

# MOVIS: Enhancing Multi-Object Novel View Synthesis for Indoor Scenes

Ruijie Lu<sup>1,2\*</sup> Yixin Chen<sup>2,\*†</sup> Junfeng Ni<sup>2,3</sup> Baoxiong Jia<sup>2</sup>  
 Yu Liu<sup>2,3</sup> Diwen Wan<sup>1</sup> Gang Zeng<sup>1†</sup> Siyuan Huang<sup>2†</sup>

\* Equal contribution † Corresponding author

<sup>1</sup> State Key Laboratory of General Artificial Intelligence, Peking University

<sup>2</sup> State Key Laboratory of General Artificial Intelligence, BIGAI <sup>3</sup> Tsinghua University




Figure 1. **Novel view synthesis and cross-view image matching.** The first row shows that MOVIS generalizes to different datasets on novel view synthesis (NVS). We also show visualizations of cross-view consistency compared with Zero-1-to-3 [46] and ground truth by applying image-matching. MOVIS can match a significantly greater number of points, closely aligned with the ground truth.

## Abstract

Repurposing pre-trained diffusion models has been proven to be effective for NVS. However, these methods are mostly limited to a single object; directly applying such methods to compositional multi-object scenarios yields inferior results, especially incorrect object placement and inconsistent shape and appearance under novel views. How to enhance and systematically evaluate the cross-view consistency of such models remains under-explored. To address this issue, we propose MOVIS to enhance the structural awareness of

the view-conditioned diffusion model for multi-object NVS in terms of model inputs, auxiliary tasks, and training strategy. First, we inject structure-aware features, including depth and object mask, into the denoising U-Net to enhance the model’s comprehension of object instances and their spatial relationships. Second, we introduce an auxiliary task requiring the model to simultaneously predict novel view object masks, further improving the model’s capability in differentiating and placing objects. Finally, we conduct an in-depth analysis of the diffusion sampling process and carefully devise a structure-guided timestep sampling

*scheduler during training, which balances the learning of global object placement and fine-grained detail recovery. To systematically evaluate the plausibility of synthesized images, we propose to assess cross-view consistency and novel view object placement alongside existing image-level NVS metrics. Extensive experiments on challenging synthetic and realistic datasets demonstrate that our method exhibits strong generalization capabilities and produces consistent novel view synthesis, highlighting its potential to guide future 3D-aware multi-object NVS tasks. Our project page is available at <https://jason-aplp.github.io/MOVIS/>.*

## 1. Introduction

Novel view synthesis (NVS) from a single image is imperative for various applications, including scene understanding [17, 21, 22, 27, 43] and interaction [10, 28, 29], interior designs [3, 20, 24, 81, 92], robotics [37–40, 49, 51, 65], etc. This is highly challenging as it requires understanding complex spatial structures from a single viewpoint while being able to extrapolate consistent and plausible content for unobserved areas. The substantial demands for comprehensive knowledge of the 3D world render it a difficult task, even for humans with rich priors of the 3D environments.

Recently, significant progress has been made in single-object image-to-3D generation [44, 45, 50, 68, 69, 73, 74, 87] empowered by powerful 2D diffusion models [19, 66]. Among them, one prominent line of research [5, 23, 41, 42, 44, 47, 63, 72, 84, 97] has achieved compelling results by building on insights from Zero-1-to-3 [46]: repurposing a pre-trained diffusion model as a novel view synthesizer by fine-tuning on large 3D object datasets can provide promising 3D-aware prior for image-to-3D tasks. However, these methods are mostly restricted to the single-object level. It remains unclear if this paradigm can be effectively extended to the multi-object level to facilitate more complex tasks like reconstructing an indoor scene. In Fig. 1, we visualize cross-view matching results of directly applying the novel view synthesizers [46] in multi-object scenarios, which showcases weak consistency with input views. Specifically, we believe that the lack of structural awareness is the primary reason for the disappearance, distortion, incorrect position and orientation of objects under novel views. While several works [67, 76] have explored training on mixed real-world scene datasets, the complexity introduced by multiple objects, such as spatial placement, per-instance geometry and appearance, and occlusion relationship, makes incorporating such awareness non-trivial.

Inspired by the discussion above, our paper seeks to address the question: *How to enhance the structural awareness of current diffusion-based novel view synthesizers?* We begin by identifying key challenges in extending single-object methods for multi-object NVS tasks. A multi-object image possesses more complicated structural information

than a single-object one. The model must first grasp hierarchical structures within, which includes both high-level object placement, e.g., position and orientation, and low-level ones like per-object geometry and appearance. High-level structural information significantly reduces the ambiguity in object composition while low-level details are essential for capturing the characteristics of each object instance. Subsequently, the model needs to retain this hierarchical information captured from the input view while synthesizing novel-view images to ensure cross-view structural consistency. These capabilities are less critical in single-object NVS tasks due to the reduced ambiguity in one-to-one mapping but are crucial for effective multi-object NVS models.

Building on these insights, our technical designs are threefold. We first propose injecting structure-aware features, i.e., depth and object mask, from the input view as additional inputs to provide information on both high-level global placement and fine-grained local details. Secondly, we utilize the prediction of novel view object masks as an auxiliary task during training for the model to differentiate object instances, laying a solid foundation for fine-grained geometry and appearance recovery. Finally, through an in-depth analysis of the model’s inference process, we highlight the importance of revising the noise timestep sampling schedule, which influences the learning focus in the training process. To be specific, larger timesteps emphasize global placement learning, while smaller timesteps focus on local fine-grained object geometry and appearance recovery. To endow the view-conditioned diffusion model with both capabilities, we propose a structure-guided timestep sampling scheduler that prioritizes larger timesteps in the initial stage, gradually decreasing over time to balance these two conflicting inductive biases. This design is fundamental to our proposed model’s effectiveness in addressing the complexity of multi-object level NVS tasks.

To systematically assess the plausibility of synthesized novel view images, we additionally evaluate novel-view object mask and cross-view structural consistency apart from the existing NVS metrics. Specifically, we employ image-matching techniques [35, 82] to compare the input-view image with both the ground-truth and synthesized novel-view images. Cross-view structural consistency evaluates how closely the matching results align, providing a measure of the accuracy in recovering object placement, shape, and appearance. On the other hand, the object mask, as measured by Intersection over Union (IoU), assesses the precision of object placement. Extensive experiments demonstrate that our method excels at multi-object level NVS in indoor scenes, achieving consistent object placement, shape, and appearance. Notably, it exhibits strong generalization capabilities for generating novel views on unseen datasets, including both synthetic ones 3D-FRONT [15], Room-Texture [54] and Objaverse [12], as well as the real-

world SUNRGB-D [71].

In summary, our main contributions are:

1. We introduce structure-aware features as model inputs and incorporate novel view mask prediction as an auxiliary task during training. This enhances the model’s understanding of hierarchical structures in multi-object scenarios, leading to improved NVS performance.
2. We present a novel noise timestep sampling scheduler to balance the learning of global object placement and fine-grained detail recovery, which is critical for addressing the increased complexity in multi-object scenarios.
3. We introduce additional metrics to systematically evaluate the novel view structural consistency. Through extensive experiments, our model demonstrates superiority in consistent object placement, geometry, and appearance recovery, showcasing strong generalization capability to unseen datasets.

## 2. Related Work

**Single object NVS with generative models** Synthesizing novel view images for single objects given a single-view image is an extremely ill-posed problem that requires strong priors. With great advances achieved in diffusion models [19, 66], research efforts [55, 73, 88] seek to distill priors [26, 62] learned from Text-to-Image (T2I) diffusion models via image captioning like [36]. However, this presents a huge gap between the image and semantics due to the ambiguity of the text, hindering the 3D consistency of these methods. On the other hand, view-conditioned diffusion models like Zero-1-to-3 [46] explore an Image-to-Image (I2I) generation paradigm that “teaches” the diffusion model to control viewpoints to synthesize plausible images under novel views, providing a more consistent 3D-aware prior. Subsequent work focuses on accelerating the generation speed [44, 72], enhancing the view consistency [5, 23, 42, 47, 84], or accelerating the training process [30]. However, all these methods deal with single and complete object novel view synthesis tasks since they usually fine-tune their model on Objaverse [11, 12], an extensive single-object level dataset, contrary to real images which normally consist of multiple or incomplete objects. The lack of specific model designs for compositional scenes also leads to significant inconsistencies when directly applying them to the multi-object scenarios, as shown in Fig. 1.

