor scene synthesis. (1) ATISS (Paschalidou et al., 2021), an autoregressive approach that employs a transformer encoder to predict object features. (2) DiffuScene (Tang et al., 2024), a diffusion-based method that utilizes a 1D U-Net to learn the denoising process of object features. (3) InstructScene (Lin and Mu, 2024), a two-stage approach that first employs a discrete diffusion method to generate a scene graph prior, and then uses it to condition a continuous diffusion model to generate the final layout. These three approaches differ in their controllability. ATISS does not support text conditioning. Although DiffuScene accepts text prompts, it struggles with object descriptions because its shape codes were trained solely on point clouds without explicit text alignment. InstructScene addresses this by extending the first two approaches with a pretrained CLIP text encoder for text conditioning and applying the same feature quantizer to the OpenCLIP ViT-bigG/14 embeddings. Building on this work, we compare our method against all three models (with extensions to (1) and (2) as described in InstructScene).

Since our focus in this paper is on efficient, offline scene synthesis, our experiments exclude VLM- and agent-based approaches.

¹ The augmented version of 3D-FRONT from InstructScene is available at https://huggingface.co/datasets/chenguolin/InstructScene_dataset.

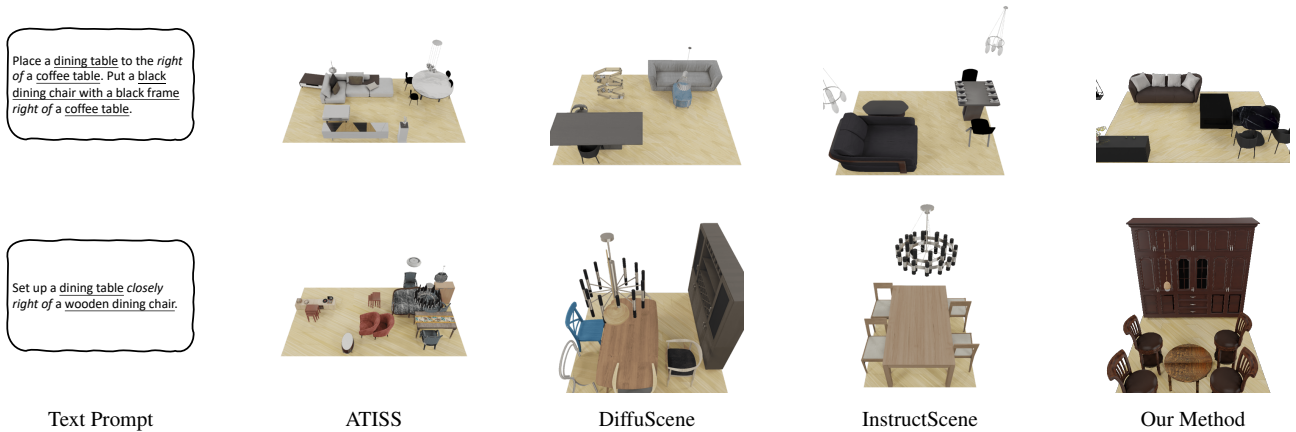


Figure 3. Visual comparison of our method and baseline approaches for text-guided scene generation on the living room (top row), and dining room (bottom row) datasets.

These methods typically rely on massive pre-trained backbones or cloud-based APIs, which fall outside the scope of lightweight, task-specific deployment. Instead, we prioritize comparisons with diffusion-based and graph-based baselines that share our objective of achieving high-fidelity 3D generation within a constrained computational footprint.

