

S-INF: Towards Realistic Indoor Scene Synthesis via Scene Implicit Neural Field

Zixi Liang¹, Guowei Xu¹, Haifeng Wu², Ye Huang¹, Wen Li^{1,2*}, Lixin Duan^{1,3}

¹Shenzhen Institute for Advanced Study, University of Electronic Science and Technology of China

²School of Computer Science and Engineering, University of Electronic Science and Technology of China

³Sichuan Provincial Key Laboratory for Human Disease Gene Study and the Center for Medical Genetics, Department of Laboratory Medicine, Sichuan Academy of Medical Sciences and Sichuan Provincial People's Hospital, UESTC {zixiliang.acad, xuguowei368, haifengwu205, liwenbnu, lxduan}@gmail.com, edward.ye.huang@qq.com

Abstract

Learning-based methods have become increasingly popular in 3D indoor scene synthesis (ISS), showing superior performance over traditional optimization-based approaches. These learning-based methods typically model distributions on simple yet explicit scene representations using generative models. However, due to the oversimplified explicit representations that overlook detailed information and the lack of guidance from multimodal relationships within the scene, most learning-based methods struggle to generate indoor scenes with realistic object arrangements and styles. In this paper, we introduce a new method, Scene Implicit Neural Field (S-INF), for indoor scene synthesis, aiming to learn meaningful representations of multimodal relationships, to enhance the realism of indoor scene synthesis. S-INF assumes that the scene layout is often related to the object-detailed information. It disentangles the multimodal relationships into scene layout relationships and detailed object relationships, fusing them later through implicit neural fields (INFs). By learning specialized scene layout relationships and projecting them into S-INF, we achieve a realistic generation of scene layout. Additionally, S-INF captures dense and detailed object relationships through differentiable rendering, ensuring stylistic consistency across objects. Through extensive experiments on the benchmark 3D-FRONT dataset, we demonstrate that our method consistently achieves state-of-the-art performance under different types of ISS.

Code — <https://github.com/ZixiLiang/S-INF>

Introduction

Synthesizing realistic and diverse 3D indoor scenes is a long-standing problem in computer vision and graphics. This research topic has received widespread attention due to its significant cost reduction in fields such as virtual reality (Yu and Chen 2024; Li et al. 2022) and 3D design (Shi et al. 2023; Yu, Yeung, and Terzopoulos 2015). Specifically, ISS can virtually rearrange existing furniture, enabling more convenient virtual interior design. Despite recent progress on this topic (Gao et al. 2023; Yang et al. 2021a; Inoue et al. 2023; Zhao et al. 2023; Zhai et al. 2024b), the nature of the underlying multimodal distribution, including scene layout

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

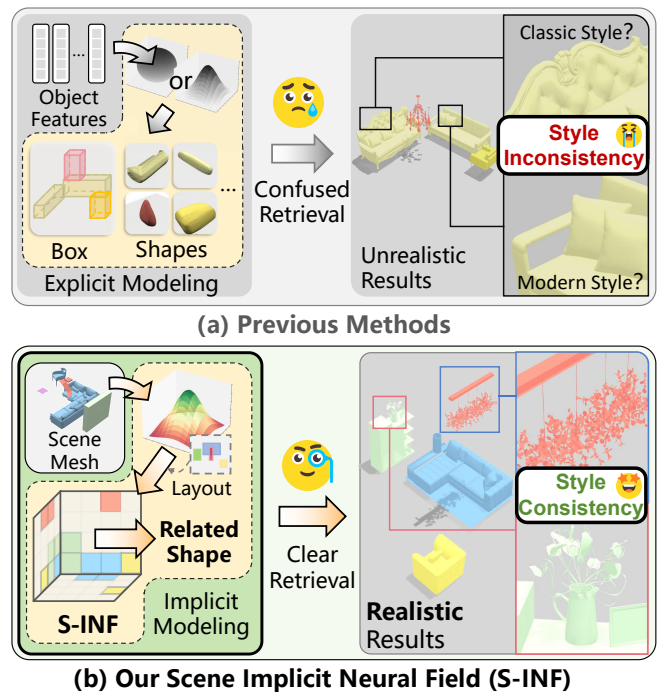


Figure 1: Unlike previous methods, we enhance the implicit modeling process from the S-INF with scene layout relationships and detailed object relationships, to achieve more layout realistic and style consistency generations.

and detailed object relationships, makes it still challenging to generate realistic and diverse 3D indoor scenes.

At the outset of this research, scene modeling and synthesis were typically formulated as an optimization problem (Zhang et al. 2019). Using scene priors like room design rules (Chang, Savva, and Manning 2014a; Chang et al. 2017) and in a human-centric manner (Qi et al. 2018a; Fisher et al. 2015a; Fu et al. 2017a), they first sample an initial scene and then refine its configuration through iterative optimization. However, defining precise rules demands significant expertise and may limit the representation of complex, diverse scenes. In recent years, learning-based methods have become popular, utilizing generative models

to learn scene distributions from data, such as Generative Adversarial Networks (GANs) (Yang et al. 2021c; Li, Li et al. 2023), Variational Autoencoders (VAEs) (Yang et al. 2021a; Purkait, Zach, and Reid 2020; Yang et al. 2021b), and diffusion models (Wu et al. 2024; Tang et al. 2023; Zhai et al. 2024a). These methods use generative models to model distributions on over-simplified and explicit-format scene representations (e.g., boxes and features). They usually map these explicit scene representations to latent distributions, constrained by prior distributions such as Gaussian or spherical distributions. When generating scenes, these methods decode sampled vectors from the prior distribution to obtain scene representations, followed by post-processing steps (Fisher and Hanrahan 2010; Shen et al. 2012; Chen et al. 2014; Zhang et al. 2023) to retrieve CAD models from the dataset and produce the final results. While this approach establishes a solid generative framework for ISS tasks, the overly simplified explicit representations overlook scene layout relationships and lack guidance of detailed object relationships within the scene, hindering the model from effectively learning multimodal scene relationships. Consequently, these methods often struggle to generate realistic indoor scenes, as they tend to focus on the major modes of the latent distribution while usually ignoring minor modes, a phenomenon known as mode collapse. Specifically, due to the limited expressiveness of the learned latent distribution, it is challenging for them to generate complex scenes that have realistic relationships or stylistic consistency.