### **Multi-object 3D reconstruction with single object priors**

Following the advance in 3D-aware single object generative prior [46, 68], a line of research work [6, 7, 14, 57, 58, 81] focuses on extending their application to compositional multi-object scenarios. The core idea is to decompose object compositions into individual objects, thereby fully leveraging the powerful generative priors of single-object

models. They first break down a multi-object composition into several components via segmentation models like SAM [33], and then complete every single object with amodal [52, 61, 89, 95] or inpainting [53, 66, 75] techniques. The object instances are lifted to 3D via image-to-3D models [46, 72, 79, 86] and finally composited into a whole utilizing spatial-aware optimization, 3D bounding box detection [1, 59] or carefully estimating the metric depth [31, 91]. However, this divide-and-conquer paradigm is limited by the user-specified spatial relations from language prompts [7] and relies heavily on the cascaded modules of detection [1, 31, 33], completion [53, 66] and 3D-aware object-level novel view synthesis (NVS) [44, 79] to provide priors for reconstruction. Unlike any of the above, our method aims to build an end-to-end image-conditioned novel view synthesis model that can directly cope with the increased complexity in multi-object compositions, especially in the multiple-object setting.

**Scene-level NVS with sparse view input** Early efforts [25, 48, 80, 94] attempted to directly perform scene-level NVS tasks by extracting image features from input-view images and inferring the underlying 3D representation [56]. With the development of Gaussian Splatting [32], recent works [2, 8] attempt to switch the underlying representation to Gaussian Splatting for efficiency. However, they mainly deal with synthesizing views near input ones with limited generative capabilities to the unseen region. Inspired by the great success of diffusion models [66] and the object-level 3D-aware novel-view synthesizer [46, 79], several recent works have also attempted to perform scene-level NVS tasks by directly conditioning the generative models on a single-view scene image or a monocular dynamic scene video [78]. ZeroNVS [67] proposes to train a view-conditioned diffusion model on a mixture of real-world datasets, MegaScenes [76] further scales up the training dataset with Internet-level data pairs for stronger generalization capabilities. However, all these works mainly deal with small view-change and simple scenarios in terms of object number, with few adaptations to tackle the multi-object complexity. In this work, we systematically examine the cross-view consistency of NVS by proposing new metrics, and explore the critical designs required to enhance the structural consistency of the view-conditioned diffusion models in the multi-object scenarios.

## 3. Method

In this section, we address the challenge of enhancing the structural awareness of diffusion-based novel view synthesizers for better cross-view consistency in multi-object scenarios. We begin with a brief introduction to diffusion models and view-conditioned diffusion models (Sec. 3.1). Next, we detail the key architectural designs of MOVIS,




Figure 2. **Overview of MOVIS.** Our model performs NVS from the input image and relative camera change. We introduce structure-aware features as additional inputs and employ mask prediction as an auxiliary task (Sec. 3.2). The model is trained with a structure-guided timestep sampling scheduler (Sec. 3.3) to balance the learning of global object placement and local detail recovery.

including how we incorporate structural-aware features as input to improve the model’s understanding of hierarchical structure information (Sec. 3.2) and how we introduce novel view mask prediction as an auxiliary task, instructing the model to differentiate the object instances with correct object placement (Sec. 3.2). Finally, we provide an in-depth analysis of the inference process and adopt a structure-guided timestep sampling scheduler (Sec. 3.3) to balance the learning of global object placement and local fine-grained object geometry and appearance recovery. We provide an overview of our model in Fig. 2.

### 3.1. Diffusion models as novel view synthesizers

Diffusion models have been recently repurposed as a novel view synthesizer. By training on posed image pairs  $\{(\mathbf{x}_0, \hat{\mathbf{x}}_0)\}$  where  $\hat{\mathbf{x}}_0 \in \mathbb{R}^{H \times W \times 3}$  denotes the input view image and  $\mathbf{x}_0 \in \mathbb{R}^{H \times W \times 3}$  denotes the target view, view-conditioned diffusion models [46, 83] use the input image  $\hat{\mathbf{x}}_0$  and camera pose transformation as conditions to predict the target view image  $\mathbf{x}_0$ . Concretely, the learning objective of view-conditioned diffusion models is:

$$\mathbb{E}[\|\epsilon_\theta(\alpha_t \mathbf{x}_0 + \sigma_t \epsilon, t, C(\hat{\mathbf{x}}_0, R, T)) - \epsilon\|_2^2], \quad (1)$$

where  $R, T$  represent the relative camera pose transformation between the target view  $\mathbf{x}_0$  and the input view  $\hat{\mathbf{x}}_0$ .  $C(\hat{\mathbf{x}}_0, R, T)$  is the view-conditioned feature, combining the

relative camera pose transformation with encoded image features to form a new ‘pose-aware’ feature map, taking the place of the origin CLIP [64] feature embedding. Moreover, input view image  $\mathbf{x}_0$  will be concatenated with the noisy image as the input of the denoising U-Net. As discussed in Sec. 2, single-image-based NVS is extremely challenging, current methods inherit natural image priors from large-scale pre-training [66] and fine-tune diffusion models on large-scale 3D object datasets like Objaverse [12] to learn the transformation between objects in the input and novel views given the relative camera pose. Despite their ability to generalize to in-the-wild objects, these view-conditioned diffusion models struggle with multi-object scenarios like multi-furniture indoor scenes due to the scarcity of similar data and increased complexity arising from intricate object compositions. Our method builds on the insight of repurposing the diffusion model as a novel view synthesizer while emphasizing the inherent properties of multi-object scenarios in both model design and training strategy to facilitate multi-object NVS.

### 3.2. MOVIS

Our proposed method extends view-conditioned diffusion models to multi-object level, as illustrated in Fig. 2. The model leverages a pre-trained Stable Diffusion [66] and concatenates the 2D structural information from the input

view with a noisy target image as input. Additionally, it integrates a pre-trained image encoder [60] to capture semantic information, which is injected into the network through cross-attention alongside the relative camera pose. Moreover, it predicts novel view mask simultaneously as an auxiliary task to aid global object placement learning.

**Structure-Aware Feature Amalgamation** To synthesize plausible images under novel viewpoints, the model must first grasp the compositional structural information from the input view, laying a solid foundation for generation. To address the innate complexity in multi-object scenarios due to the intricate object relationship, we propose to leverage structure-aware features to facilitate model’s comprehension. Specifically, we use depth maps and object masks as proxies for image-level structural information. Object masks provide a rough concept of object placement and shape as well as distinguishing distinct object instances, while depth maps encode the rough relative position and shape of the visible objects. Together with input-view images, these conditions provide both global structural information like object placement and local fine-grained details like object shape. Concretely speaking, we normalize the image rendered with object instance IDs of the input view to create a continuous object mask image  $\widehat{\mathbf{M}}$ . We then replicate the depth map  $\widehat{\mathbf{D}}$  and object mask image  $\widehat{\mathbf{M}}$  into three channels to simulate RGB images. These two structure-aware feature images, along with the input image  $\widehat{\mathbf{x}}_0$ , are passed into a VAE to obtain latent features, which will be later concatenated with the noisy target view image  $\mathbf{x}_t$  as input to the denoising U-Net. Note that both object mask and depth can be obtained with off-the-shelf detectors during the inference stage, such as SAM [33] and Marigold [31]. After introducing these additional conditions, the learning objective of MOVIS becomes:

$$\mathbb{E}[\|\epsilon_\theta(\alpha_t \mathbf{x}_0 + \sigma_t \epsilon, t, C_{SA}(\widehat{\mathbf{x}}_0, R, T, \widehat{\mathbf{D}}, \widehat{\mathbf{M}})) - \epsilon\|_2^2]. \quad (2)$$

We use  $C_{SA}(\cdot)$  as a shorthand for the structure-aware view-conditioned feature throughout the paper.