Evaluation Metrics. To assess how well synthesized scenes adhere to user instructions, we employ the "Instruction Recall" (*iRecall*) metric introduced by InstructScene (Lin and Mu, 2024). This metric calculates the proportion of predicted relation triplets ($o_{pred}, r_{pred}, s_{pred}$), derived via geometric rules during dataset construction, that match the ground-truth triplets (o_{gt}, r_{gt}, s_{gt}) extracted from the text prompts. Following previous works for 3D scene generation (Paschalidou et al., 2021, Tang et al., 2024, Lin and Mu, 2024), we evaluate generative quality by rendering 256×256 top-view projections of both real and synthesized scenes. From these, we compute standard image-based metrics: Fréchet Inception Distance (FID) (Heusel et al., 2017), FID-CLIP (Kynkäänniemi et al., 2022), Kernel Inception Distance (KID) (Bińkowski et al., 2018), and Scene Classification Accuracy (SCA). All scenes are rendered using the Blender engine (Community, 2018).

3.2 Text-guided 3D Scene Synthesis

In Tables 1 and 2, we present the experimental results for both InstructScene and our method. Note that there are discrepancies between our reproduced metrics and those reported in the InstructScene paper (Lin and Mu, 2024). Based on our review of their code, it appears that the reported metrics in the InstructScene paper were computed on the output of the first stage (i.e., the semantic scene graph) rather than the final layout. Since this metric can only be computed for their method, we also include these results in parentheses for comparison. Because the official extensions of ATISS and DiffuScene are not publicly available, we were unable to reproduce them and instead report their results directly from the InstructScene paper. Table 1 provides a quantitative comparison of our method and the baselines on the task of 3D scene synthesis guided by text instructions.

We first evaluate scene quality using FID, FID-CLIP, KID, and SCA. The results show that methods incorporating a scene graph into their pipeline, specifically InstructScene and our GeoScene-Graph, tend to produce higher-quality scenes. This improvement is particularly evident in more complex environments such

as living and dining rooms, which can contain up to 21 objects and more than 10 objects on average. In contrast, the bedroom dataset is comparatively simpler, containing at most 12 objects and typically around 5 per scene. Although InstructScene reports slightly stronger performance, our method achieves consistently competitive results across all datasets. For controllability, measured by the *iRecall* metric, the results are more varied. Our method outperforms all others on the dining room dataset and remains competitive on the living room dataset, but is surpassed on the bedroom dataset, where DiffuScene and ATISS achieve substantially higher scores. We hypothesize two contributing factors: (1) the bedroom dataset exhibits lower layout diversity, with only about five objects per scene (on average) and many layouts being structurally similar, making it easier for methods that do not incorporate scene graph structure to memorize these patterns. In such cases, the inherent scene graph becomes less informative, which may advantage these methods since they do not rely on the graph structure of the 3D scene; and (2) there may have been modifications to the evaluation procedure in the InstructScene implementation to enable comparisons using semantic scene graphs, which could have influenced the reported results.

Figure 3 presents qualitative comparisons between our method and the baselines on the living room, and dining room datasets.

3.3 Zero-shot Applications

In addition to the core task discussed in Section 3.2, we evaluate our method on four zero-shot applications proposed in InstructScene (Lin and Mu, 2024): stylization, re-arrangement, completion, and unconditional generation. Stylization modifies an input scene based on a text prompt specifying a style change. Re-arrangement repositions misplaced objects to match the text description. Completion augments a partial scene with additional objects that match the text prompt. Conversely, unconditional generation produces scenes with no text control or prior scene state. For stylization, we use the difference cosine similarity Δ , proposed in InstructScene, to evaluate how well the style matches the text prompt. For a scene S_k with N_k objects, it is defined as:

$$\Delta_k = \frac{1}{N_k} \sum_{i=0}^{N_k} S_C(e_{ki}^{obj}, e_{ki}^{style}) - S_C(e_{ki}^{obj}, e_{ki}^{class}), \quad (3)$$




Figure 4. Comparison visualizations of our method with DiffuScene(Tang et al., 2024) and InstructScene(Lin and Mu, 2024) for different zero-shot tasks. Due to space constraints, we omit ATISS (Paschalidou et al., 2021) from the first three visualizations.

