

# Learning Unified Decompositional and Compositional NeRF for Editable Novel View Synthesis

Yuxin Wang<sup>1</sup>Wayne Wu<sup>2</sup>Dan Xu<sup>1\*</sup><sup>1</sup>Hong Kong University of Science and Technology <sup>2</sup>Shanghai Artificial Intelligence Laboratory

ywangom@connect.ust.hk, wuwenyan0503@gmail.com, danxu@cse.ust.hk

## Abstract

Implicit neural representations have shown powerful capacity in modeling real-world 3D scenes, offering superior performance in novel view synthesis. In this paper, we target a more challenging scenario, i.e., joint scene novel view synthesis and editing based on implicit neural scene representations. State-of-the-art methods in this direction typically consider building separate networks for these two tasks (i.e., view synthesis and editing). Thus, the modeling of interactions and correlations between these two tasks is very limited, which, however, is critical for learning high-quality scene representations. To tackle this problem, in this paper, we propose a unified Neural Radiance Field (NeRF) framework to effectively perform joint scene decomposition and composition for modeling real-world scenes. The decomposition aims at learning disentangled 3D representations of different objects and the background, allowing for scene editing, while scene composition models an entire scene representation for novel view synthesis. Specifically, with a two-stage NeRF framework, we learn a coarse stage for predicting a global radiance field as guidance for point sampling, and in the second fine-grained stage, we perform scene decomposition by a novel one-hot object radiance field regularization module and a pseudo supervision via inpainting to handle ambiguous background regions occluded by objects. The decomposed object-level radiance fields are further composed by using activations from the decomposition module. Extensive quantitative and qualitative results show the effectiveness of our method for scene decomposition and composition, outperforming state-of-the-art methods for both novel-view synthesis and editing tasks<sup>1</sup>.

## 1. Introduction

The reconstruction and rendering of natural scenes are important for computers to understand the 3D physical

\*Corresponding author

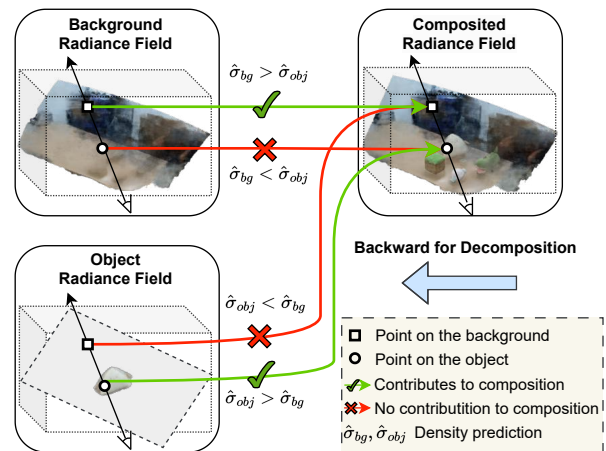
<sup>1</sup>Project: <https://w-ted.github.io/publications/udc-nerf>

Figure 1. Illustration of our decomposition-composition design. It enables scene editing and novel view synthesis in a unified NeRF framework. The proposed decomposition targets learning disentangled 3D representations of different foreground objects and the background, allowing for scene editing, while scene composition constructs an entire scene representation for novel view synthesis.

world. Fine-grained object-level representations within the scene can bring significant benefits in various applications, such as scene understanding, novel content generation via object editing, and robotic manipulation. The emerging neural rendering techniques [17, 18, 24] allow learning object-level or scene-level representations from multi-view posed images and enable rendering high-quality images from novel views. However, most implicit neural methods model an entire scene and lack fine-grained representations of objects in the scene, which severely limits object-level scene representation and understanding.

Learning effective 3D object-aware implicit neural scene representations is still in its infancy. Existing works that tackle this challenging problem typically utilize additional object-level semantic supervision as a priori knowledge in the model optimization [35, 37, 42]. For instance, ObjectSDF [35] considers learning object-level geometries and semantics with an implicit signed distance function (SDF), which achieves effective scene composition from object-

level representations. However, it remains challenging for this framework to perform joint object editing and novel view synthesis. On the other hand, ObjectNeRF [37] explores scene modeling with Neural Radiance Fields (NeRF) via introducing learnable object activation codes as switches to condition the radiance field prediction among different objects. This framework can perform joint scene editing and novel view synthesis benefiting from the radiance field representations. However, it learns two separate networks for these two tasks, a global scene branch for novel view synthesis, and a local object branch for scene editing, without modeling any interaction or correlation among them, which, however, is critical for learning more effective scene representations, due to the simple fact that, global scene and local object representations are two important perspectives of the same scene. The global entire scene representation can model the scene’s overall structures and appearance consistency, and the local object-specific representation can learn more fine-grained object details. They are highly complementary for learning a high-quality scene representation.