**Auxiliary Novel View Mask Prediction Task** Input-view depth maps and mask images are intended to help the model indirectly understand the structure of multi-object compositions by incorporating additional structure-aware information into the input. To encourage the model to better grasp overall structure, particularly its ability to generate it, we propose leveraging structural information (*i.e.*, mask image) prediction under the target view as an auxiliary task, providing more direct supervision. Our approach draws inspiration from classifier guidance [13], where a classifier  $p_\phi(y|x_t, t)$  guides the denoising process of image  $x_t$  to meet the criterion  $y$  via incorporating the gradient  $\nabla_{x_t} \log p_\phi(y|x_t, t)$  during the inference process as an auxiliary guidance. Similarly, to improve the model’s ability to

learn compositional structure, particularly in synthesizing novel view plausible object placement (position and orientation), we introduce an auxiliary task during training: predicting object mask images  $\mathbf{M}_t \sim p(\mathbf{M}_t|\mathbf{x}_t, t, C_{SA}(\cdot))$  under target view. This prediction is conditioned on the noisy target-view image  $\mathbf{x}_t$ , timestep  $t$  and input-view structure-aware feature  $C_{SA}(\cdot)$ , derived from the final layer of the denoising U-Net. We jointly train the mask predictor and denoising U-Net following:

$$\mathbb{E}[\|\epsilon_\theta(\alpha_t \mathbf{x}_0 + \sigma_t \epsilon, t, C_{SA}(\cdot)) - \epsilon\|_2^2 + \gamma \|\mathbf{M}_{tgt} - \mathbf{M}_t\|_2^2], \quad (3)$$

where we use  $\mathbf{M}_{tgt}$  to denote the ground-truth target-view image, and we use the weight  $\gamma = 0.1$  to balance the diffusion loss and mask prediction loss.

### 3.3. Structure-Guided Sampling Scheduler

Inspired by previous works [4, 30] that identify the importance of different scheduling strategies, we provide an in-depth analysis of the inference process of multi-object novel view synthesis, where we adopt a DDIM [70] sampler:

$$\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \mathbf{x}'_0 + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \mathbf{F} + \sigma_t \epsilon_t. \quad (4)$$

where  $\mathbf{x}'_0 = (\mathbf{x}_t - \sqrt{1 - \alpha_t} \cdot \mathbf{F}) / \sqrt{\alpha_t}$ . We use  $\mathbf{F}$  as a shorthand for  $\epsilon_\theta(\mathbf{x}_t, t, C_{SA}(\cdot))$  and  $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . We examine the predicted  $\mathbf{x}'_0$  (as in Eq. (4)) and the predicted mask image  $\mathbf{M}_t$  at various timesteps during the inference process as they offer direct visualizations for analysis. These visualized results are presented in Fig. 3.

In Fig. 3, we observe that a blurry image, which indicates the approximate placement of each object, is quickly restored in the early stages (*i.e.*, larger  $t$ ) of the inference process. This suggests that global structural information is prioritized for the model to learn during this stage. Accurate object placements are crucial for synthesizing reasonable novel view images, as incorrect placement predictions indicate a fundamental misunderstanding of the compositional structure. This underscores the importance of training the model with a larger  $t$  during the initial training periods, which is even more important for multi-object NVS scenarios considering the increased compositional complexity compared with a single object. Conversely, a mask with a clear boundary is not predicted until a later stage of the sampling process (*i.e.*, smaller  $t$ ). This is because accurate mask prediction depends heavily on a relatively noiseless image. Therefore, to capture fine-grained geometry and appearance details of objects, it is essential to train the model with a smaller  $t$  during later training periods.

Recognizing the importance of timestep  $t$  in balancing the learning of global placement information and local fine-grained details, we propose to adjust the original timestep sampling process to:

$$t \sim \mathcal{U}(1, 1000) \rightarrow t \sim \mathcal{N}(\mu(s), \sigma), \quad (5)$$




Figure 3. **Visualization of inference.** The early stage of the denoising process focuses on restoring global object placements, while the prediction of object masks requires a relatively noiseless image to recover fine-grained geometry. This motivates us to seek a balanced timestep sampling scheduler during training. The model trained *w/ shift* yields better mask prediction and thus recovers an image with more details and sharp object boundary. The *w/o shift* here refers to not shifting the  $\mu$  value.

where  $\mu(s) = \mu_{\text{local}} + (\mu_{\text{global}} - \mu_{\text{local}}) \cdot \frac{s}{T_s}$  and  $s$  denotes the model training iteration,  $T_s$  denotes the total number of training steps,  $\sigma = 200$  is a constant variance. We sample the timestep  $t$  from a Gaussian distribution with mean  $\mu(s)$  following a linear decay from a large value  $\mu_{\text{global}} = 1000$  to a small value  $\mu_{\text{local}} = 500$ . This approach allows the model to initially learn correct global object placement information and gradually turn its focus to refining detailed object geometry in later training stages. In practice, we include a warmup period with 4000 training steps sampling  $t$  with a fixed  $\mu(s) = \mu_{\text{global}}$ . After the warmup, we use the linear decay schedule over 2000 steps, and then stabilize the learning for fine-grained details after 6000 steps where we use  $\mu(s) = \mu_{\text{local}}$ .

## 4. Experiments

### 4.1. Experiment Setup

We focus on multi-object composite NVS tasks in indoor scenes, with an emphasis on foreground objects, examining novel view structural plausibility regarding object placement, geometry, appearance, and cross-view consistency

with input view. This choice stems from the recent advancements in object segmentation [33], while we leave the background modeling for future work.

**Datasets.** To facilitate the training and evaluation of our proposed method, we curate a scalable synthetic dataset Compositional 3D-FUTURE (C3DFS), comprising 100k composites for training and 5k for testing. Each composite is created by composing pre-filtered furniture items from 3D-FUTURE [16] using a heuristic strategy to avoid collision and penetration. Beyond C3DFS, we emphasize testing the generalization capability by benchmarking our method on Room-Texture [54] and Objaverse [12]. We also evaluate our model on diverse indoor scenes from both the synthetic dataset 3D-FRONT [15] and the real-world dataset SUNRGB-D [71]. Refer to *supplementary materials* for more details.

**Baselines.** We compare our method against three recent novel view synthesis methods including Zero-1-to-3 [46], ZeroNVS [67], and Free3D [97]. The original Zero-1-to-3 is trained on extensive object-level datasets. Therefore, we also re-train Zero-1-to-3 on our synthetic dataset C3DFS, denoted as Zero-1-to-3<sup>†</sup>, for a fair comparison. ZeroNVS is



Figure 4. **Qualitative results of NVS and cross-view matching.** Our method generates plausible novel-view images across various datasets, surpassing baselines regarding object placement, shape, and appearance. In cross-view matching, points of the same color indicate correspondences between the input and target views. We achieve a higher number of matched points with more precise locations.

Table 1. **Quantitative results of multi-object NVS, Object Placement, and Cross-view Consistency.** We evaluate on C3DFS test set, along with *generalization experiments* on Room-Texture [54] and Objaverse [12]. † indicates re-training on C3DFS.

Dataset	Method	Novel View Synthesis			Placement	Cross-view Consistency	
		PSNR(↑)	SSIM(↑)	LPIPS(↓)	IoU(↑)	Hit Rate(↑)	Dist(↓)
C3DFS	ZeroNVS	10.704	0.533	0.481	21.6	1.4	135.2
	Zero-1-to-3	14.255	0.771	0.302	33.7	4.4	86.7
	Zero-1-to-3 <sup>†</sup>	14.811	0.794	0.283	34.4	1.3	117.9
	Free3D	14.390	0.774	0.297	34.2	4.8	83.6
	Ours	<b>17.432</b>	<b>0.825</b>	<b>0.171</b>	<b>58.1</b>	<b>19.3</b>	<b>44.9</b>
Objaverse	ZeroNVS	10.557	0.513	0.486	17.3	1.3	135.8
	Zero-1-to-3	15.850	0.810	0.259	34.7	6.0	80.2
	Zero-1-to-3 <sup>†</sup>	15.433	0.815	0.273	29.7	1.2	124.1
	Free3D	15.980	0.813	0.254	35.6	6.1	77.0
	Ours	<b>17.749</b>	<b>0.840</b>	<b>0.169</b>	<b>51.3</b>	<b>17.0</b>	<b>47.2</b>
Room-Texture	ZeroNVS	8.217	0.647	0.487	8.2	1.0	138.0
	Zero-1-to-3	9.860	0.712	0.406	13.9	2.3	104.2
	Zero-1-to-3 <sup>†</sup>	8.342	0.657	0.452	13.5	0.4	156.1
	Free3D	9.623	0.705	0.415	19.7	2.1	104.6
	Ours	<b>10.014</b>	<b>0.718</b>	<b>0.366</b>	<b>24.2</b>	<b>4.4</b>	<b>78.1</b>

trained on a mixture of real-world images with background, so we use images with backgrounds as its input if possible for a fair comparison.