Type of Room	Method	↑ iRecall (%)	↓ FID	↓ FID-CLIP	↓ KID ($\times 10^{-3}$)	SCA (%)
Bedroom	ATISS†	48.13	119.73	6.95	0.39	59.17
	DiffuScene†	56.43	123.09	7.13	0.39	60.49
	InstructScene	33.06 (66.53*)	119.32	6.66	0.30	60.49
	GeoSceneGraph	30.61	121.21	6.91	0.92	53.70
	ATISS†	29.50	117.67	6.08	17.60	69.38
Living room	DiffuScene†	31.15	122.20	6.10	16.49	72.92
	InstructScene	28.91 (51.02*)	109.21	5.06	3.96	58.33
	GeoSceneGraph	25.17	111.20	5.01	3.81	59.37
	ATISS†	37.58	137.10	8.49	23.60	67.61
Dining room	DiffuScene†	37.87	145.48	8.63	24.08	70.57
	InstructScene	37.92 (58.74*)	122.22	6.20	0.41	55.68
	GeoSceneGraph	40.89	123.04	6.41	0.49	56.81
	ATISS†					

Table 1. Quantitative comparisons of scene synthesis guided by text prompts. The † symbol indicates that these methods could not be reproduced, and their metrics were taken from the InstructScene (Lin and Mu, 2024) paper. Results marked with * in parentheses denote the *iRecall* computed on the semantic scene graph of the first stage of InstructScene. For the SCA metric, results close to 50% indicate better performance.

where S_C is the cosine similarity, e_{ki}^{obj} are CLIP visual features of object o_i , e_{ki}^{style} are CLIP text features from the stylization prompt, and e_{ki}^{class} are text features of the object class. The final score is the average across all M scenes: $\Delta = \frac{1}{M} \sum_{k=1}^M \Delta_k$. For the re-arrangement and completion tasks, we report both *iRecall* and FID metrics, while for unconditional generation, we report only the FID metric as it is not condition on anything. Although our method does not top the zero-shot tasks, it shows

competitive performance in controllability and generation quality. Notably, for unconditional generation, it ranks second, outperforming ATISS and DiffuScene by a wide margin. Figure 4 shows qualitative results on the zero-shot tasks. The visualizations indicate that our method performs strongly, especially in rearrangement and unconditional generation scenarios.

Type of Room	Method	Stylization	Re-arrangement		Completion		Uncond.
		$\uparrow \Delta \times 1e^{-3}$	\uparrow iRecall	\downarrow FID	\uparrow iRecall	\downarrow FID	\downarrow FID
Bedroom	ATISS [†]	3.44	61.22	107.67	64.90	89.77	134.51
	DiffuScene [†]	1.08	68.57	106.15	48.57	96.28	135.46
	InstructScene	1.59	38.37	113.96	65.71	84.41	124.97
	GeoSceneGraph	0.5	41.63	120.01	46.53	84.14	129.52
Living room	ATISS [†]	-3.57	31.97	117.97	43.20	106.48	129.23
	DiffuScene [†]	-2.69	41.50	115.30	19.73	95.94	129.75
	InstructScene	-0.95	37.76	109.46	48.64	81.39	111.34
	GeoSceneGraph	-1.94	39.46	133.67	20.75	81.09	116.74
Dining room	ATISS [†]	-1.11	36.06	134.54	57.99	122.44	147.52
	DiffuScene [†]	-2.98	46.84	133.73	32.34	115.08	150.81
	InstructScene	-1.75	46.84	123.95	56.13	96.13	126.97
	GeoSceneGraph	-2.22	47.96	129.02	35.69	83.19	134.90

Table 2. Quantitative comparisons on four zero-shot applications for scene synthesis. Symbol † indicates that these methods could not be reproduced, and their metrics were taken from the InstructScene (Lin and Mu, 2024) paper.