To target the above-mentioned problem, this paper proposes a unified decompositional and compositional neural radiance field framework (see Fig. 1), to learn 3D scene representation for joint scene editing and novel view synthesis. The decomposition can provide the functionality of learning disentangled 3D representations of different objects and the background, allowing for scene editing, while scene composition models an entire scene representation for novel view synthesis. These two can be united to facilitate the consistency constraint in the unified optimization framework. Specifically, we design a new two-stage framework for the scene decomposition and composition, consisting a coarse and a fine stage. In the coarse stage, we learn to predict a global scene radiance field as guidance for point sampling, and in the second fine-grained stage, we perform joint scene decomposition and composition. To perform effective scene decomposition, in the fine-stage, we apply a set of learnable object codes to predict distinct object-level radiance fields, and also propose two novel decomposition strategies: **(i)** 3D one-hot object radiance activation and regularization. For the object or background branches, only one branch is activated during training using a designed Gumbel-Softmax activation function directly applied on the point density predicted from different object/background branches with also a corresponding one-hot regularization; **(ii)** an in-painting pseudo supervision strategy. A challenge in scene rendering is modeling the appearance and geometry of regions occluded by objects. This causes generation ambiguity especially when the occluded regions are unseen in all the training views. To address this issue, we propose to use a pre-trained inpainting model to provide additional pseudo-color supervision for those ambiguous areas. Note that although the 2D inpainting may bring new ambiguity for the regions

seen in other views, the supervision from multi-view consistency is strong enough to suppress most ambiguities. The composition is further performed by utilizing the learned one-hot activation weights for different object-level radiance fields.

In summary, this paper has the following contributions:

- We propose a novel NeRF framework for joint scene decomposition and composition to effectively learn object-level and scene-level implicit representations for effective scene modeling, allowing object editing and novel view synthesis in a unified pipeline.
- To learn a robust scene decomposition, we design two novel strategies, *i.e.*, 3D one-hot radiance regularization and 2D in-painting pseudo supervision, to significantly improve the rendering and editing qualities.
- Extensive experiments demonstrate the effectiveness of our approach for scene decomposition and show clear improvements upon state-of-the-art object-compositional methods on two important downstream tasks, *i.e.* novel-view synthesis and object editing.

## 2. Related Work

**Neural Implicit Representations.** As pioneering methods, [17, 21, 24, 28] introduce neural networks to implicitly parameterize 3D scenes. These works can learn continuous shape and texture from differentiable rendering with [12, 24] or without [21, 28] 3D supervision. Neural Radiance Fields (NeRF) [18] optimize a 3D continuous neural scene representation by using only 2D multi-view posed images for training and produce high-quality novel view synthesis with 3D structure awareness. Recent works also extend the generalization of NeRF to more significant scenes [1, 2, 41], or model scenes with fewer perspectives [5, 33, 39]. However, all these works only focus on the entire scene representation, neglecting fine-grained representations of objects/background within the scene. Other researches represent object surfaces with implicit neural function to realize multi-view 3D reconstruction for single objects [22, 32, 38]. The decomposed implicit representation for independently characterizing components within the scene has been studied in not only object scenes [3, 14, 8, 23, 25, 27], but also in the face or body tasks [9, 11, 34]. In contrast to these works, our method targets both scene synthesis and editing, considering joint object- and scene-level implicit representation learning.

**Object-Compositional Implicit Representation.** Representing scenes in an object-compositional manner can benefit both scene synthesis and editing [3, 4, 26, 29, 30, 35, 40, 36]. Compared to decomposed object representations represented by generative models [7, 19, 20], implicit models present more 3D physical meaning during the modeling process, which facilitates image rendering from novel