**Metrics.** We utilize PSNR, SSIM, and LPIPS as met-

rics for evaluating the quality of *Novel View Synthesis*. To assess global object *Placement*, we compute the foreground-background IoU with ground-truth masks. Finally, we propose metrics to evaluate *Cross-view Consis-*

tenacy with image-matching. More specifically, we first apply MAST3R [35] to acquire the image matching between the input-view image and target-view image for both ground truth and model predictions. With the ground-truth matching as references, we compute each method’s *Hit Rate* and the nearest matching distance (*Dist.*). *Hit Rate* measures the proportion of predicted matches that align with the ground truth matches. *Dist.* quantifies the distance between the predicted matching and ground-truth matching in the target view. Please refer to *supplementary materials* for more details about the metrics.

## 4.2. Results and Discussions

Fig. 4 presents qualitative results of multi-object NVS and cross-view matching visualization on different datasets, with quantitative results in Tab. 1. We summarize the following key observations:

1. Our method realizes the highest PSNR and generates high-quality images under novel views, closely aligned with the ground truth images, especially regarding novel-view object placement (position and orientation), shape, and appearance. In contrast, the baseline models struggle to accurately capture the compositional structure under novel views. For example, in the first row, the red bed is incorrectly oriented in Zero-1-to-3 and is either missing or distorted in other baselines.
2. From the visualized cross-view matching results and the metrics in Tab. 1, it is evident that our method significantly outperforms the baseline approaches in *Cross-view Consistency*. It achieves a much higher IoU and *Hit Rate* while exhibiting a considerably lower matching distance. The visualized results are consistent with the metrics, further validating our method’s accuracy in capturing cross-view structural consistency, which cannot be reflected by existing NVS metrics.
3. Our model exhibits strong generalization capabilities on unseen datasets, *e.g.*, Room-Texture and Objaverse. We demonstrate more qualitative results on in-the-wild datasets, including 3D-FRONT, SUNRGB-D, ScanNet++ [93], and RealEstate10K [98] in *supplementary materials*. We showcase potential applications, including object removal and reconstruction in *supplementary materials*. Further discussion about limitation and failure cases are presented in *supplementary materials*.

## 4.3. Ablation Study

To verify the efficacy of each component, we perform an ablation study on our key technical designs, including the depth input (*w/o depth*), mask prediction auxiliary task (*w/o mask*), and the scheduler (*w/o sch.* learns with a uniform sampler  $t \sim \mathcal{U}(1, 1000)$ ). Results in Tab. 2 show that the auxiliary mask prediction task and the timestep sampler are the most critical components, significantly affecting all the



Figure 5. **Qualitative comparison for ablation study.** Excluding mask predictions or the scheduler reduces the model’s ability to learn object placement, as shown by the brown cabinet example.

Table 2. **Ablation results on C3DFS.**

Method	Novel View Synthesis			Placement
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	IoU( $\uparrow$ )
w/o depth	17.080	0.819	0.178	57.2
w/o mask	16.914	0.818	0.187	54.7
w/o sch.	16.166	0.808	0.212	49.1
Ours	<b>17.432</b>	<b>0.825</b>	<b>0.171</b>	<b>58.1</b>

metrics and the realistic object recovery as demonstrated by the misoriented brown cabinet in the example from Fig. 5. Without the scheduler, the model produces less accurate object positions, evident both qualitatively and quantitatively. Furthermore, removing depth or mask predictions weakens the model’s understanding of spatial relationships and object existence. This also shows incorporating structure-aware features as inputs, though seemingly intuitive, offers the most straightforward approach to enhancing the model’s structural awareness, particularly given recent advancements in monocular predictors [31, 33]. We present a more comprehensive discussion on the scheduler strategy and ablations in *supplementary materials*.

## 5. Conclusion

We extend diffusion-based novel view synthesizers to handle multi-object compositions in indoor scenes. Our proposed model generalize well across diverse datasets with more accurate object placement, shape, and appearance, showing a stronger cross-view consistency with input view. The core of our approach lies in integrating structure-aware features as additional inputs, an auxiliary mask prediction task, and a structure-guided timestep sampling scheduler. These components enhance the model’s awareness of compositional structure while balancing the learning of global object placement and fine-grained local shape and appearance. Given the prevalence of multi-object compositions in real-world scenes, we believe that our model designs and comprehensive evaluations can offer valuable insights for advancing NVS models in more complex environments.

## Acknowledgment

This work is supported by the Sichuan Science and Technology Program (2023YFSY0008), China Tower-Peking University Joint Laboratory of Intelligent Society and Space Governance, National Natural Science Foundation of China (61632003, 61375022, 61403005), Grant SC-ITLAB20017 of Intelligent Terminal Key Laboratory of SiChuan Province, Beijing Advanced Innovation Center for Intelligent Robots and Systems (2018IRS11), and PEK-SenseTime Joint Laboratory of Machine Vision.

## References

- [1] Garrick Brazil, Abhinav Kumar, Julian Straub, Nikhila Ravi, Justin Johnson, and Georgia Gkioxari. Omni3d: A large benchmark and model for 3d object detection in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [2] David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [3] Dave Zhenyu Chen, Haoxuan Li, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner. Scenetex: High-quality texture synthesis for indoor scenes via diffusion priors. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [4] Ting Chen. On the importance of noise scheduling for diffusion models. *arXiv preprint arXiv:2301.10972*, 2023. 5
- [5] Yabo Chen, Jiemin Fang, Yuyang Huang, Taoran Yi, Xiaopeng Zhang, Lingxi Xie, Xinggang Wang, Wenrui Dai, Hongkai Xiong, and Qi Tian. Cascade-zero123: One image to highly consistent 3d with self-prompted nearby views. In *European Conference on Computer Vision (ECCV)*, 2024. 2, 3
- [6] Yixin Chen, Junfeng Ni, Nan Jiang, Yaowei Zhang, Yixin Zhu, and Siyuan Huang. Single-view 3d scene reconstruction with high-fidelity shape and texture. In *International Conference on 3D Vision (3DV)*, 2024. 3, 14
- [7] Yongwei Chen, Tengfei Wang, Tong Wu, Xingang Pan, Kui Jia, and Ziwei Liu. Comboverse: Compositional 3d assets creation using spatially-aware diffusion guidance. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [8] Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [9] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. 14
- [10] Jieming Cui, Tengyu Liu, Nian Liu, Yaodong Yang, Yixin Zhu, and Siyuan Huang. Anyskill: Learning open-vocabulary physical skill for interactive agents. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [11] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 3
- [12] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3, 4, 6, 7, 15, 16
- [13] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 5
- [14] Andreea Dogaru, Mert Özer, and Bernhard Egger. Generalizable 3d scene reconstruction via divide and conquer from a single view. *arXiv preprint arXiv:2404.03421*, 2024. 3
- [15] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Binqiang Zhao, et al. 3d-front: 3d furnished rooms with layouts and semantics. In *International Conference on Computer Vision (ICCV)*, 2021. 2, 6, 16
- [16] Huan Fu, Rongfei Jia, Lin Gao, Mingming Gong, Binqiang Zhao, Steve Maybank, and Dacheng Tao. 3d-future: 3d furniture shape with texture. *International Journal of Computer Vision (IJCV)*, 129:3313–3337, 2021. 6, 14, 15, 17
- [17] Ran Gong, Jiangyong Huang, Yizhou Zhao, Haoran Geng, Xiaofeng Gao, Qingyang Wu, Wensi Ai, Ziheng Zhou, Demetri Terzopoulos, Song-Chun Zhu, et al. Arnold: A benchmark for language-grounded task learning with continuous states in realistic 3d scenes. *arXiv preprint arXiv:2304.04321*, 2023. 2
- [18] Ming Gui, Johannes S Fischer, Ulrich Prestel, Pingchuan Ma, Dmytro Kotovenko, Olga Grebenkova, Stefan Andreas Baumann, Vincent Tao Hu, and Björn Ommer. Depthfm: Fast monocular depth estimation with flow matching. *arXiv preprint arXiv:2403.13788*, 2024. 14, 15
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3
- [20] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *International Conference on Computer Vision (ICCV)*, 2023. 2
- [21] Jiangyong Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiong Jia, and Siyuan Huang. An embodied generalist agent in 3d world. In *International Conference on Machine Learning (ICML)*, 2024. 2
- [22] Jiangyong Huang, Baoxiong Jia, Yan Wang, Ziyu Zhu, Xiongkun Linghu, Qing Li, Song-Chun Zhu, and Siyuan Huang. Unveiling the mist over 3d vision-language understanding: Object-centric evaluation with chain-of-analysis. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2
- [23] Zehuan Huang, Hao Wen, Junting Dong, Yaohui Wang, Yangguang Li, Xinyuan Chen, Yan-Pei Cao, Ding Liang, Yu