Method	\uparrow iRecall (%)	\downarrow FID	\downarrow FID-CLIP	\downarrow KID ($\times 1e^{-3}$)	SCA (%)
Our Method	25.17	111.20	5.01	3.81	59.37
Our Method (non-equivariant)	27.89	111.48	4.89	2.84	51.56
Concat. Conditioning (Hooeboom et al., 2022)	16.33	131.04	7.55	9.62	63.42
Cross Attention	27.21	115.88	5.40	4.48	59.37

Table 3. Results for our ablation studies on the living room dataset. Concat. Conditioning refers to concatenate the text and time step embeddings to the node features. Cross Attention refers to using a cross attention module with the node features and the text embeddings.

3.4 Ablation Studies

We focus our ablation studies on three distinct conditioning strategies for the Equivariant Diffusion Model (EDM). The first strategy concatenates the text and diffusion time-step embeddings directly with the node features prior to processing by the EGCL layers, following the original EDM implementation (Hooeboom et al., 2022). The second strategy employs standard cross-attention between node features and text embeddings. The third strategy, introduced in our method, first incorporates the time condition via ResNet and self-attention Transformer blocks, followed by the injection of text embeddings at each message-passing step of the graph neural network.

Table 3 reports the experimental results on the living room dataset. The results confirm that simply concatenating the temporal and textual context to the node features is insufficient for this task. Furthermore, even a more sophisticated cross-attention strategy underperforms compared to our proposed approach. By contrast, our multi-stage conditioning, which decouples the time-step processing and integrates text embeddings directly into the message-passing layers, achieves the best performance across all metrics.

We also evaluated the impact of the equivariance property by comparing our 3DS-EGNN against a standard graph neural network (GNN) baseline. This baseline utilizes the same geometric graph structure and message-passing scheme but lacks the strict equivariance constraints. Results in Table 3 show that the non-equivariant variant of our method outperforms the equivariant version across most metrics. While the 3DS-EGNN framework provides strict $E(3)$ equivariance, this global symmetry is partially broken in the context of indoor 3D scenes, which possess a canonical *up* orientation. Technically, indoor environments exhibit symmetry primarily in the XY -plane (i.e., rotation about the Z -axis), whereas the vertical axis is constrained by gravity. By enforcing full $SO(3)$ rotational equivariance, the model may be over-constrained, potentially preventing it from leveraging gravity-aligned spatial priors present in the data that

a non-equivariant or Z -axis-equivariant model can more effectively learn. We maintain the full equivariant formulation to ensure consistency with existing equivariant diffusion frameworks and to establish a rigorous baseline for such models in the context of 3D scene generation; however, we identify *gravity-aware* $SO(2)$ -equivariant architectures as a promising direction for future research.

4. Conclusion

In this work, we presented GeoSceneGraph, a geometric graph-based diffusion approach for synthesizing 3D indoor scenes from text prompts. Our method leverages the inherent geometric graph structure of the scene and does not rely on ground-truth semantic edges for object relationships, thereby enabling edge-level open-vocabulary training. Despite not using semantic edge annotations, our method achieves competitive results in both visual quality and text controllability, as demonstrated in our experiments. To handle high-dimensional conditioning for the equivariant graph neural network (EGNN), we introduce a strategy that first incorporates the time condition using a ResNet and self-attention Transformer blocks, and then integrates the text context within the message-passing steps of the EGNN. Our ablation studies corroborate these design choices.

References

- Berdoz, F., Lanzendörfer, L. A., Tuninga, N., Wattenhofer, R., 2026. Text-to-scene with large reasoning models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40number 4, 2435–2443.
- Bińkowski, M., Sutherland, D. J., Arbel, M., Gretton, A., 2018. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*.
- Bronstein, M. M., Bruna, J., Cohen, T., Velicković, P., 2021. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*.