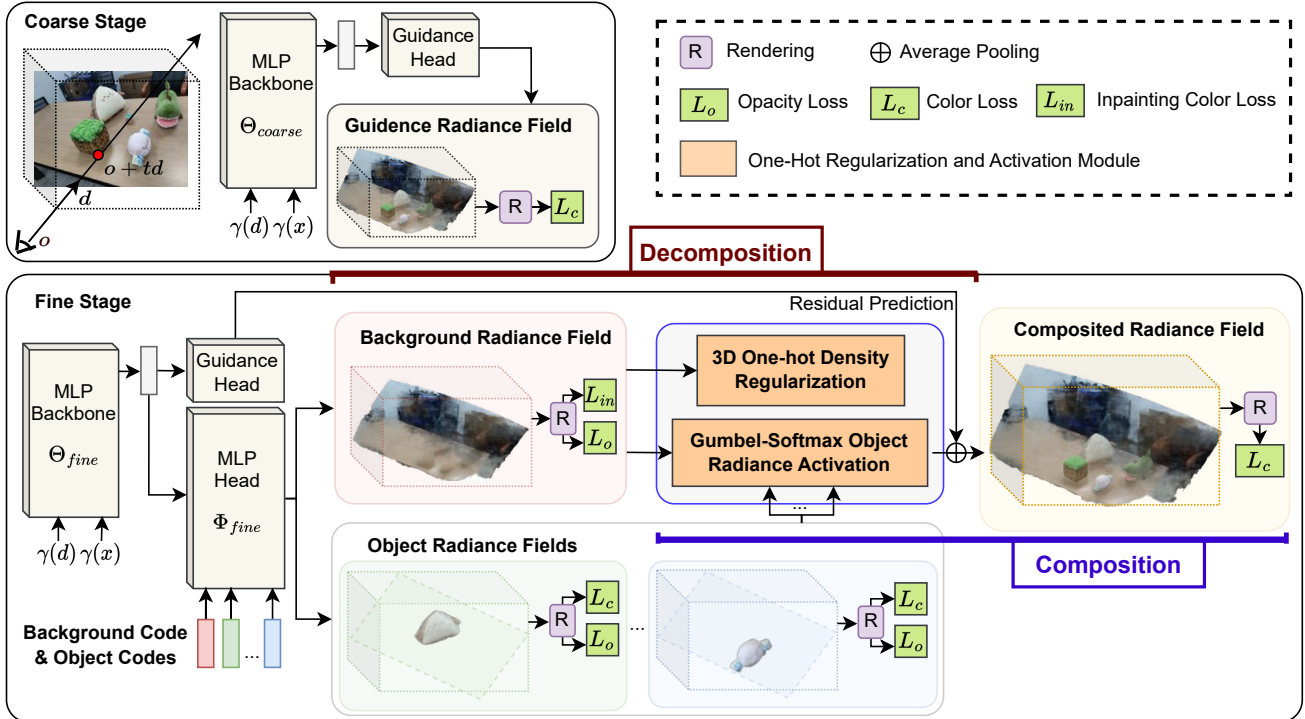


Figure 2. An overview of the proposed unified decompositional and compositional NeRF framework for joint novel view synthesis and scene editing. It has two stages. In the first stage (*i.e.* the coarse stage), it learns a guidance radiance field for guiding point sampling. In the second stage (*i.e.* the fine stage), we learn scene decomposition via learnable object codes and two novel decomposition schemes: (i) the 3D one-shot object radiance activation regularization and (ii) color inpainting handling ambiguous generation in occluded background areas. The scene composition is achieved by using one-hot activation weights for different object-level radiance fields learned in the decomposition stage. The decomposition allows scene editing and the composition enables novel view synthesis in the unified framework.

views. We can categorize the works of object-compositional implicit representations into supervised and unsupervised methods. The unsupervised methods localize objects with different aspects of prior information, such as multiple similar scenes [26, 29, 40], temporal consistency [13], or clustering-based initialization [15]. Specifically, uORF [40] involves a slot-attention [13] module in learning object-level iterative attention masks for each object decomposition. RFP [15] utilizes clustering predictions as initialization and applies a bidirectional color loss to decompose individual objects. However, lacking precise annotation, although these methods can obtain object position clues, they usually suffer from low-resolution [26, 40] issues that restrict generalization ability to high-resolution real-world scenes. The supervised methods usually use object-level annotations, such as semantic masks [35, 37, 42] or scene graphs [23]. Among these works, ObjectNeRF [37] utilizes a two-branch network to learn a scene radiance field with an activation code-conditioned object radiance field supervised by multi-view images and semantic masks. Similarly, ObjectSDF [35] encodes the entire scene by combining the individual Signed Distance Functions (SDF) represented with semantic labels. The most related work to us is ObjectNeRF [37]. In contrast to ObjectNeRF, we model

an the scene representation via joint scene decomposition and composition, which can effectively model global consistency and rich local details in a unified NeRF framework for a high-quality scene representation.

### 3. Decompositional and Compositional NeRF

Given a set of multi-view posed images  $\{x_i, i \in [1, N]\}$  and their corresponding instance semantic masks  $\{s_{ij}, i \in [1, N], j \in [1, K + 1]\}$  of a static real-world scene, our goal is to learn an entire scene representation based on the composition of a global scene radiance field and explicitly decomposed local radiance fields of different objects and the background. The decomposition allows scene editing and the composition enables novel view synthesis.  $N$  is the number of multi-view images, and  $K$  is the number of object instances, where we regard the background as an additional instance. We propose a joint decomposition and composition framework to allow simultaneous performing these two tasks.

#### 3.1. Framework Overview

An overall framework illustration is shown in Fig. 2. It consists of two rendering stages. In the first stage (*i.e.* the

coarse stage), it learns a global scene guidance radiance field for guiding point sampling. In the second stage (*i.e.* the fine stage), we learn scene decomposition and composition. The decomposition is performed via first applying a set of learnable object codes as conditions to predict  $K + 1$  local object/background radiance field components, where  $K$  is the number of objects, and the additional 1 represents the background. Following NeRF [18], we use continuous neural radiance fields to represent each local object/background. Considering the composition nature that the whole scene consists of all the  $K + 1$  components, we composite all the object and background radiance fields to an entire scene radiance field based on a proposed one-hot object radiance activation module. We also propose an effective color-inpainting strategy to handle the ambiguity in generating occluded background regions. The residual prediction from the first coarse stage is also utilized to benefit from a global scene representation. In the following subsections, we first give a general introduction to the proposed unified framework. Then, we introduce our novel decomposition and composition design and how the method can help disentangle different objects explicitly. After that, we present how our unified framework is optimized and utilized for both novel-view synthesis and editing tasks.

### 3.2. Scene Decomposition

**Two-branch network architecture design.** For scene decomposition in the fine-stage, we design a two-branch network architecture with a shared MLP backbone to learn generic implicit scene representation from input point coordinates (see Fig. 3). Specifically, the shared backbone takes a sampled 3D point coordinate  $\mathbf{p} = (x, y, z)$  and view direction  $\mathbf{d}$  of a camera ray as input, and predicts a generic scene feature vector  $\mathbf{f}$ . Similar to [37], we also apply a set of learnable codes, each representing the background or an object. We further pass  $\mathbf{f}$  accompanying with background or object codes to the following MLP heads. For better background learning, we use separate the background head and the object head, predicting an RGB color  $\mathbf{c}_{\text{object}}^i$  and a volume density  $\sigma_{\text{object}}^i$  for the corresponding background or object  $i$ . For simplicity, we refer to the background as a particular object with index  $i = 1$ . For efficient point sampling, we apply a coarse-to-fine sampling technique. In the coarse stage, we only adopt several MLPs after the backbone for predicting RGB color  $\mathbf{c}_{\text{guidance}}$  and volume density  $\sigma_{\text{guidance}}$  of the global scene. We use the distribution of coarse prediction for the fine point coordinate sampling, so we call these MLPs guidance head. In the fine stage, we also use the prediction of the guidance head as a residual radiance field in the following composition module.