In this paper, we aim to address the limitations mentioned above by proposing the S-INF to model multimodal relationships. S-INF assumes that the scene layout is often related to the object-detailed information. It disentangles the multimodal relationships into scene layout relationships and detailed object relationships. To generate indoor scenes, we first decode the scene layout relationships and detailed object relationships into the layout and the INF, then project the layout into the INF to obtain refined related shapes for retrieval. The disentangling construction offers several advantages: 1) we directly extract more advantageous multimodal information from the entire scene in a multiscale manner and map them into the S-INF, effectively modeling the multimodal relationship within the scene. 2) Unlike previous methods that sample directly from the prior distribution and decode it into explicit scene representations, our latent space also learns detailed object relationships, leading to a style-consistency generation. 3) During the learning process of the S-INF, we use differentiable rendering to capture dense and detailed object relationships, thereby ensuring realistic and stylistic consistency across objects (see Figure 1). Based on this, we retrieve CAD models from the dataset according to refined meshes sampled from the S-INF to obtain the final result. This multimodal relationship-based related shapes helps achieve diverse and realistic indoor scene synthesis. In summary, our main contributions are as follows:

- We uncover that the overly simplified explicit representations in current scene generation frameworks overlook detailed information and lack the necessary guidance for latent scene space modeling, making it difficult to effectively learn meaningful multimodal relationships, which

leads to challenges in generating realistic indoor scenes.

- We introduce a novel approach called Scene Implicit Neural Field (S-INF), which models wide and scene layout relationships as well as detailed object relationships, resulting in more realistic and style consistency ISS.
- Through extensive experiments conducted on the 3D-FRONT dataset, we demonstrate that our method attains state-of-the-art performance in ISS.

Related Work

3D Indoor Scene Synthesis

Early research treated the ISS task as an optimization problem, promoting the optimization process by introducing scene priors (Zhang et al. 2019). These priors typically included interior design guidelines, object frequency distributions, and scene arrangement examples. Guided by scene priors, new scenes can be generated from the formulation using various optimization methods, such as iterative approaches (Fisher et al. 2015b; Fu et al. 2017b), nonlinear optimization (Chang, Savva, and Manning 2014b; Fisher et al. 2012; Qi et al. 2018b), or manual interaction (Qi et al. 2018a; Fisher et al. 2015a; Fu et al. 2017a).

Recently, many learning-based methods have been proposed to synthesize complex scene compositions. These methods process the scenes to obtain simple, explicit scene representations (e.g., boxes and shapes), assuming that the scene representations obey a latent distribution. Generative models, such as feed-forward networks (Nie et al. 2022), recurrent networks (Li et al. 2019; Paschalidou et al. 2021; Zhang et al. 2020; Wang et al. 2019; Wang, Yeshwanth, and Nießner 2021; Inoue et al. 2023; Yi et al. 2023), GANs (Yang et al. 2021c; Li, Li et al. 2023), VAEs (Yang et al. 2021a; Purkait, Zach, and Reid 2020; Yang et al. 2021b), and diffusion models (Wu et al. 2024; Tang et al. 2023; Zhai et al. 2024a; Lin and Mu 2024), are then designed to learn this latent distribution, followed by retrieving CAD models from datasets based on the output of the generative models, such as shape-based retrieval or box-based retrieval. While the choice of generative models for ISS is crucial, the specific properties of the latent distribution that are advantageous for ISS remain unclear. Many works incorporated Gaussian or spherical distributions into latent representations and focused on developing robust generative models. However, directly modeling latent distributions based on explicit scene representations and lacking appropriate guidance on the latent distribution, these methods struggle to learn meaningful object relationship representations within scenes, thus achieving diverse and realistic 3D indoor scene synthesis. Therefore, the goal of this paper is to achieve meaningful object relationship representations in the latent space, enabling the model to autonomously learn the local relationships between different objects and thereby achieve accurate object relationships.

Implicit Neural Field

Our S-INF is inspired by INFs, which have shown promising results in representing high-fidelity geometry and appearance in 3D. Early work like DeepSDF (Park et al. 2019)