- Qiao, Bo Dai, et al. Epidiff: Enhancing multi-view synthesis via localized epipolar-constrained diffusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3
- [24] Zhipeng Huang, Wangbo Yu, Xinhua Cheng, ChengShu Zhao, Yunyang Ge, Mingyi Guo, Li Yuan, and Yonghong Tian. Roompainter: View-integrated diffusion for consistent indoor scene texturing. *arXiv preprint arXiv:2412.16778*, 2024. 2
- [25] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [26] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [27] Baoxiong Jia, Yixin Chen, Huangyue Yu, Yan Wang, Xuesong Niu, Tengyu Liu, Qing Li, and Siyuan Huang. Sceneverse: Scaling 3d vision-language learning for grounded scene understanding. In *European Conference on Computer Vision (ECCV)*, 2024. 2
- [28] Nan Jiang, Zimo He, Zi Wang, Hongjie Li, Yixin Chen, Siyuan Huang, and Yixin Zhu. Autonomous character-scene interaction synthesis from text instruction. In *SIGGRAPH Asia 2024 Conference Papers*, 2024. 2
- [29] Nan Jiang, Zhiyuan Zhang, Hongjie Li, Xiaoxuan Ma, Zan Wang, Yixin Chen, Tengyu Liu, Yixin Zhu, and Siyuan Huang. Scaling up dynamic human-scene interaction modeling. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [30] Yifan Jiang, Hao Tang, Jen-Hao Rick Chang, Liangchen Song, Zhangyang Wang, and Liangliang Cao. Efficient-3dim: Learning a generalizable single-image novel-view synthesizer in one day. *arXiv preprint arXiv:2310.03015*, 2023. 3, 5, 13
- [31] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3, 5, 8
- [32] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 2023. 3
- [33] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *International Conference on Computer Vision (ICCV)*, 2023. 3, 5, 6, 8, 14, 15
- [34] Xin Kong, Shikun Liu, Xiaoyang Lyu, Marwan Taher, Xiaojuan Qi, and Andrew J Davison. Eschernet: A generative model for scalable view synthesis. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 13, 19
- [35] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision (ECCV)*, 2024. 2, 8, 15
- [36] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning (ICML)*, 2023. 3
- [37] Puhao Li, Tengyu Liu, Yuyang Li, Yiran Geng, Yixin Zhu, Yaodong Yang, and Siyuan Huang. Gendexgrasp: Generalizable dexterous grasping. In *International Conference on Robotics and Automation (ICRA)*, 2023. 2
- [38] Puhao Li, Tengyu Liu, Yuyang Li, Muzhi Han, Haoran Geng, Shu Wang, Yixin Zhu, Song-Chun Zhu, and Siyuan Huang. Ag2manip: Learning novel manipulation skills with agent-agnostic visual and action representations. In *International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [39] Yishu Li, Wen Hui Leng, Yiming Fang, Ben Eisner, and David Held. Flowbothd: History-aware diffuser handling ambiguities in articulated objects manipulation. *arXiv preprint arXiv:2410.07078*, 2024.
- [40] Yuyang Li, Bo Liu, Yiran Geng, Puhao Li, Yaodong Yang, Yixin Zhu, Tengyu Liu, and Siyuan Huang. Grasp multiple objects with one hand. *IEEE Robotics and Automation Letters (RA-L)*, 2024. 2
- [41] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [42] Yukang Lin, Haonan Han, Chaoqun Gong, Zunnan Xu, Yachao Zhang, and Xiu Li. Consistent123: One image to highly consistent 3d asset using case-aware diffusion priors. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 2024. 2, 3
- [43] Xiongkun Linghu, Jiangyong Huang, Xuesong Niu, Xiaojian Ma, Baoxiong Jia, and Siyuan Huang. Multi-modal situated reasoning in 3d scenes. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2
- [44] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 2, 3
- [45] Minghua Liu, Ruoxi Shi, Linghao Chen, Zhuoyang Zhang, Chao Xu, Xinyue Wei, Hansheng Chen, Chong Zeng, Jiayuan Gu, and Hao Su. One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [46] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *International Conference on Computer Vision (ICCV)*, 2023. 1, 2, 3, 4, 6, 13
- [47] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023. 2, 3, 19
- [48] Yu Liu, Baoxiong Jia, Yixin Chen, and Siyuan Huang. Slotlifter: Slot-guided feature lifting for learning object-

- centric radiance fields. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [49] Yu Liu, Baoxiong Jia, Ruijie Lu, Junfeng Ni, Song-Chun Zhu, and Siyuan Huang. Building interactable replicas of complex articulated objects via gaussian splatting. *arXiv preprint arXiv:2502.19459*, 2025. 2
- [50] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [51] Guanxing Lu, Shiyi Zhang, Ziwei Wang, Changliu Liu, Jiwen Lu, and Yansong Tang. Manigaussian: Dynamic gaussian splatting for multi-task robotic manipulation. In *European Conference on Computer Vision (ECCV)*, 2024. 2
- [52] Ruijie Lu, Yixin Chen, Yu Liu, Jiayang Tang, Junfeng Ni, Diwen Wan, Gang Zeng, and Siyuan Huang. Taco: Taming diffusion for in-the-wild video amodal completion. *arXiv preprint arXiv:2503.12049*, 2025. 3
- [53] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [54] Rundong Luo, Hong-Xing Yu, and Jiajun Wu. Unsupervised discovery of object-centric neural fields. *arXiv preprint arXiv:2402.07376*, 2024. 2, 6, 7, 16, 17
- [55] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [56] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [57] Junfeng Ni, Yixin Chen, Bohan Jing, Nan Jiang, Bin Wang, Bo Dai, Puhao Li, Yixin Zhu, Song-Chun Zhu, and Siyuan Huang. Phyrecon: Physically plausible neural scene reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [58] Junfeng Ni, Yu Liu, Ruijie Lu, Zirui Zhou, Song-Chun Zhu, Yixin Chen, and Siyuan Huang. Decompositional neural scene reconstruction with generative diffusion prior. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [59] Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, and Jian Jun Zhang. Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [60] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 5, 13
- [61] Ege Ozguroglu, Ruoshi Liu, Dídac Surís, Dian Chen, Achal Dave, Pavel Tokmakov, and Carl Vondrick. pix2gestalt: Amodal segmentation by synthesizing wholes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3, 16
- [62] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 3
- [63] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 2
- [64] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, 2021. 4
- [65] Adam Rashid, Satvik Sharma, Chung Min Kim, Justin Kerr, Lawrence Yunliang Chen, Angjoo Kanazawa, and Ken Goldberg. Language embedded radiance fields for zero-shot task-oriented grasping. In *Conference on Robot Learning (CoRL)*, 2023. 2
- [66] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3, 4
- [67] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, et al. Zeronvs: Zero-shot 360-degree view synthesis from a single image. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3, 6
- [68] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. 2, 3, 19
- [69] Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv:2308.16512*, 2023. 2, 19
- [70] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2020. 5
- [71] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3, 6, 16, 17
- [72] Jiayang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 2, 3
- [73] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d