**3D one-hot density regularization.** To achieve effective decomposition of objects/background radiance fields from the scene, we also design a one-hot regularization term for

the uncertain regions in the background radiance field. Our motivation is that the 3D points belonging to the foreground objects should not have any volume density in the background radiance field. We thus apply a one-hot regularization loss term, which acts as an additional constraint to the object-specific radiance field optimization objective  $\mathcal{L}_{\text{object}}$ . In detail, we first obtain the 3D point positions that foreground objects occupy by determining whether the foreground objects take the largest volume density at the points. We thus define a binary occupancy mask  $M_c$ , and apply a regularization term for the occupied points as follows:

$$M_c = [\text{argmax}([\sigma_{\text{object}}^1, \sigma_{\text{object}}^2, \sigma_{\text{object}}^3, \dots, \sigma_{\text{object}}^{K+1}]) > 1] \quad (1)$$

$$\mathcal{L}_{\text{one-hot}} = M_c \|\sigma_{\text{object}}^1 - \sigma_0\|_2^2$$

$\sigma_{\text{object}}^1$  here means we only want to constrain the background radiance field (with index 1). Note that we implement a ReLU operation on the volume density  $\sigma$  to make it positive during the volume rendering process, so if the raw prediction  $\sigma_{\text{object}}^1$  is negative at some points, it will also have no contribution to the rendering results. Therefore, we can set  $\sigma_0$  as any non-positive number here. We put it -0.01 in our experiments from experience.

**2D in-painting pseudo supervision.** To model a better background-specific radiance field, we need to learn what the scene looks like even if all the foreground objects are removed from the scene. Several occluded regions in the background can be seen in other views, while regions covered by the objects are invisible from all views, which are uncertain regions. We thus do not have any pixel-wise supervision for those regions in the training data, resulting in unsatisfactory extrapolation results on those regions when we learn the background radiance field. To tackle this problem, we involve a pre-trained in-painting model to fill those uncertain regions on the 2D images. Our motivation is that the 2D in-painting model can utilize local context information on an image to fill the regions, providing effective 2D pseudo color supervision. Specifically, we use a famous pre-trained lama [31] model to fill the uncertain 2D background regions and then apply a color loss for the rendered background pixels with a foreground mask  $M_1(\mathbf{r})^1$ :

$$\mathcal{L}_{\text{color-pseudo}} = M_1(\mathbf{r})^1 \|\hat{C}(\mathbf{r})_{\text{object}}^1 - C(\mathbf{r})_{\text{pseudo}}\|_2^2 \quad (2)$$

**Intermediate local object rendering.** To enforce the decomposition of all object-specific radiance fields, we apply a color consistent loss and an opacity loss with semantic masks as [18, 37] via rendering both 2D color and opacity:

$$\hat{C}(\mathbf{r})_{\text{object}}^k = \sum_{i=1}^N T_i^k (1 - \exp(-\sigma_{\text{object}}^k \delta_i)) \mathbf{c}_{\text{object}}^k \quad (3)$$

$$\hat{O}(\mathbf{r})_{\text{object}}^k = \sum_{i=1}^N T_i^k (1 - \exp(-\sigma_{\text{object}}^k \delta_i))$$

where  $T_i^k = \exp(-\sum_{j=1}^{i-1} \sigma_{\text{object}}^k \delta_j)$  is the transmittance representing how much light is transmitted along the ray for the  $k$ -th object;  $N$  is the number of sampled points along the camera ray, and  $\delta_i$  is the distance between two adjacent points. Then we apply the following losses:

$$\mathcal{L}_{\text{object}} = \sum_{k=1}^{K+1} M_o(\mathbf{r})^k \|\hat{C}(\mathbf{r})_{\text{object}}^k - C(\mathbf{r})\|_2^2 + \lambda M_b(\mathbf{r})^1 \|\hat{O}(\mathbf{r})_{\text{object}}^1 - M_o(\mathbf{r})^1\|_2^2, \quad (4)$$

where  $M_o(\mathbf{r})$  and  $M_b(\mathbf{r})$  are 2D spatial weight maps determined by object-level and background-level semantic masks, respectively, which indicates that rendered object-specific pixels are given supervision at the corresponding semantic mask region. The background region in the binary semantic mask can be either real background or occluded by a foreground object, so we use a small weight value (*i.e.* 0.05) in  $M_b$  as a balance weight for the uncertain region.

### 3.3. Scene Composition

This section presents how the learned global and the local object/background radiance fields are composited for an entire scene radiance field.

**Gumbel-Softmax one-hot object radiance activation.** As shown in Fig. 3, for each sampled point, there are  $K + 1$  predictions in the local radiance field decomposition. Considering that each point in the 3D space can only belong to either one object or the background, we use the predicted  $(\mathbf{c}_{\text{object}}^i, \sigma_{\text{object}}^i)$  with the maximum volume density  $\sigma_{\text{object}}^i$  among all the object instances to construct an activated object-specific radiance field. More specifically, to address the gradient backpropagation problem of the argmax operation, following [10, 16], We employ the Gumbel-Softmax function on volume density as a continuous and differentiable approximation. First, we calculate a soft probability  $p_i$  for the  $i$ -th object branch as

$$p_i = \frac{\exp((\sigma_{\text{object}}^i + g_i)/\text{tau})}{\sum_j \exp((\sigma_{\text{object}}^j + g_j)/\text{tau})} \quad (5)$$