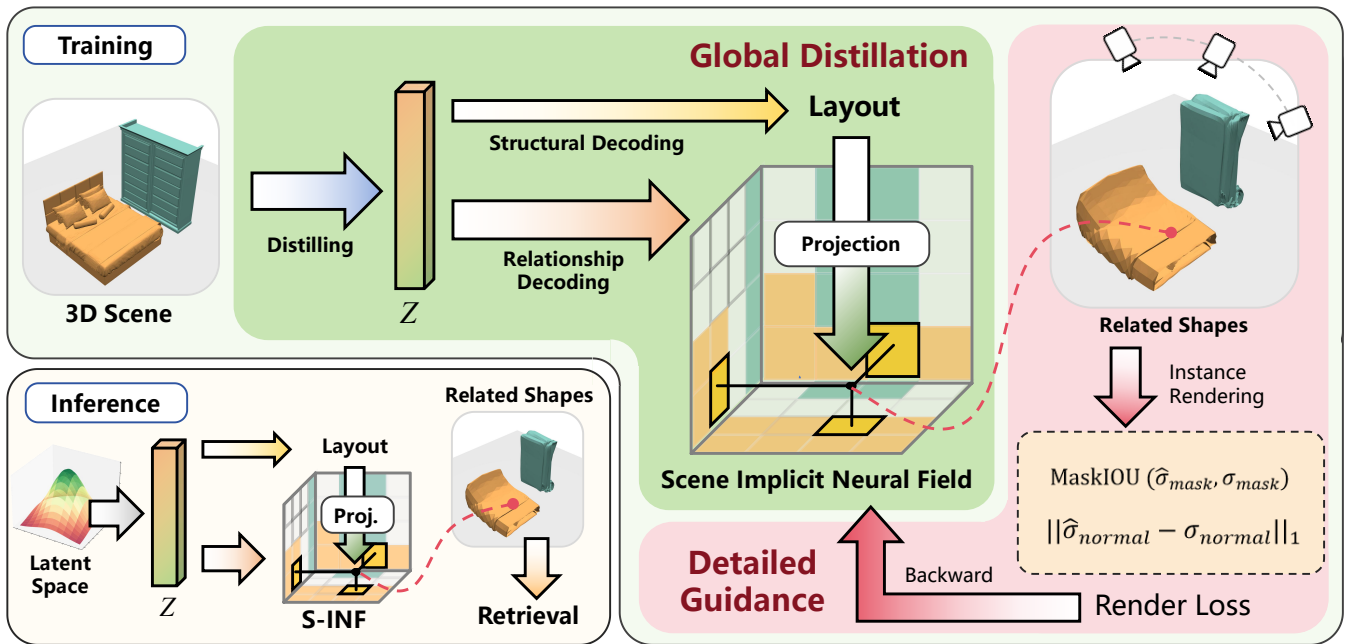


Figure 2: Our approach focuses on developing the S-INF to enable efficient capture of multimodal relationships and generate realistic and reliable 3D indoor scenes. We utilize the scene encoder to distill the realistic multimodal relationships into the S-INF. We also use differentiable rendering to enhance the S-INF style consistency information in detailed style relationships. Those optimized equip the S-INF with genuine multimodal relationship understanding capabilities, facilitating the generation of realistic, and style-invariant 3D indoor scenes.

represents the shape of a class of objects using signed distance functions, allowing for high-quality interpolation and completion from partial and noisy 3D input data. Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) leverage an MLP to model a coordinate-based radiance field, generating photo-realistic 2D renderings from novel views through volumetric rendering. DualOGNN (Wang, Liu, and Tong 2022) employs octrees to store the volumetric field of 3D shapes, effectively capturing shape details and demonstrating superior performance in various 3D shape and scene reconstruction tasks. NKField (Williams et al. 2022; Huang et al. 2023) uses data-dependent neural kernels to encode INFs, demonstrating strong generalization capabilities in 3D scene completion tasks. However, INF have not yet been fully explored in ISS. This paper constructs an INF in the latent space to achieve meaningful multimodal relationship representations. This enables the model to autonomously learn the detailed object relationships between different objects and thereby achieve accurate object relationships.

Methodology

The objective of ISS is to generate a sequence of object meshes, $\hat{X} = \{\hat{x}_j\}_{j=1}^m$, where m denotes the number of objects. The objects $\{\hat{x}_j\}_{j=1}^m$ should maintain object detailed consistency and a reasonable layout, and the modeled distribution $p(\hat{X})$ should closely match the diversity GT distribution $p(X)$. In this section, we first review the general formulation of the Learning-based ISS. Based on the gen-

eral framework, we then introduce our core S-INF equipped with global layout distillation and detailed style guidance. Finally, we summarize the training and inference processes.

General Formulation of Learning-based ISS

The core problem of learning-based ISS is to model the distribution $p(X)$. It is nontrivial to directly parameterize this distribution with a neural network. To this end, existing methods usually assume the scenes can be decoupled into scene-wise layout information and object-wise detailed information, modeling the distribution $p(X)$ through a two-step sequential process. The distribution $p(X)$ can be reformulated as $p(X) = p(\{D_{IS}(b_i)\}_{i=1}^n | b_i \in D_{RS}(z)\}_{i=1}^n)$, where the Relationship Decoder D_s captures scene relationships, the Instance Decoder D_{IS} learns object representations, and z is noise sampled from a prior distribution, such as Gaussian or Spherical distribution. The formulation of the above process can be summarized as follows:

$$\{b_i\}_{i=1}^n = D_{RS}(z), \quad z \sim p(z), \quad (1)$$

$$\{x_i\}_{i=1}^n = \{D_{IS}(b_i)\}_{i=1}^n. \quad (2)$$

The incorporation of the scene layout representation $\{b_i\}_{i=1}^n$ streamline the scene generation learning process. While generative models theoretically have the potential to incorporate a wide variety of scene attributes for z to learn, the discrete modeling of explicit object representation hinders information transfer between objects, making it difficult to accurately capture multimodal relationships, including layout-modal and object-modal. Specifically, the $\{b_i\}_{i=1}^n$ in Eq. (1)

are often with the over-simplified explicit format, and discretized, lacking expressiveness in the details of objects. The independent decoding of objects in Eq. (2) further leads to inconsistencies in object-detailed style due to the lack of constraints on detailed object relationships (Nie et al. 2023).