- Bucher, M. J., Armeni, I., 2025. ReSpace: Text-Driven 3D Indoor Scene Synthesis and Editing with Preference Alignment. *arXiv preprint arXiv:2506.02459*.
- Çelen, A., Han, G., Schindler, K., Van Gool, L., Armeni, I., Obukhov, A., Wang, X., 2024. I-design: Personalized 3D interior designer. *European Conference on Computer Vision*, Springer, 217–234.
- Chen, W., Zhang, S., Zhang, Y., Zhou, T., Li, R., 2025. RoomPilot: Controllable Synthesis of Interactive Indoor Environments via Multimodal Semantic Parsing. *arXiv preprint arXiv:2512.11234*.
- Community, B. O., 2018. Blender - a 3D modelling and rendering package. Blender Foundation, Stichting Blender Foundation, Amsterdam.
- Feng, W., Zhu, W., Fu, T.-j., Jampani, V., Akula, A., He, X., Basu, S., Wang, X. E., Wang, W. Y., 2023. Layoutgpt: Compositional visual planning and generation with large language models. *Advances in Neural Information Processing Systems*, 36, 18225–18250.
- Fu, H., Cai, B., Gao, L., Zhang, L.-X., Wang, J., Li, C., Zeng, Q., Sun, C., Jia, R., Zhao, B. et al., 2021. 3d-front: 3d furnished rooms with layouts and semantics. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10933–10942.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S., 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33, 6840–6851.
- Hoogeboom, E., Satorras, V. G., Vignac, C., Welling, M., 2022. Equivariant diffusion for molecule generation in 3d. *International conference on machine learning*, PMLR, 8867–8887.
- Hu, S., Arroyo, D. M., Debats, S., Manhardt, F., Carlone, L., Tombari, F., 2024. Mixed Diffusion for 3D Indoor Scene Synthesis. *arXiv preprint arXiv:2405.21066*.
- Kingma, D. P., 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kynkäänniemi, T., Karras, T., Aittala, M., Aila, T., Lehtinen, J., 2022. The role of imagenet classes in fr¹chet inception distance. *arXiv preprint arXiv:2203.06026*.
- Li, J., Li, D., Xiong, C., Hoi, S., 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *International conference on machine learning*, PMLR, 12888–12900.
- Lin, C., Mu, Y., 2024. Instructscene: Instruction-driven 3d indoor scene synthesis with semantic graph prior. *arXiv preprint arXiv:2402.04717*.
- Ling, L., Lin, C.-H., Lin, T.-Y., Ding, Y., Zeng, Y., Sheng, Y., Ge, Y., Liu, M.-Y., Bera, A., Li, Z., 2025. Scenethesis: A language and vision agentic framework for 3d scene generation. *arXiv preprint arXiv:2505.02836*.
- Liu, M., Shi, R., Kuang, K., Zhu, Y., Li, X., Han, S., Cai, H., Porikli, F., Su, H., 2023. Openshape: Scaling up 3d shape representation towards open-world understanding. *Advances in neural information processing systems*, 36, 44860–44879.
- Maillard, L., Sereyjol-Garros, N., Durand, T., Ovsjanikov, M., 2025. DeBaRA: Denoising-Based 3D Room Arrangement Generation. *Advances in Neural Information Processing Systems*, 37, 109202–109232.
- Nie, Y., Dai, A., Han, X., Nießner, M., 2023. Learning 3d scene priors with 2d supervision. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 792–802.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A. et al., 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35, 27730–27744.
- Paschalidou, D., Kar, A., Shugrina, M., Kreis, K., Geiger, A., Fidler, S., 2021. Atiss: Autoregressive transformers for indoor scene synthesis. *Advances in Neural Information Processing Systems*, 34, 12013–12026.
- Pun, H. I. D., Tam, H. I. I., Wang, A. T., Huo, X., Chang, A. X., Savva, M., 2025. Hsm: Hierarchical scene motifs for multi-scale indoor scene generation. *arXiv preprint arXiv:2503.16848*.
- Purkait, P., Zach, C., Reid, I., 2020. Sg-vae: Scene grammar variational autoencoder to generate new indoor scenes. *European Conference on Computer Vision*, Springer, 155–171.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J. et al., 2021. Learning transferable visual models from natural language supervision. *International conference on machine learning*, PmlR, 8748–8763.
- Ran, X., Li, Y., Xu, L., Yu, M., Dai, B., 2025. Direct numerical layout generation for 3D indoor scene synthesis via spatial reasoning. *arXiv preprint arXiv:2506.05341*.
- Satorras, V. G., Hoogeboom, E., Welling, M., 2021. E (n) equivariant graph neural networks. *International conference on machine learning*, PMLR, 9323–9332.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S., 2015. Deep unsupervised learning using nonequilibrium thermodynamics. *International conference on machine learning*, PMLR, 2256–2265.
- Sun, F.-Y., Liu, W., Gu, S., Lim, D., Bhat, G., Tombari, F., Li, M., Haber, N., Wu, J., 2025. Layoutvln: Differentiable optimization of 3d layout via vision-language models. *Proceedings of the Computer Vision and Pattern Recognition Conference*, 29469–29478.