The above  $g_1, g_2, \dots, g_j$  are the independent and identically distributed samples drawn from  $\text{Gumbel}(0, 1)^1$  [10]. The symbol *tau* is the temperature ratio set to 0.1 in our experiments. Then, we can obtain one-hot composition weights from  $p_i$  for combining the predictions of all the object/background branches, which ensures one-hot activation for all the object/background radiance fields, and enables effective backpropagation for learning object-specific radiance fields and their composition. Specifically, we calculate the one-hot weights  $\mathbf{M}_{\text{obj}}$  as Eq. 6. Let us denote  $\mathbf{P}$  as the soft probability vector, containing  $K + 1$  predictions calculated by Eq. 5.  $\mathbf{M}_{\text{obj}}$  is a one-hot vector of the  $K + 1$  elements. The one-hot vector is calculated with a stop gradient operation  $f_{\text{sp}}$  on  $\mathbf{P}$ , which not only reserves the one-

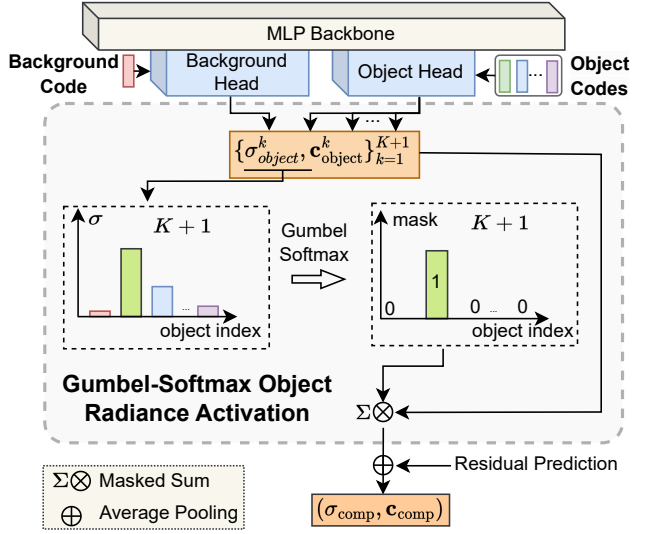


Figure 3. Illustration of the Composition Module. Our composition module consists of one-hot object radiance activation and global-local radiance field composition (the two green blocks). We first apply Gumbel-Softmax on object-level volume density values to obtain the one-hot mask for object radiance fields. Then, the local compositional radiance field is generated by a weighted sum of the object and background color and density predictions. Finally, the composited radiance is combined with a residual prediction to produce the final scene radiance field.

hot activation but also allows the gradient backpropagated through  $\mathbf{P}$ . Then, the computation of  $\mathbf{M}_{\text{obj}}$  is written as:

$$\mathbf{M}_{\text{obj}} = f_{\text{h}}(\arg \max_i [g_i + \log \sigma_{\text{object}}^i]) - f_{\text{sp}}(\mathbf{P}) + \mathbf{P}, \quad (6)$$

where  $f_{\text{h}}(\cdot)$  denotes a one-hot function; Then, we obtain a composed scene radiance field  $(\mathbf{c}_{\text{comp}}, \sigma_{\text{comp}})$  of all decomposed branches  $\{\mathbf{c}_{\text{object}}^k, \sigma_{\text{object}}^k\}_{k=1}^{K+1}$ , by multiplying the above one-hot weights as follows:

$$\mathbf{c}_{\text{comp}} = \sum_{K+1} \mathbf{M}_{\text{obj}}^k \mathbf{c}_{\text{object}}^k, \quad \sigma_{\text{comp}} = \sum_{K+1} \mathbf{M}_{\text{obj}}^k \sigma_{\text{object}}^k, \quad (7)$$

Our novel one-hot object radiance activation can not only ensure a reasonable composition of all the object radiance fields but also force that only one object volume density prediction can be activated in the gradient backpropagation, which can also effectively help the object/background decomposition.

**Residual radiance field composition.** After the one-hot object radiance activation, we further combine the composited radiance fields from objects/background branches with the predicted residual radiance field, to produce a final composited scene radiance field, using point-wise average pooling operation applied on the color and density. Then we apply a color loss  $\mathcal{L}_{\text{comp}}$  for the pixels rendered from the composited radiance field.

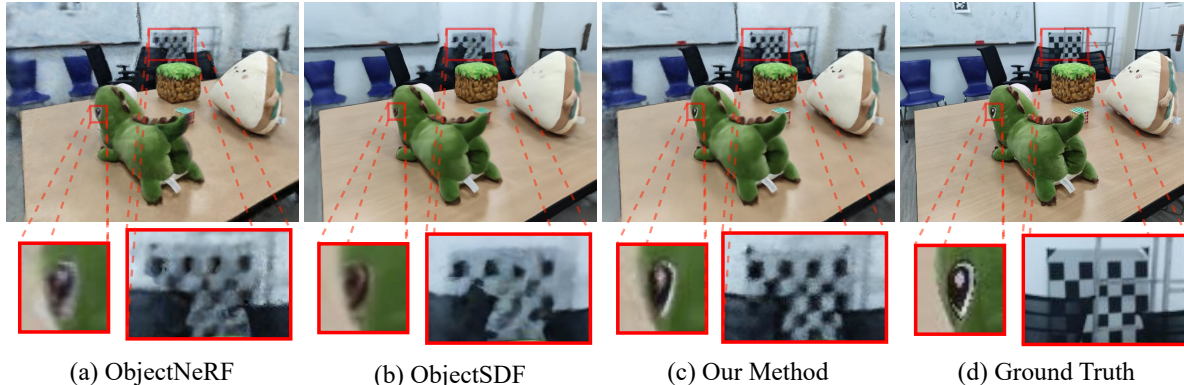


Figure 4. Qualitative comparison of novel view synthesis on the ToyDesk dataset. We compared with ObjectNeRF [37] and ObjectSDF [35]. The results show that the rendered images of our framework achieve better generation quality for both the foreground objects (e.g., the green dinosaur) and the background (e.g., the chessboard).