Scene Implicit Neural Field

To address the above issues, we construct the S-INF in the latent scene space, benefiting from its implicit modeling capabilities for potentially advantageous features, which is depicted in Fig. 2. The latent vector is decomposed into two factors: the scene layout representation l_b and the INF f_b . The scene representation, enriched with multimodal relationships, is assumed to be generated by sampling the scene layout within the INF. This process is formalized as follows:

$$f_b, l_b = D_{RS}(z), \quad z \sim p(z), \quad (3)$$

$$\{x_i\}_{i=1}^n = (D_{IS} \circ f_b)(l_b), \quad (4)$$

where operation $g \circ h : X \rightarrow Z$ denotes the function composition of $g : Y \rightarrow Z$ and $h : X \rightarrow Y$. In implementation, we draw from current works (Gao et al. 2022; Chan et al. 2022; Karras et al. 2020) and employ a 2D CNN to map the latent vector z to the INF of dimensions $N \times N \times (C \times 3)$. Here, N denotes the spatial resolution, and C represents the feature channels in the field, the output vector $f_b(p)$ of a point p can be interpolated as $f_b(p) = \sum_e \rho[\pi_e(p)]$, where π_e denotes the projection of the point p onto feature plane e , and ρ represents the bilinear interpolation operation. We employ a Transformer (Vaswani et al. 2017) decoder and generate the transformation of template spheres for deformation.

As shown in Eq. (3), S-INF disentangles multimodal relationships into f_b , which captures detailed object relationships, and l_b , which captures scene layout relationships. Through such explicit disentangling, S-INF demonstrates a superior capability in representing multimodal relationships compared to simply using Tri-Plane INF. In addition, by integrating global distillation and detailed guidance, S-INF can provide realistic multimodal relationships and style-consistent related shapes $\{x_i\}_{i=1}^n$ for retrieval. The following sections will introduce these two approaches.

Global Distillation. We utilize the global distillation to extra multimodal relationships globally from the Scene Encoder E_S , storing the extracted scene layout relationships and detailed object relationships into the layout l_b and INF f_b of S-INF, respectively. These are then combined into the related shapes $\{x_i\}_{i=1}^n$ through the process described in Eq. (4). Current methods typically rely on manually defining the over-simple explicit object representations, such as boxes, to represent the scene layout, then learning the scene layout relationships from it. However, we contend that these overly simplistic explicit representations lose detailed object relationships, leading to unrealistic layouts, such as overlaps, error arrangements, and misalignments (see Fig. 4). Additionally, these approaches make it difficult for the model to learn complex scene layout relationships, such as embedded relationships.

Fortunately, the implicit modeling capability of INFs allows for the accommodation of detailed object structural information and facilitates learning of complex scene layout

relationships (Peng et al. 2020; Chan et al. 2022). This enables us to combine INFs with explicit layouts to generate more realistic related shapes $\{x_i\}_{i=1}^n$ for retrieval. Specifically, we disentangle f_b and l_b , where f_b effectively extracts global-level structural features from the Scene Encoder E_S through distillation to guide the INF in learning complex detailed object representations. Next, we constrain l_b to multiple spherical layout attributes (such as position and scale) and use Layout Loss (Nie et al. 2023) to guide l_b in capturing scene layout relationships from the Scene Encoder E_S , as illustrated on the left of Fig. 3. Instead of using traditional 3D-CNNs, we implement the Scene Encoder E_S with sparse CNNs in distillation, which robustly encode and compress the scene through an efficient hierarchical structure. The scene is compressed into a latent vector z , then mapped into the INF f_b and layout l_b using the Relationship Decoder D_{RS} . To facilitate sampling, we constrain the latent distribution $p(Z)$ to approximate a Gaussian distribution.

Detailed Guidance. In addition to learning realistic multimodal relationships, ISS must further consider style consistency when capturing detailed object relationships. However, current methods, which rely on oversimplifying explicit representations, struggle to establish enough connections in detailed object relationships, leading to inconsistent object detailed styles.

To encourage style-consistent detailed object relationships, we utilize differentiable rendering to generate dense, detailed style guidance, enhancing style-detailed awareness of object components within the S-INF, as shown on the right side of Fig. 3. Specifically, we randomly initialize a camera pose on the spherical space. Then, we render each object within the scene individually (related shapes or ground truths) to obtain their normals $\hat{\sigma}_n$ or σ_n and masks $\hat{\sigma}_m$ or σ_m . Note that the normals $\hat{\sigma}_n$ or σ_n equip $\hat{\sigma}_{normal}$ or σ_{normal} in Fig. 2, and the mask $\hat{\sigma}_m = \hat{\sigma}_{mask}$ and $\sigma_m = \sigma_{mask}$. We render each object independently to reduce occlusion effects while preserving their positional information within the scene.

Through this rendering process, we implicitly acquire style information in detailed object relationships, while encouraging the S-INF to adaptively integrate detailed object relationships with scene layout relationships at all levels, creating realistic and style-consistent related shapes for retrieval. We employ the Render Loss L_{RD} as follows:

$$L_{RD} = \|\hat{\sigma}_n - \sigma_n\|_1 + \text{MaskIOU}(\hat{\sigma}_m, \sigma_m), \quad (5)$$

where the $\text{MaskIoU}(\cdot, \cdot)$ is the Mask IoU loss.