- creation from a single image with diffusion prior. In *International Conference on Computer Vision (ICCV)*, 2023. 2, 3
- [74] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. In *European Conference on Computer Vision (ECCV)*, 2024. 2
- [75] Jiaxiang Tang, Ruijie Lu, Xiaokang Chen, Xiang Wen, Gang Zeng, and Ziwei Liu. Intex: Interactive text-to-texture synthesis via unified depth-aware inpainting. *arXiv preprint arXiv:2403.11878*, 2024. 3
- [76] Joseph Tung, Gene Chou, Ruojin Cai, Guandao Yang, Kai Zhang, Gordon Wetzstein, Bharath Hariharan, and Noah Snavely. Megascenes: Scene-level view synthesis at scale. In *European Conference on Computer Vision (ECCV)*, 2024. 2, 3
- [77] Basile Van Hoorick, Pavel Tokmakov, Simon Stent, Jie Li, and Carl Vondrick. Tracking through containers and occluders in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 16
- [78] Basile Van Hoorick, Rundi Wu, Ege Ozguroglu, Kyle Sargent, Ruoshi Liu, Pavel Tokmakov, Achal Dave, Changxi Zheng, and Carl Vondrick. Generative camera dolly: Extreme monocular dynamic novel view synthesis. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [79] Peng Wang and Yichun Shi. Imagedream: Image-prompt multi-view diffusion for 3d generation. *arXiv preprint arXiv:2312.02201*, 2023. 3, 19
- [80] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Sriniwasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [81] Qi Wang, Ruijie Lu, Xudong Xu, Jingbo Wang, Michael Yu Wang, Bo Dai, Gang Zeng, and Dan Xu. Roomtex: Texturing compositional indoor scenes via iterative inpainting. In *European Conference on Computer Vision (ECCV)*, 2024. 2, 3
- [82] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 16
- [83] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. *arXiv preprint arXiv:2210.04628*, 2022. 4
- [84] Haohan Weng, Tianyu Yang, Jianan Wang, Yu Li, Tong Zhang, CL Chen, and Lei Zhang. Consistent123: Improve consistency for one image to 3d object synthesis. *arXiv preprint arXiv:2310.08092*, 2023. 2, 3
- [85] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 13
- [86] Kailu Wu, Fangfu Liu, Zhihan Cai, Runjie Yan, Hanyang Wang, Yating Hu, Yueqi Duan, and Kaisheng Ma. Unique3d: High-quality and efficient 3d mesh generation from a single image. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [87] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv preprint arXiv:2412.01506*, 2024. 2
- [88] Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [89] Katherine Xu, Lingzhi Zhang, and Jianbo Shi. Amodal completion via progressive mixed context diffusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9099–9109, 2024. 3, 16
- [90] Jiayu Yang, Ziang Cheng, Yunfei Duan, Pan Ji, and Hongdong Li. Consistnet: Enforcing 3d consistency for multi-view images diffusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 19
- [91] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [92] Yandan Yang, Baoxiong Jia, Peiyuan Zhi, and Siyuan Huang. Physcene: Physically interactable 3d scene synthesis for embodied ai. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [93] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *International Conference on Computer Vision (ICCV)*, 2023. 8
- [94] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [95] Guanqi Zhan, Chuanxia Zheng, Weidi Xie, and Andrew Zisserman. Amodal ground truth and completion in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3, 16
- [96] Xiaohang Zhan, Xingang Pan, Bo Dai, Ziwei Liu, Dahua Lin, and Chen Change Loy. Self-supervised scene deocclusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 16
- [97] Chuanxia Zheng and Andrea Vedaldi. Free3d: Consistent novel view synthesis without 3d representation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 6
- [98] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. In *SIGGRAPH*, 2018. 8
- [99] Yan Zhu, Yuandong Tian, Dimitris Metaxas, and Piotr Dollár. Semantic amodal segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 16

# MOVIS: Enhancing Multi-Object Novel View Synthesis for Indoor Scenes

## Supplementary Material

### A. Model details

#### A.1. DINO patch feature and camera view embedding

The original image encoder of Stable Diffusion is CLIP, which excels at aligning images with text. Other image encoders like DINO-v2 [60] or ConvNeXtv2 [85] may provide denser image features that may benefit generation tasks as mentioned by previous works [30, 34]. Therefore, we opt to use the DINO feature instead of the original CLIP feature in our network following [30]. To inject the DINO patch feature into our network, we encode the input view image using DINO-v2 [60] “norm patchtokens”, whose shape dimension is  $(b, 16, 16, 1024)$ . We will simply flatten it into  $(b, 256, 1024)$  to apply cross-attention, and  $b$  means batch size here.

As for the camera view embedding, we choose to embed it using a 6 degrees of freedom (6DoF) representation. To be specific, let  $E_i$  be the extrinsic matrix under the input view and  $E_j$  be the extrinsic matrix under the output view, we represent relative camera pose change as  $E_i^{-1}E_j$ . We will also flatten it into 16 dimensions to concatenate it to the image feature. Afterwards, we will replicate the 16-dimension embedding 256 times to concatenate the embedding to every channel of the DINO feature map. A projection layer will later be employed to project the feature map into  $(b, 256, 768)$  to match the dimension of the CLIP encoder, which was originally used by Stable Diffusion so that we can fine-tune the pre-trained checkpoint. It is worth noting that we also tried other novel view synthesizer’s camera embedding like Zero-1-to-3 [46] using a 3DoF spherical coordinates in early experiments, but we found that it does not make much of a difference.

#### A.2. Depth and mask condition

In this section, we will explain how input view depth and mask are incorporated as additional conditioning inputs. For depth maps, regions with infinite depth values are assigned a value equal to twice the maximum finite depth value in the rest of the image. After this adjustment, we apply a normalization technique to scale the depth values to the range of  $[-1, 1]$ , enabling the use of the same VAE architecture as for images.

For mask images, we assign unique values to different object instances in the input view. For instance, if there are four objects in the multi-object composite, they will be labeled as 1, 2, 3, and 4, respectively, while the background will be assigned a value of 0. The same normalization technique used for depth maps is applied to these mask images.

These mask images, like all other inputs, are processed by the VAE, with all images set to a resolution of  $256 \times 256$ .

#### A.3. Supervision for auxiliary mask prediction task

To implement the auxiliary mask prediction task, we encode the output view mask images into the same latent space as the input view mask images. Object instances viewed from different angles will be assigned the same value, which is ensured during the curation of our compositional dataset. Supervision is directly applied to the latent mask features extracted from the final layer of the denoising U-Net. Only the input view mask images are required during inference, simplifying the process while preserving consistency across views.

#### A.4. Timestep scheduler




Figure S.6. Illustration of different timestep sampling strategies.

Table S.3. Ablation on different strategies. Incorporating sampling strategies significantly improves the model performance, while the linear decline (LDC) achieves the best.

Dataset	Method	Novel View Synthesis		
		PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
C3DFS	w/o sch.	16.166	0.808	0.212
	KMS	17.148	0.823	0.175
	LIND	17.279	0.824	0.172
	LDC	<b>17.432</b>	<b>0.825</b>	<b>0.171</b>

Though we finally employed a linearly declining strategy, we experimented with several alternatives. Specifically, we tested linearly declining the mean of the Gaussian




Figure S.7. **Comparison of different strategies.** The predicted images and mask images under novel views using different strategies are visualized. We can observe that images predicted by the KMS strategy possess weird and blurry color while LDC strategy seems to be slightly better than LIND.

distribution (LDC), linearly increasing the mean after a sudden drop (LIND), and keeping the mean constant (KMS). These strategies are illustrated in Fig. S.6. The metrics on the test set of our C3DFS are provided in Tab. S.3, with some visual comparisons in Fig. S.7. *w/o sch.* in Tab. S.3 refers to applying a uniform sampler, same as the one in the main paper. From the results, we observe that LDC achieves slightly better performance than LIND and KMS, largely outperforming *w/o sch.*

However, we observed significant visual artifacts such as weird colors and extremely blurry mask images when combining the auxiliary mask prediction task with the KMS sampling strategy, as demonstrated in Fig. S.7. For example, the bed in the second example possesses unclear object boundaries and distorted texture. We believe this is due to KMS focusing primarily on denoising at larger timesteps, which provides limited guidance for recovering mask images and refining fine-grained geometry and appearance. Consequently, without a dedicated period for denoising smaller timesteps, the per-object shape and appearance appear distorted and unrealistic.