Methods	ToyDesk2		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Our Full-Model	<b>25.7560</b>	<b>0.8126</b>	<b>0.4488</b>
w/o 3D One-hot Density Regularization	25.7362	0.8116	0.4524
w/o Gumbel-Softmax Activation	22.9972	0.7217	0.5452
w/o 2D In-painting pseudo Supervision	25.6881	0.8127	0.4516
w/o Residual Radiance Composition	25.2288	0.8003	0.4589

Table 1. Quantitative comparison of different variants of the proposed unified decompositional and compositional NeRF framework on the ToyDesk2.

### 3.4. Optimization and Inference

**Overall framework optimization.** The overall optimization loss for training the joint framework consists of the following four different loss functions:

$$\mathcal{L}_{\text{overall}} = \mathcal{L}_{\text{comp}} + \lambda_1 \mathcal{L}_{\text{object}} + \lambda_2 \mathcal{L}_{\text{color-pseudo}} + \lambda_3 \mathcal{L}_{\text{one-hot}}, \quad (8)$$

where,  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  are the weights for balancing the learning with the different loss functions.

**Novel View Rendering and Scene Editing** After training the joint framework, we can render images of arbitrary perspectives based on the composited radiance field. By re-organizing the object radiance fields, we can generate images with object manipulations, including removal, addition, duplication, and changing position (e.g., rotation).

## 4. Experiments

We aim at performing joint scene novel view synthesis and editing based on our proposed unified decompositional and compositional NeRF framework. We thus evaluate our method on both tasks using challenging public benchmarks.

### 4.1. Experimental Setup

**Datasets.** We use ToyDesk [37] and ScanNet [6] datasets for the evaluation. **ToyDesk** consists of two sets of indoor multi-view images. In each scene, several toys are put on a desk in the center of the room. The multi-view images are captured by looking at the desk center with 360 degrees




Figure 5. Qualitative results of the ablation study: (a) Our full model; (b) w/o 3D One-hot Density Regularization; (c) w/o Gumbel-Softmax Activation; (d) w/o 2D In-painting pseudo Supervision. (e) w/o Residual Radiance Composition.

around the desk, and the camera poses and meshes are provided. The instance semantic labels are annotated on object meshes and then projected onto each perspective to obtain their corresponding 2D semantic label maps. **ScanNet** is an RGB-D video dataset that contains more than 1500 scans of indoor scenes for multiple 3D understanding tasks. Hand-held RGB-D scanners are used for scanning various indoor scenes. Rich annotations such as camera pose, semantic segmentation, and surface reconstruction are also provided. Following ObjectNeRF [37], we choose four scenes, i.e., 0024\_00, 0038\_00, 0113\_00, and 0192\_00, from the whole dataset for the experiment.

**Implementation details.** Our joint framework consists of a shared backbone, background head, object head, and guidance head. The shared backbone is a 4-layer Multi-Layer Perceptron (MLP), and the output generic implicit scene feature has a dimension of 256. The follow-up heads are all three-layer MLPs with a color and density prediction layer to generate corresponding radiance fields. Note that our one-hot regularization and activation module has no parameters to be optimized. We also utilize positional encoding to increase the representation ability for high-frequency details. Specifically, we adopt ten levels for coordinates and four levels for ray directions. The framework is trained end-to-end from scratch, without voxel feature initialization as in [37] or geometric initialization as in [35].




Figure 6. Qualitative comparison of novel view synthesis on the ScanNet dataset. Compared with ObjectNeRF [37] and ObjectSDF [35], our results also achieve better fidelity. We highlight the details in the red boxes.




Figure 7. The Decomposition results on the ToyDesk dataset compared with ObjectNeRF [37]. We show the rendered results of our composited radiance field (the first column) and each background/object radiance field (the second and third columns).

## 4.2. Model Analysis

**Baseline Models.** To verify the effectiveness of our designs for scene decomposition and composition, We conduct several ablation studies considering different variants: (i) ‘w/o 3D One-hot Density Regularization’ means that we remove the one-hot regularization term for the uncertain regions in the background radiance field. (ii) ‘w/o Gumbel-Softmax Activation’ indicates that we do not use the Gumbel-Softmax one-hot weight for compositing the object radiance fields. (iii) ‘w/o In-painting Supervision’ means that we do not apply the in-painting pseudo supervision for the background radiance field. (iv) ‘w/o Residual Radiance Composition’ means that we do not use the guidance radiance field for fine sampling and remove the residual radiance composition in the fine stage.

**Effect of 3D one-hot density regularization.** The quantitative results based on different evaluation metrics, *i.e.*, PSNR, SSIM, and LPIPS, are shown in Table 1. Several

Dataset	Methods	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
ToyDesk	Object-NeRF [37]	22.1736	0.7723	0.4711
	Object-SDF [35]	22.0298	-	-
	Our Method	<b>23.1435</b>	<b>0.7899</b>	<b>0.4691</b>
ScanNet	Object-NeRF [37]	25.2640	0.8047	0.4094
	Object-SDF [35]	25.2325	-	-
	Our Method	<b>26.1349</b>	<b>0.8249</b>	<b>0.3951</b>

Table 2. Quantitative comparison on novel view synthesis. We compared with ObjectNeRF [37] and ObjectSDF [35]. ‘-’ indicates the metrics are not reported by the authors in the paper. The evaluation results show that our synthesis results quantitatively outperform these two SOTA methods on all the shown metrics.

examples of rendered images from the learned background radiance field are presented in Fig. 5. We evaluate the effectiveness of 3D one-hot density regularization by disabling the loss term. As can be observed from Table 1, our joint full mode with 3D one-hot density regularization achieves better results than the model without it.