Training and Inference

Training. In our implementation, the Scene Encoder E_S compresses the scene X into a latent vector z . Then, the Relationship Decoder D_{RS} decomposes the latent vector z into the layout l_b and the INF f_b , where the layout l_b is implemented through the initial transformation of template spheres for all objects. To fine-tune the layout l_b into detailed related shapes $\{x_i\}_{i=1}^n$, the Instance Decoder D_{IS} projects the layout l_b into the INF f_b and obtains refined deformations at each point, applying a secondary transformation to

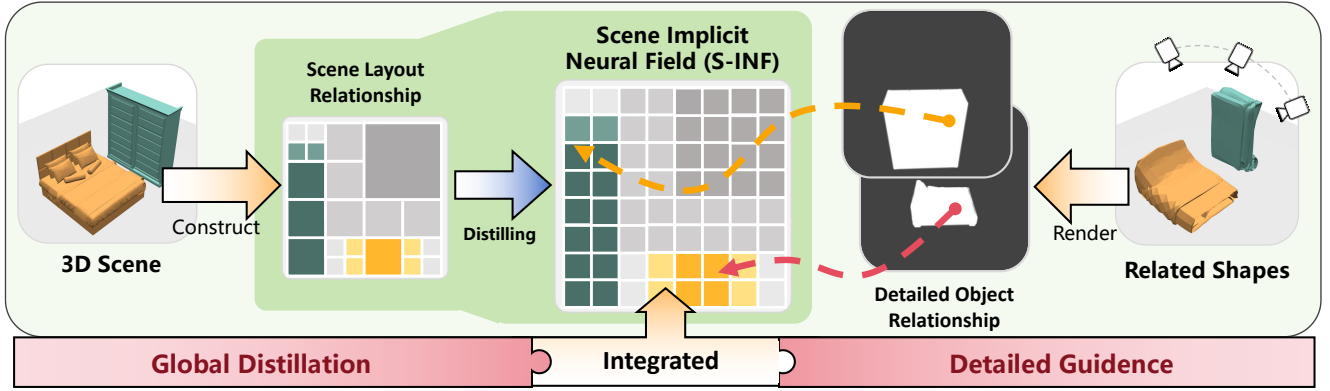


Figure 3: Integration between scene layout relationships and detailed object relationships. We leverage a specialized scene encoder to construct scene layout relationships (and little detailed object relationships) and distill multimodal relationships into S-INF. On the other hand, the rendered image provides dense and detailed style information, enriching the detailed object relationships in S-INF.

the initially transformed template spheres to generate the related shapes $\{x_i\}_{i=1}^n$. Following previous work (Nie et al. 2023), we then use Chamfer distance for shape retrieval to obtain the final scene \hat{X} . We train our network end-to-end with the following loss:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{KL}} + \mathcal{L}_{\text{RD}} + \mathcal{L}_{\text{LO}}, \quad (6)$$

where \mathcal{L}_{LO} represents the Layout Loss, and \mathcal{L}_{KL} is the Kullback-Leibler Loss, with $\alpha = 1 \times 10^{-4}$.

Inference. During the inference stage, we directly sample the latent vector z from the latent space $p(Z)$, then follow the training process and decode the latent vector z to the related shapes $\{x_i\}_{i=1}^n$, then finally use the Chamfer distance to retrieval shapes and replaced with the final object meshes.

Experiments

Experiment Setup

Datasets. To verify the performance in different scene types, we choose three scene types: Bedroom, Living Room, and Dining Room from the 3D-FRONT (Fu et al. 2021a).

Baselines. Apart from previous state-of-the-art methods such as Sync2Gen (SG), ATISS (AT), ScenePrior (SP), DifuScene (DS), EchoScene (ES), and InstructScene (IS), we also provide variants methods for a full comparison, designed based on the mentioned baselines. We design the ScenePrior-NN (NN) to utilize a MLP as D_{RS} , passing the code through three fully connected layers, projecting it onto a concatenation vector, and dividing it into equal lengths. On the other hand, the ScenePrior-Tr (Tr) employs a Transformer (Vaswani et al. 2017) as D_{RS} to parallel generate permutation-invariant layout representations. For a fair comparison, we optimize all the only-2D-supervision baselines using 3D IoU as layout loss, which has been admitted with a more stable performance in training (Nie et al. 2023).

Implementation. We provided all baselines with groundtruth meshes, categories, positions, render masks, and normal images in training. Given variations in scene sizes, we normalize all scene sizes to the range of $[-0.5, 0.5]$. Optimization is carried out using the AdamW (Loshchilov and Hutter 2017) optimizer with a batch size of 4 and a learning rate of 1×10^{-4} . All experiments are conducted on one RTX3090 GPU.

Evaluation Metrics. Similar to (Paschalidou et al. 2021), we utilize the Fréchet Inception Distance (FID), category KL divergence (KL), and scene classification accuracy (SCA) for quantitative comparisons. We normalize all the scenes first, then select a top-down projection to render 1024^2 images. Rather than generating category-color maps, we generate normal images for more detailed structure comparison. Also, to compare the diversity (Div) of different methods, we calculated the average \mathcal{L}_2 distance between all pairs of scene-rendered images, which is computed as follows: $\frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j=1}^N \|\hat{\sigma}_n - \sigma_n\|_2$. Note that lower KL and FID values indicate better performance, higher Div means higher diversity and the best SCA score is 0.5.

Experiment Results

Quantitative Comparisons. Table 1 provides a quantitative comparison in 3D-FRONT. Without retrieval, the compared methods rely on over-simply explicit representations to model scenes, discarding the multimodal relationships within the scenes. As a result, they fail to generate sufficiently complex scene representations, reflected in poor FID and SCA performance after retrieval. In contrast, by decoupling and learning the multimodal relationships within scenes, S-INF leverages scene layout relationships to generate more realistic and diverse scenes, while utilizing detailed object relationships to ensure the stylistic consistency of related shapes. This makes the retrieval results more diverse and has a better FID, SCA, and diversity.




Figure 4: Qualitative comparisons of bedroom with different baselines on ISS.