## B. Experiment Details

### B.1. Implementation Details

We solely utilize the data from C3DFS as the training source for our model. The training process takes around 2 days on 8 NVIDIA A100 (80G) GPUs, employing a batch size of 172 per GPU. The exact training steps are 8,000 steps. During the inference process, we apply 50 DDIM steps and set the guidance scale to 3.0. We use DepthFM [18] and SAM [33] to extract the depth maps and object masks when

they are not available, as well as for all real-world images.

### B.2. Datasets

**C3DFS** We use the furniture models from the 3D-FUTURE dataset [16] to create our synthetic multi-object compositional data. The 3D-FUTURE dataset contains 9,992 detailed 3D furniture models with high-resolution textures and labels. Following previous work [6], we categorize the furniture into seven groups: bed, bookshelf, cabinet, chair, nightstand, sofa, and table. To ensure unbiased evaluation, we further split the furniture into distinct training and test sets, ensuring that none of the test set items are seen during training.

After filtering the furniture, we first determine the number of pieces to include in each composite, which is randomly selected to be between 3 and 6. Next, we establish a probability distribution based on the different types of furniture items and sample each piece according to this distribution. To prevent collisions and penetration between furniture items, we employ a heuristic strategy. Specifically, for each furniture item to be added, we apply a random scale adjustment within the range of  $[0.95, 1.05]$ , as the inherent scale of the furniture models accurately reflects real-world sizes. We also rotate each model by a random angle to introduce additional variability. Once these adjustments are complete, we begin placing the furniture items in the scene. The first item is positioned at the center of the scene at coordinates  $(0, 0, 0)$ . Subsequent objects are added one by one, initially placed at the same central location. Since this results in inevitable collisions, we randomly sample a direction and gradually move the newly added item along this vector until there is no intersection between the bounding boxes of the objects. By following these steps, we generate a substantial number of multiple furniture items composites, ultimately creating a training set of 100,000 composites and a test set of 5,000 to evaluate the capabilities of our network.

After placing all the furniture items, we render multi-view images to facilitate training, using Blender [9] as our renderer due to its high-quality output. We first normalize each composite along its longest axis. To simulate real-world camera poses and capture meaningful multi-object compositions, we employ the following method for sampling camera views.

Cameras are randomly sampled using spherical coordinates, with a radius range of  $[1.3, 1.7]$  and an elevation angle range of  $[2^\circ, 40^\circ]$ . There are no constraints on the azimuth angle, allowing the camera to rotate freely around multiple objects. The chosen ranges for the radius and elevation angles are empirical. In addition to determining the camera positions, we establish a "look-at" point to compute the camera pose. This point is randomly selected on a spherical shell with a radius range of  $[0.01, 0.2]$ .

To enhance the model's compositional structural aware-

Table S.4. Availability of conditions in different datasets.

	C3DFS	Room-Texture	Objaverse	SUNRGB-D	3D-FRONT
depth	✓	×	✓	×	×
mask	✓	✓	✓	×	×

ness, we also render depth maps and instance masks (both occluded and unoccluded) from 12 different viewpoints. The unoccluded instance mask ensures that if one object is blocked by another, the complete amodal mask of the occluded object is still provided, regardless of any obstructions. Although these unoccluded instance masks are not currently necessary for our network, we render them for potential future use.

**Objaverse** To evaluate our network’s generalization capability, we create a small dataset comprising 300 composites sourced from Objaverse [12]. Specifically, we utilize the provided LVIS annotations to select categories that are commonly found in indoor environments, such as beds, chairs, sofas, dressers, tables, and others. Since the meshes from Objaverse vary in scale, we rescale each object based on reference object scales from the 3D-FUTURE dataset [16]. The composition and rendering processes follow the same strategy employed in C3DFS.

---

#### Algorithm 1 Hit Rate Computation

```

1: // Obtain image-matching pairs using MAST3R and
   // save in a list
2: Pairsgt = MAST3R(GT)
3: Pairsours = MAST3R(Ours)
4: // Each element in the list is a four-element tuple p =
   // (x0, y0, x1, y1)
5: // (x0, y0) refers to the point in the input view image
   // and (x1, y1) the point in output view image
6: Hits = 0
7: For pgt in Pairsgt
8:   // poursi is the i-th element of Pairsours
9:   // p[:2] refers to the first two element in the tuple
   // and p[2:] the last two
10:  i* = arg mini(L2(pgt[:2], poursi[:2]))
11:  IF L2(pgt[:2], poursi*[:2]) < 20 and
   // L2(pgt[2:], poursi*[2:]) < 20
12:  // Successfully hit one, delete it from gt pairs
   // and ours pairs
13:  Hits ← Hits + 1
14:  POP(Pairsours, poursi*)
15: return Hits/len(Pairsgt)

```

---

**Inference Details** Since our model requires input-view depth map and mask images as additional inputs, we need to

---

#### Algorithm 2 Nearest Matching Distance Computation

```

1: // The notations are the same as the one in Algorithm 1
2: Pairsgt = MAST3R(GT)
3: Pairsours = MAST3R(Ours)
4: Dist = EmptyList()
5: For pgt in Pairsgt
6:  i* = arg mini(L2(pgt[:2], poursi[:2]))
7:  IF L2(pgt[:2], poursi*[:2]) < 20
8:    Append(Dist, L2(pgt[2:], poursi*[2:]))
9:    POP(Pairsours, poursi*)
10: return Mean(Dist)

```

---

use DepthFM [18] and SAM [33] to extract the depth maps and object masks when they are not available, as well as for all real-world images. We show whether all the used datasets have provided depth maps and mask images in Tab. S.4. ‘×’ means they do not provide such conditions while ‘✓’ means they do provide such conditions.

### B.3. Metrics

**Intersection over Union (IoU)** Since all baseline methods do not possess the concept of every object instance, we compute a foreground-background IoU for comparison. This metric can provide a rough concept of the overall placement alignment with ground truth images. We extract the foreground object mask by converting the generated image to grayscale ( $I_L$ ). Given that the generated image has a white background, we compute the foreground mask  $M$  as  $M = I_L < \beta_{th}$ , where  $\beta_{th}$  is a threshold that is fixed as 250.

**Cross-view Matching** As outlined in the main paper, we introduce two metrics to systematically evaluate cross-view consistency with the input view: **Hit Rate** and **Nearest Matching Distance**. Since direct assessment of cross-view consistency is not feasible by merely evaluating the success matches between each method’s predicted novel view images and the input view image, we opt to how far the predicted matches deviate from the ground-truth matches.

We first compute ground-truth matching points and every model’s matching points using MAST3R [35] upon the input view image and the output view image (ground truth or predicted). Each matching pair is represented as a four-element tuple  $(x^0, y^0, x^1, y^1)$ , where  $(x^0, y^0)$  corresponds



Figure S.8. Visualized results on *in-the-wild* datasets

to the point on the input-view image, and  $(x^1, y^1)$  corresponds to the point on the output-view image.

For each ground-truth matching pair  $(x_{gt}^0, y_{gt}^0, x_{gt}^1, y_{gt}^1)$ , we find the nearest predicted matching pair in each model’s results, denoted as  $(x^0, y^0, x^1, y^1)$ , based on the Euclidean distance between points in the input view image. If both  $L_2\|(x_{gt}^0, y_{gt}^0), (x^0, y^0)\|$  and  $L_2\|(x_{gt}^1, y_{gt}^1), (x^1, y^1)\|$  is smaller than a fixed threshold 20, the match is considered a successful hit. The **Hit Rate** is then calculated as the ratio of successful hits to the total number of ground-truth matches.

For Nearest Matching Distance, we examine whether  $L_2\|(x_{gt}^0, y_{gt}^0), (x^0, y^0)\|$  is within the threshold. For those passing this check, we compute the mean distance  $L_2\|(x_{gt}^1, y_{gt}^1), (x^1, y^1)\|$  as the **Nearest Matching Distance**, averaging over all successful hits. A detailed pseudocode explanation can be found in Algorithm 1 and Algorithm 2.