### Effect of Gumbel-Softmax object radiance activation.

We also investigate the effect of using Gumbel-Softmax one-hot activation for the composition of background/object radiance fields. We use an alternative composition method, in which all the local object/background radiance fields that contribute to the same 3D points in the composited field, are directly added instead of using the Gumbel-Softmax one-hot activation weights, as introduced in the method part. As shown in Fig. 5c, without the Gumbel-Softmax one-hot activation, the model cannot successfully learn a decomposed background radiance field (all black), confirming our motivation of using the Gumbel-Softmax activation for a more effective scene representation in a unified framework.

### Effect of 2D in-painting pseudo supervision and residual radiance composition.

We also show the contribution of the proposed 2D in-painting pseudo-supervision. If we directly disable the pseudo supervision in the framework, the qualitative results are slightly lower, as shown in the




Figure 8. The Decomposition results on the ScanNet dataset compared with ObjectNeRF [37]. Our background has fewer artifacts when removing the foreground object.

next to last row in Table. 1. At the same time, as can be seen in Fig. 5, lacking pseudo supervision suffers from significantly downgraded extrapolation in the region occupied by the foreground objects in the 3D space, which indicates that the proposed scheme can effectively help the local object/background disentanglement, especially for learning the radiance field representation of the background on uncertain regions. As shown in Table. 1 and Fig. 5, if we remove the residual radiance composition, the overall qualitative metrics and background decomposition both decline, which demonstrates the importance of using a global radiation field to provide coarse sampling guidance and also further merged as residual information in the fine stage.

### 4.3. Results on Novel View Synthesis

We compare our methods on the task of novel view synthesis with two state-of-the-art methods, *i.e.*, ObjectNeRF [37] and ObjectSDF [35]. Our train-test split follows ObjectNeRF [37]. We measure the quality of synthesized images on the evaluation metrics of PSNR, SSIM, and LPIPS. In detail, for our joint framework, we use the finally composited entire scene radiance field for the image rendering. For ObjectNeRF and ObjectSDF on the Toydesk dataset, we use the rendered images provided by the authors for comparison. ObjectSDF adopts a different train-test split and label mapping strategy on ScanNet, so we re-train the ObjectSDF on ScanNet following our setting.

In Table 2, our synthesis results quantitatively outperform the comparison methods on all the shown metrics. We also present the image examples rendered from novel views in Fig. 4 and Fig. 6. As can be observed in Fig. 4, the rendered images of our framework achieve better generation quality for both the foreground objects (*e.g.*, the green dinosaur) and the background (*e.g.*, the chessboard) than both ObjectNeRF [37] and ObjectSDF [35], proving the capability of our method in modeling both global scene consistency and local scene details. Fig. 6 shows the synthesized images



(a) Original view (b) Edited image

Figure 9. Illustration of the editing results. (a) shows the image in the original view, and (b) shows the editing result from the same perspective. We do the rotation and duplication for the white triangle and candy by re-compositing the learned radiance fields.

on ScanNet. We can also observe that our results outperform others in the overall fidelity and local details marked in the red boxes.

### 4.4. Results of Object Decomposition and Editing

We also conduct experiments to show the decomposition effectiveness of our method, especially for the background radiance field. As shown in Fig. 7 and Fig. 8, we rendered the composited radiance field and each local object/background radiance field to show the decomposition results. The background images in Fig. 7 clearly show that our method can generate a clearer table without unsatisfactory artifacts such as strange colors and black shadow effects when removing all the foreground objects on the table. Although with some remained black contours, the rendered pixel colors of the regions below the objects are significantly more reasonable compared to ObjectNeRF [37]. Note that the region under the toys is unseen in any view. Our decomposition module enables clear background and object-specific details, while the composition design helps the rendered entire scene achieve higher quality than the baseline. In Fig. 8, we further show that our pipeline can also provide better background rendering when removing foreground objects on the images from the ScanNet dataset. The better decomposition naturally benefits downstream object editing tasks. Fig. 9 shows that, after learning composited radiance fields, our method can perform rotation and duplication for the white triangle and candy by re-compositing the object and background radiance field representations. We add more editing results in the supplementary materials.

## 5. Conclusion

We propose a unified NeRF framework with a novel decomposition and composition design to jointly model real-world scenes with both global and local implicit representations. To solve the ambiguous extrapolation problem for unseen areas, we also propose a 3D one-hot density regularization strategy and a 2D in-painting pseudo supervision. Our experimental evaluation demonstrates that our method outperforms the state-of-the-art object-compositional methods for both novel-view synthesis and scene editing tasks.



## 6. Acknowledgments

This research is supported in part by HKUST-SAIL joint research funding, the Early Career Scheme of the Research Grants Council (RGC) of the Hong Kong SAR under grant No. 26202321, and HKUST Startup Fund No. R9253.