In addition, during retrieval, although existing methods like DiffuScene(Tang et al. 2023) and InstructScene(Lin and Mu 2024) achieve good FID and SCA performance, they overlook detailed object relationships in scenes and focus solely on modeling scene layout relationships. This results in a lack of diversity in their generated outcomes and leads to style-inconsistent results. From the performance in all metrics, our S-INF has a more outstanding performance in both realistic modeling and style-consistency.

Qualitative Comparisons. We present the qualitative comparative results in Fig. 4. All methods employ Chamfer Distance to gauge the similarity between x_i and CAD models in 3D-FUTURE (Fu et al. 2021b). Unlike previous

baselines that handle objects independently, we fully consider the multimodal relationships in the scene, enabling S-INF can generate realistic and style-consistent results.

We also visualize the S-INF generation capabilities in Fig. 5 and Fig. 6. The results show that the S-INF efficiently captures realistic multimodal relationships, highlighting that S-INF can generate tighter scenes without any overlap, misalignment, unaligned, or confused arrangement. Fig. 6 shows scenes with detailed object relationships. After retrieval, each scene presents objects unified by similar object details. This visualization highlights the ability of the S-INF to create realistic layout and uniform style-related detail scenes, ensuring realistic and consistent ISS.

Met.	Re?	Bedroom				Living Room				Dining Room			
		FID \downarrow	KL \downarrow	SCA \sim	Div \uparrow	FID \downarrow	KL \downarrow	SCA \sim	Div \uparrow	FID \downarrow	KL \downarrow	SCA \sim	Div \uparrow
NN	×	161.39	0.2521	0.8832	0.7284	188.25	0.1455	0.8565	0.7354	203.34	0.4252	0.9163	0.7346
Tr	×	156.42	0.2082	0.8742	0.7271	156.48	0.1339	0.8583	0.7414	162.82	0.2528	0.8821	0.7456
SP	×	164.70	0.2165	0.8596	0.7338	172.31	0.1357	0.8205	0.7390	168.85	0.2640	0.8812	0.7441
Ours	×	131.47	0.1985	0.8393	0.7379	155.47	0.1334	0.8034	0.7431	157.93	0.2342	0.8476	0.7465
SG	✓	119.68	0.2403	0.8128	0.7083	179.32	0.2190	0.8669	0.7260	99.97	0.2866	0.7062	0.7146
AT	✓	118.38	0.2069	0.7632	0.7161	174.13	0.3024	0.8871	0.7320	131.16	0.2397	0.7802	0.7063
NN	✓	114.63	0.2521	0.7105	0.7072	81.34	0.1455	0.6455	0.7188	84.46	0.4252	0.6437	0.7174
Tr	✓	117.02	0.2082	0.7865	0.7161	85.46	0.1339	0.6582	0.7106	97.55	0.2528	0.7020	0.7094
SP	✓	110.74	0.2165	0.7994	0.7223	84.36	0.1357	0.6708	0.7269	131.20	0.2640	0.8792	0.7036
ES	✓	107.27	-	0.6994	0.7259	109.30	-	0.6792	0.7275	119.30	-	0.7500	0.7208
DS	✓	84.40	0.2060	0.6319	0.7118	82.00	0.1771	0.6433	0.7319	86.28	0.2464	0.6533	0.7206
IS	✓	89.68	-	0.6202	0.7231	80.15	-	0.6435	0.7233	86.55	-	0.6706	0.7285
Ours	✓	79.52	0.1985	0.6128	0.7319	78.98	0.1334	0.6371	0.7322	79.95	0.2342	0.6236	0.7297

Table 1: We conducted a quantitative comparison of our method with state-of-the-art approaches on the 3D-FRONT dataset, where our method consistently demonstrated superior performance. Note that since EchoScene (ES) and InstructScene (IS) are class-conditional, their category-KL divergence is excluded, "Re" is retrieval and " \sim " in SCA indicates that 0.5 is optimal.




Figure 5: Realistic ISS from multimodal relationships.




Figure 6: Style consistency ISS from detailed object relationships.

Ablation Study. Table 2 compares the effects of different configurations of the relationship decoder D_{RS} , where the " $l_b + f_b$ " setting refers to the scene layout representation l_b (which learns scene layout relationships) and INF f_b (which learns detailed object relationships). Note that l_b and f_b imply that we ignored disentangling and modeling the INF and layout during training, but strived to maintain the parameter

Type	D_{RS}	FID \downarrow	KL \downarrow	SCA \sim	Div \uparrow
Bed	$l_b + f_b$	131.47	0.1985	0.8393	0.7309
	l_b	156.42	0.2082	0.8742	0.7271
	f_b	184.86	0.2408	0.8526	0.7137
Living	$l_b + f_b$	155.47	0.1334	0.8034	0.7431
	l_b	156.48	0.1339	0.8583	0.7414
	f_b	176.85	0.5819	0.8690	0.7193

Table 2: Evaluation of the configurations of D_{RS} . Note that " $l_b + f_b$ " denotes both l_b and f_b are both provided.

amount the same with the " $l_b + f_b$ ".

The results indicate that modeling the two relationships in a disentangled manner achieves the best performance; omitting either leads to a significant performance drop.

Conclusion

In this paper, we propose a novel 3D ISS method called Scene Implicit Neural Field (S-INF). S-INF effectively captures multimodal relationships within scenes, enhancing the realistic and style-consistency of ISS. It directly distillates more advantageous multimodal relationships from the entire scene, effectively capturing both scene layout relationships and detailed object relationships. For generating realistic related shapes for retrieval, S-INF achieves realistic multimodal relationship learning by disentangling and modeling scene layout relationships into the layout and detailed object relationships into the INF. For style-consistency, differentiable rendering is employed to enrich style information across objects. Extensive experiments on widely used benchmarks show that our method consistently achieves state-of-the-art performance in ISS tasks.