Tang, J., Nie, Y., Markhasin, L., Dai, A., Thies, J., Nießner, M., 2024. Diffuscene: Denoising diffusion models for generative indoor scene synthesis. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20507–20518.

Van Den Oord, A., Vinyals, O. et al., 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, X., Yeshwanth, C., Nießner, M., 2021. Sceneformer: Indoor scene generation with transformers. *2021 International Conference on 3D Vision (3DV)*, IEEE, 106–115.

Xia, H., Li, X., Li, Z., Ma, Q., Xu, J., Liu, M.-Y., Cui, Y., Lin, T.-Y., Ma, W.-C., Wang, S. et al., 2026. SAGE: Scalable Agentic 3D Scene Generation for Embodied AI. *arXiv preprint arXiv:2602.10116*.

Xu, M., Powers, A. S., Dror, R. O., Ermon, S., Leskovec, J., 2023. Geometric latent diffusion models for 3d molecule generation. *International Conference on Machine Learning*, PMLR, 38592–38610.

Yang, H., Zhang, Z., Yan, S., Huang, H., Ma, C., Zheng, Y., Bajaj, C., Huang, Q., 2021a. Scene synthesis via uncertainty-driven attribute synchronization. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5630–5640.

Yang, M.-J., Guo, Y.-X., Zhou, B., Tong, X., 2021b. Indoor scene generation from a collection of semantic-segmented depth images. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15203–15212.

Yang, Y., Jia, B., Zhang, S., Huang, S., 2025. SceneWeaver: All-in-One 3D Scene Synthesis with an Extensible and Self-Reflective Agent. *arXiv preprint arXiv:2509.20414*.

Yang, Y., Sun, F.-Y., Weihs, L., VanderBilt, E., Herrasti, A., Han, W., Wu, J., Haber, N., Krishna, R., Liu, L. et al., 2024. Holodeck: Language guided generation of 3d embodied ai environments. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16227–16237.

Zhai, G., Örnek, E. P., Chen, D. Z., Liao, R., Di, Y., Navab, N., Tombari, F., Busam, B., 2024. Echoscene: Indoor scene generation via information echo over scene graph diffusion. *European Conference on Computer Vision*, Springer, 167–184.

Zhai, G., Örnek, E. P., Wu, S.-C., Di, Y., Tombari, F., Navab, N., Busam, B., 2023. Commonsences: Generating common-sense 3d indoor scenes with scene graph diffusion. *arXiv preprint arXiv:2305.16283*.

Zhang, Z., Yang, Z., Ma, C., Luo, L., Huth, A., Vouga, E., Huang, Q., 2020. Deep generative modeling for scene synthesis via hybrid representations. *ACM Transactions on Graphics (TOG)*, 39(2), 1–21.