## References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, 2021. 2
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, 2022. 2
- [3] Sagie Benaim, Frederik Warburg, Peter Ebert Christensen, and Serge Belongie. Volumetric disentanglement for 3d scene manipulation. *arXiv*, 2022. 2
- [4] Boming Zhao and Bangbang Yang, Zhenyang Li, Zuoyue Li, Guofeng Zhang, Jiashu Zhao, Dawei Yin, Zhaopeng Cui, and Hujun Bao. Factorized and controllable neural re-rendering of outdoor scene for photo extrapolation. In *ACM Multimedia*, 2022. 2
- [5] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo radiance fields (srf): Learning view synthesis from sparse views of novel scenes. In *CVPR*, 2021. 2
- [6] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. 6
- [7] Dave Epstein, Taesung Park, Richard Zhang, Eli Shechtman, and Alexei A. Efros. Blobgan: Spatially disentangled scene representations. *ECCV*, 2022. 2
- [8] Michelle Guo, Alireza Fathi, Jiajun Wu, and Thomas Funkhouser. Object-centric neural scene rendering. *arXiv preprint arXiv:2012.08503*, 2020. 2
- [9] Yudong Guo, Keyu Chen, Sen Liang, Yongjin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *ICCV*, 2021. 2
- [10] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017. 5
- [11] Zhang Jiakai, Liu Xinhang, Ye Xinyi, Zhao Fuqiang, Zhang Yanshun, Wu Minye, Zhang Yingliang, Xu Lan, and Yu Jingyi. Editable free-viewpoint video using a layered neural representation. In *ACM SIGGRAPH*, 2021. 2
- [12] Chiyu Max Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, and Thomas Funkhouser. Local implicit grid representations for 3d scenes. In *CVPR*, 2020. 2
- [13] Thomas Kipf, Gamaleldin F. Elsayed, Aravindh Mahendran, Austin Stone, Sara Sabour, Georg Heigold, Rico Jonschkowski, Alexey Dosovitskiy, and Klaus Greff. Conditional Object-Centric Learning from Video. In *ICLR*, 2022. 3
- [14] Amit Pal Singh Kohli, Vincent Sitzmann, and Gordon Wetzstein. Semantic implicit neural scene representations with semi-supervised training. In *3DV*, 2020. 2
- [15] Xinhang Liu, Jiaben Chen, Huai Yu, Yu-Wing Tai, and Chi-Keung Tang. Unsupervised multi-view object segmentation using radiance field propagation. In *NeurIPS*, 2022. 3
- [16] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *ICLR*, 2017. 5
- [17] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019. 1, 2
- [18] 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 *ECCV*, 2020. 1, 2, 4, 11
- [19] Thu Nguyen-Phuoc, Christian Richardt, Long Mai, Yong-Liang Yang, and Niloy Mitra. Blockgan: Learning 3d object-aware scene representations from unlabelled images. In *NeurIPS*, 2020. 2
- [20] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *CVPR*, 2021. 2
- [21] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020. 2
- [22] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *ICCV*, 2021. 2
- [23] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. In *CVPR*, 2021. 2, 3
- [24] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 1, 2
- [25] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. In *CVPR*, 2021. 2
- [26] Mehdi S. M. Sajjadi, Daniel Duckworth, Aravindh Mahendran, Sjoerd van Steenkiste, Filip Pavetić, Mario Lučić, Leonidas J. Guibas, Klaus Greff, and Thomas Kipf. Object Scene Representation Transformer. *NeurIPS*, 2022. 2, 3
- [27] Mehdi S. M. Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lucic, Daniel Duckworth, Alexey Dosovitskiy, Jakob Uszkoreit, Thomas Funkhouser, and Andrea Tagliasacchi. Scene Representation Transformer: Geometry-Free Novel View Synthesis Through Set-Latent Scene Representations. In *CVPR*, 2022. 2
- [28] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *NeurIPS*, 2019. 2

- [29] Cameron Smith, Hong-Xing Yu, Sergey Zakharov, Fredo Durand, Joshua B Tenenbaum, Jiajun Wu, and Vincent Sitzmann. Unsupervised discovery and composition of object light fields. *arXiv preprint arXiv:2205.03923*, 2022. [2](#), [3](#)
- [30] Karl Stelzner, Kristian Kersting, and Adam R Kosiorek. Decomposing 3d scenes into objects via unsupervised volume segmentation. *arXiv preprint arXiv:2104.01148*, 2021. [2](#)
- [31] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. *arXiv preprint arXiv:2109.07161*, 2021. [4](#), [11](#)
- [32] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NeurIPS*, 2021. [2](#)
- [33] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. [2](#)
- [34] Ziyang Wang, Timur Bagautdinov, Stephen Lombardi, Tomas Simon, Jason Saragih, Jessica Hodgins, and Michael Zollhofer. Learning compositional radiance fields of dynamic human heads. In *CVPR*, 2021. [2](#)
- [35] Qianyi Wu, Xian Liu, Yuedong Chen, Kejie Li, Chuanxia Zheng, Jianfei Cai, and Jianmin Zheng. Object-compositional neural implicit surfaces. In *ECCV*, 2022. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#), [11](#), [13](#)
- [36] Qianyi Wu, Kaisiyuan Wang, Kejie Li, Jianmin Zheng, and Jianfei Cai. Objectsdf++: Improved object-compositional neural implicit surfaces. In *ICCV*, 2023. [2](#)
- [37] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *ICCV*, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#), [11](#), [13](#)
- [38] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *NeurIPS*, 2021. [2](#)
- [39] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. [2](#)
- [40] Hong-Xing Yu, Leonidas J. Guibas, and Jiajun Wu. Unsupervised discovery of object radiance fields. In *ICLR*, 2022. [2](#), [3](#)
- [41] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020. [2](#)
- [42] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J. Davison. In-place scene labelling and understanding with implicit scene representation. In *ICCV*, 2021. [1](#), [3](#)