Acknowledgements

This research was partially funded by the National Natural Science Foundation of China (No. 82121003, and No. 62176047), the Shenzhen Fundamental Research Program (No. JCYJ20220530164812027).

References

- Chan, E. R.; Lin, C. Z.; Chan, M. A.; Nagano, K.; Pan, B.; De Mello, S.; Gallo, O.; Guibas, L. J.; Tremblay, J.; Khamis, S.; et al. 2022. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16123–16133.
- Chang, A.; Savva, M.; and Manning, C. D. 2014a. Learning spatial knowledge for text to 3D scene generation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2028–2038.
- Chang, A.; Savva, M.; and Manning, C. D. 2014b. Learning spatial knowledge for text to 3D scene generation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2028–2038.
- Chang, A. X.; Eric, M.; Savva, M.; and Manning, C. D. 2017. SceneSeer: 3D scene design with natural language. *arXiv preprint arXiv:1703.00050*.
- Chen, K.; Lai, Y.-K.; Wu, Y.-X.; Martin, R.; and Hu, S.-M. 2014. Automatic semantic modeling of indoor scenes from low-quality RGB-D data using contextual information. *ACM Transactions on Graphics*, 33(6).
- Fisher, M.; and Hanrahan, P. 2010. Context-based search for 3d models. In *ACM SIGGRAPH Asia 2010 papers*, 1–10.
- Fisher, M.; Ritchie, D.; Savva, M.; Funkhouser, T.; and Hanrahan, P. 2012. Example-based synthesis of 3D object arrangements. *ACM Transactions on Graphics (TOG)*, 31(6): 1–11.
- Fisher, M.; Savva, M.; Li, Y.; Hanrahan, P.; and Nießner, M. 2015a. Activity-centric scene synthesis for functional 3D scene modeling. *ACM Transactions on Graphics (TOG)*, 34(6): 1–13.
- Fisher, M.; Savva, M.; Li, Y.; Hanrahan, P.; and Nießner, M. 2015b. Activity-centric scene synthesis for functional 3D scene modeling. *ACM Transactions on Graphics (TOG)*, 34(6): 1–13.
- Fu, H.; Cai, B.; Gao, L.; Zhang, L.-X.; Wang, J.; Li, C.; Zeng, Q.; Sun, C.; Jia, R.; Zhao, B.; et al. 2021a. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10933–10942.
- Fu, H.; Jia, R.; Gao, L.; Gong, M.; Zhao, B.; Maybank, S.; and Tao, D. 2021b. 3d-future: 3d furniture shape with texture. *International Journal of Computer Vision*, 1–25.
- Fu, Q.; Chen, X.; Wang, X.; Wen, S.; Zhou, B.; and Fu, H. 2017a. Adaptive synthesis of indoor scenes via activity-associated object relation graphs. *ACM Transactions on Graphics (TOG)*, 36(6): 1–13.
- Fu, Q.; Chen, X.; Wang, X.; Wen, S.; Zhou, B.; and Fu, H. 2017b. Adaptive synthesis of indoor scenes via activity-associated object relation graphs. *ACM Transactions on Graphics (TOG)*, 36(6): 1–13.
- Gao, J.; Shen, T.; Wang, Z.; Chen, W.; Yin, K.; Li, D.; Litany, O.; Gojcic, Z.; and Fidler, S. 2022. Get3d: A generative model of high quality 3d textured shapes learned from images. *Advances In Neural Information Processing Systems*, 35: 31841–31854.
- Gao, L.; Sun, J.-M.; Mo, K.; Lai, Y.-K.; Guibas, L. J.; and Yang, J. 2023. SceneHGN: Hierarchical Graph Networks for 3D Indoor Scene Generation with Fine-Grained Geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Huang, J.; Gojcic, Z.; Atzmon, M.; Litany, O.; Fidler, S.; and Williams, F. 2023. Neural kernel surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4369–4379.
- Inoue, N.; Kikuchi, K.; Simo-Serra, E.; Otani, M.; and Yamaguchi, K. 2023. Layoutdm: Discrete diffusion model for controllable layout generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10167–10176.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8110–8119.
- Li, C.; Li, W.; Huang, H.; and Yu, L.-F. 2022. Interactive augmented reality storytelling guided by scene semantics. *ACM Transactions on Graphics (TOG)*, 41(4): 1–15.
- Li, M.; Patil, A. G.; Xu, K.; Chaudhuri, S.; Khan, O.; Shamir, A.; Tu, C.; Chen, B.; Cohen-Or, D.; and Zhang, H. 2019. Grains: Generative recursive autoencoders for indoor scenes. *ACM Transactions on Graphics (TOG)*, 38(2): 1–16.
- Li, S.; Li, H.; et al. 2023. Deep Generative Modeling Based on VAE-GAN for 3D Indoor Scene Synthesis. *International Journal of Computer Games Technology*, 2023.
- Lin, C.; and Mu, Y. 2024. InstructScene: Instruction-Driven 3D Indoor Scene Synthesis with Semantic Graph Prior. In *International Conference on Learning Representations (ICLR)*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Nie, Y.; Dai, A.; Han, X.; and Nießner, M. 2022. Pose2room: understanding 3d scenes from human activities. In *European Conference on Computer Vision*, 425–443. Springer.
- Nie, Y.; Dai, A.; Han, X.; and Nießner, M. 2023. Learning 3d scene priors with 2d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 792–802.