## B.4. Results

We show more visualized results of our own methods along with ground truth on C3DFS in Fig. S.16, on Objaverse [12] in Fig. S.17, and on Room-Texture [54] in Fig. S.18. More visualized comparisons with baselines on Room-Texture [54], SUNRGB-D [71] and 3D-FRONT [15] are shown in Fig. S.9. More results on in-the-wild datasets are shown in Fig. S.8. A more complete ablation study on other datasets including Objaverse and Room-Texture is shown in Tab. S.5. Some continuous rotation examples on SUNRGB-D and 3D-FRONT are shown in Fig. S.10, and more cross-view matching results without ground-truth pairs as reference are shown in Fig. S.11.

## B.5. Applications

**Object Removal** Since we can predict mask images under novel views, we can support simple image editing tasks like novel view object removal by simply setting a threshold value in the mask image and mask out corresponding pixels to achieve object removal. An example is shown in Fig. S.14.

**Reconstruction** The capability to synthesize novel view images that are consistent with the input view image demonstrates that the model possesses 3D-awareness, which can assist downstream tasks such as reconstruction. We leverage an off-the-shelf reconstruction method DUST3R [82] using the input-view image and novel view images predicted by our method. Two visualized examples are shown in Fig. S.15.

## B.6. Mutual Occlusion

In multi-object compositions, mutual occlusion between objects is very common. Although we did not specifically design the method to make the model aware of mutual occlusion, the model has learned some understanding of these occlusion relationships. A series of research efforts [61, 77, 89, 95, 96, 99] specifically focus on addressing mutual occlusion relationships by predicting the amodal masks or synthesizing amodal appearance, but these models typically do not consider scenarios involving camera view change. Moreover, there may not be a well-established metric to measure how well the model understands mutual occlusion from novel viewpoints. We provide a simple experiment and discussion in this section to illustrate model’s comprehension of mutual occlusion.

First, in the context of novel view synthesis, the com-




Figure S.9. Visualized comparison on Room-Texture [54], SUNRGB-D [71], and 3D-FRONT [16].

Table S.5. Ablation study on various datasets.

Dataset	Method	Novel View Synthesis			Placement	Cross-view Consistency	
		PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	IoU( $\uparrow$ )	Hit Rate( $\uparrow$ )	Dist( $\downarrow$ )
C3DFS	w/o depth	17.080	0.819	0.178	57.2	18.6	<b>42.6</b>
	w/o mask	16.914	0.818	0.187	54.7	14.7	49.6
	w/o sch.	16.166	0.808	0.212	49.1	9.7	48.6
	Ours	<b>17.432</b>	<b>0.825</b>	<b>0.171</b>	<b>58.1</b>	<b>19.3</b>	44.9
Room-Texture	w/o depth	9.829	0.705	<b>0.365</b>	<b>25.7</b>	<b>5.0</b>	<b>76.1</b>
	w/o mask	9.576	0.699	0.384	24.2	2.9	91.4
	w/o sch.	9.173	0.689	0.392	22.4	2.8	92.0
	Ours	<b>10.014</b>	<b>0.718</b>	0.366	24.2	4.4	78.1
Objaverse	w/o depth	17.457	0.835	0.178	50.5	15.9	<b>45.5</b>
	w/o mask	17.176	0.834	0.187	47.3	12.8	53.6
	w/o sch.	16.642	0.825	0.210	43.2	9.6	51.7
	Ours	<b>17.749</b>	<b>0.840</b>	<b>0.169</b>	<b>51.3</b>	<b>17.0</b>	47.2

prehension of occlusion relationships can be divided into two parts. The first is the ability to synthesize parts that were occluded in the input view. The second is the ability to synthesize new occlusion relationships under the novel view. We show several examples of synthesizing occluded parts and synthesizing new occlusions in Fig. S.13. We believe this capability is learned in a data-driven way since the

multi-object composites are physically plausible regarding these occlusion relationships.

Secondly, we now propose a new metric to evaluate the capability of understanding mutual occlusion under this setting. We first use visible mask and amodal mask in the input-view image to determine how heavily an object is occluded:



Figure S.10. **Continuous rotation examples on SUNRGB-D and 3D-FRONT.** We rotate the camera around the multi-object composites, successfully synthesizing plausible novel-view images across a wide range of camera pose variations. This first five examples are from SUNRGB-D, and the last three examples are from 3D-FRONT.




Figure S.11. **Visualized cross-view matching results.** Since we do not have ground truth image for 3D-FRONT and SUNRGB-D, we only visualize cross-view matching results using our predicted images. But we can still observe a strong cross-view consistency from the accurate matching results.

1. If an object’s visible mask is exactly its full mask, there exists no occlusion.
2. If an object’s visible mask is more than 70% of its full mask, the object is occluded.
3. If an object’s visible mask is less than 70% of its full mask, the object is heavily occluded.

Afterward, we segment the predicted view image with ground truth per-object visible mask. We calculate the specific region’s PSNR, SSIM, and LPIPS metrics as shown in Tab. S.6. It can reflect how well our model and baseline models are at synthesizing novel view plausible images that are originally occluded under the input view. There are 10903 fully visible objects, 6058 occluded objects, and 2215 heavily occluded objects. This experiment is conducted on our own C3DFS.




Figure S.12. **Failure Cases.** It is hard for our model to learn extremely fine-grained consistency on objects with delicate structure and texture.



Figure S.13. **Occlusion Synthesis Capability.** Our proposed method can synthesize new occlusion relationship under novel views as shown in the highlighted area of sofa or cabinet in (a). Our method can also hallucinate occluded parts as shown in the highlighted area of chairs in (b).

### C. Failure Cases and Limitations

**Failure Cases** We showcase two failure cases in Fig. S.12. We can observe that delicate structure and texture like colorful pillows on the sofa or slim legs of chairs are hard for our model to learn. Though object placement is approximately accurate, more fine-grained consistency is not quite ideal in these cases. We believe training with a higher resolution and incorporating epipolar constraints will mitigate this problem in the future.

**Limitations** We identify two limitations of our work. Firstly, though we achieve stronger cross-view consistency with the input view image, our model does not guarantee the multi-view consistency between our synthesized images. It is plausible to synthesize any results in areas with ambiguity, leading to potential multi-view inconsistency in our model. We believe incorporating multi-view awareness techniques [34, 47, 68, 69, 79, 90] can mitigate this problem. Secondly, we do not model background texture in our framework due to difficulty of realistically mimicking real-

Table S.6. **Evaluation on objects with varying extents of occlusion.**

Method	Visible			Occluded			Heavily Occluded		
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
Ours	<b>11.45</b>	<b>0.56</b>	<b>0.13</b>	<b>11.33</b>	<b>0.55</b>	<b>0.14</b>	<b>10.57</b>	<b>0.55</b>	<b>0.14</b>
Zero-1-to-3	9.46	0.54	0.16	9.33	0.52	0.17	9.00	0.53	0.16
Zero-1-to-3 <sup>†</sup>	9.68	0.55	0.14	9.54	0.52	0.15	9.26	0.53	0.15




Figure S.14. **Object Removal Example.** We can remove an object under novel views by setting a threshold to the predicted mask image and delete corresponding pixels.




Figure S.15. **Reconstruction results using DUST3R.** We rotate our camera around the multi-object composite and use the predicted images along with the input-view image for reconstruction.

world background texture, making it less convenient to directly apply our method to in-the-wild images. We believe training on more realistic data with background in the future can make our model more convenient to use.

## D. Potential Negative Impact

The use of diffusion models to generate compositional assets can raise ethical concerns, especially if used to create realistic yet fake environments. This could be exploited for

misinformation or deceptive purposes, potentially leading to trust issues and societal harm. Additionally, hallucinations from diffusion generation models can produce misleading or false information within generated images. This is particularly concerning in applications where accuracy and fidelity to the real world are critical.




Figure S.16. More visualized results on C3DFS dataset.




Figure S.17. More visualized results on Objaverse dataset.




Figure S.18. More visualized results on Room-Texture dataset.