- Park, J. J.; Florence, P.; Straub, J.; Newcombe, R.; and Lovegrove, S. 2019. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 165–174.
- Paschalidou, D.; Kar, A.; Shugrina, M.; Kreis, K.; Geiger, A.; and Fidler, S. 2021. Atiss: Autoregressive transformers for indoor scene synthesis. *Advances in Neural Information Processing Systems*, 34: 12013–12026.
- Peng, S.; Niemeyer, M.; Mescheder, L.; Pollefeys, M.; and Geiger, A. 2020. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, 523–540. Springer.
- Purkait, P.; Zach, C.; and Reid, I. 2020. Sg-vae: Scene grammar variational autoencoder to generate new indoor scenes. In *European Conference on Computer Vision*, 155–171. Springer.
- Qi, S.; Zhu, Y.; Huang, S.; Jiang, C.; and Zhu, S.-C. 2018a. Human-centric indoor scene synthesis using stochastic grammar. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5899–5908.
- Qi, S.; Zhu, Y.; Huang, S.; Jiang, C.; and Zhu, S.-C. 2018b. Human-centric indoor scene synthesis using stochastic grammar. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5899–5908.
- Shen, C.-H.; Fu, H.; Chen, K.; and Hu, S.-M. 2012. Structure recovery by part assembly. *ACM Transactions on Graphics (TOG)*, 31(6): 1–11.
- Shi, Y.; Gao, T.; Jiao, X.; and Cao, N. 2023. Understanding design collaboration between designers and artificial intelligence: A systematic literature review. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW2): 1–35.
- Tang, J.; Nie, Y.; Markhasin, L.; Dai, A.; Thies, J.; and Nießner, M. 2023. Diffuscene: Scene graph denoising diffusion probabilistic model for generative indoor scene synthesis. *arXiv preprint arXiv:2303.14207*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, K.; Lin, Y.-A.; Weissmann, B.; Savva, M.; Chang, A. X.; and Ritchie, D. 2019. Planit: Planning and instantiating indoor scenes with relation graph and spatial prior networks. *ACM Transactions on Graphics (TOG)*, 38(4): 1–15.
- Wang, P.-S.; Liu, Y.; and Tong, X. 2022. Dual octree graph networks for learning adaptive volumetric shape representations. *ACM Transactions on Graphics (TOG)*, 41(4): 1–15.
- Wang, X.; Yeshwanth, C.; and Nießner, M. 2021. Sceneformer: Indoor scene generation with transformers. In *2021 International Conference on 3D Vision (3DV)*, 106–115. IEEE.
- Williams, F.; Gojcic, Z.; Khamis, S.; Zorin, D.; Bruna, J.; Fidler, S.; and Litany, O. 2022. Neural fields as learnable kernels for 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18500–18510.
- Wu, Z.; Feng, M.; Wang, Y.; Xie, H.; Dong, W.; Miao, B.; and Mian, A. 2024. External Knowledge Enhanced 3D Scene Generation from Sketch. *arXiv preprint arXiv:2403.14121*.
- Yang, H.; Zhang, Z.; Yan, S.; Huang, H.; Ma, C.; Zheng, Y.; Bajaj, C.; and Huang, Q. 2021a. Scene synthesis via uncertainty-driven attribute synchronization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5630–5640.
- Yang, H.; Zhang, Z.; Yan, S.; Huang, H.; Ma, C.; Zheng, Y.; Bajaj, C.; and Huang, Q. 2021b. Scene synthesis via uncertainty-driven attribute synchronization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5630–5640.
- Yang, M.-J.; Guo, Y.-X.; Zhou, B.; and Tong, X. 2021c. Indoor scene generation from a collection of semantic-segmented depth images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15203–15212.
- Yi, H.; Huang, C.-H. P.; Tripathi, S.; Hering, L.; Thies, J.; and Black, M. J. 2023. MIME: Human-aware 3D scene generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12965–12976.
- Yu, H.; and Chen, C. 2024. Automatic generation of civil engineering structure model based on network virtual reality. *International Journal of Global Energy Issues*, 46(1-2): 69–89.
- Yu, L.-F.; Yeung, S.-K.; and Terzopoulos, D. 2015. The clutterpalette: An interactive tool for detailing indoor scenes. *IEEE transactions on visualization and computer graphics*, 22(2): 1138–1148.
- Zhai, G.; Örnek, E. P.; Chen, D. Z.; Liao, R.; Di, Y.; Navab, N.; Tombari, F.; and Busam, B. 2024a. EchoScene: Indoor Scene Generation via Information Echo over Scene Graph Diffusion. *arXiv preprint arXiv:2405.00915*.
- Zhai, G.; Örnek, E. P.; Wu, S.-C.; Di, Y.; Tombari, F.; Navab, N.; and Busam, B. 2024b. Commonsences: Generating commonsense 3d indoor scenes with scene graphs. *Advances in Neural Information Processing Systems*, 36.
- Zhang, B.; Yuan, J.; Shi, B.; Chen, T.; Li, Y.; and Qiao, Y. 2023. Uni3d: A unified baseline for multi-dataset 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9253–9262.
- Zhang, S.-H.; Zhang, S.-K.; Liang, Y.; and Hall, P. 2019. A survey of 3d indoor scene synthesis. *Journal of Computer Science and Technology*, 34: 594–608.
- Zhang, Z.; Yang, Z.; Ma, C.; Luo, L.; Huth, A.; Vouga, E.; and Huang, Q. 2020. Deep generative modeling for scene synthesis via hybrid representations. *ACM Transactions on Graphics (TOG)*, 39(2): 1–21.
- Zhao, Y.; Zhao, Z.; Li, J.; Dong, S.; and Gao, S. 2023. RoomDesigner: Encoding Anchor-latents for Style-consistent and Shape-compatible Indoor Scene Generation. *arXiv preprint arXiv:2310.10027*.