

# Object-Compositional Neural Implicit Surfaces

Qianyi Wu<sup>1</sup>, Xian Liu<sup>2</sup>, Yuedong Chen<sup>1</sup>, Kejie Li<sup>3</sup>,  
Chuanxia Zheng<sup>1</sup>, Jianfei Cai<sup>1</sup>, and Jianmin Zheng<sup>4</sup>

<sup>1</sup> Monash University

<sup>2</sup> The Chinese University of Hong Kong

<sup>3</sup> University of Oxford

<sup>4</sup> Nanyang Technological University

{qianyi.wu, yuedong.chen, jianfei.cai}@monash.edu, alvinliu@ie.cuhk.edu.hk  
kejie.li@outlook.com, {chuanxia001@e., asjmzheng}@ntu.edu.sg

**Abstract.** The neural implicit representation has shown its effectiveness in novel view synthesis and high-quality 3D reconstruction from multi-view images. However, most approaches focus on holistic scene representation yet ignore individual objects inside it, thus limiting potential downstream applications. In order to learn object-compositional representation, a few works incorporate the 2D semantic map as a cue in training to grasp the difference between objects. But they neglect the strong connections between object geometry and instance semantic information, which leads to inaccurate modeling of individual instance. This paper proposes a novel framework, *ObjectSDF*, to build an object-compositional neural implicit representation with high fidelity in 3D reconstruction and object representation. Observing the ambiguity of conventional volume rendering pipelines, we model the scene by combining the Signed Distance Functions (SDF) of individual object to exert explicit surface constraint. The key in distinguishing different instances is to revisit the strong association between an individual object’s SDF and semantic label. Particularly, we convert the semantic information to a function of object SDF and develop a unified and compact representation for scene and objects. Experimental results show the superiority of *ObjectSDF* framework in representing both the holistic object-compositional scene and the individual instances. Code can be found at <https://qianyiwu.github.io/object sdf/>

**Keywords:** Neural implicit representation · Object compositionality · Volume rendering · Signed distance function

## 1 Introduction

This paper studies the problem of efficiently learning an object-compositional 3D scene representation from posed images and semantic masks, which defines the geometry and appearance of the whole scene and individual objects as well. Such a representation characterizes the compositional nature of scenes and provides additional inherent information, thus benefiting 3D scene understanding [8,20,11]

and context-sensitive application tasks such as robotic manipulation [30,16], object editing, and AR/VR [34,37]. Learning this representation yet imposes new challenges beyond those arising in the conventional 3D scene reconstruction.

The emerging neural implicit representation rendering approaches provide promising results in novel view synthesis [18] and 3D reconstruction [22,35,33]. A typical neural implicit representation encodes scene properties into a deep network, which is trained by minimizing the discrepancies between the rendered and real RGB images from different viewpoints. For example, NeRF [18] represents the volumetric radiance field of a scene with a neural network trained from images. The volume rendering method is used to compute pixel color, which samples points along each ray and performs  $\alpha$ -composition over the radiance of the sampled points. Despite not having direct supervision on the geometry, it is shown that neural implicit representations often implicitly learn the 3D geometry to render photorealistic images during training [18]. However, the scene-based neural rendering in these works is mostly *agnostic to individual object identities*.

To enable the model’s object-level awareness, several works are developed to encode objects’ semantics into the neural implicit representation. Zhi *et al.* propose an in-place scene labeling scheme [41], which trains the network to render not only RGB images but also 2D semantic maps. Decomposing a scene into objects can then be achieved by painting the scene-level geometric reconstruction using the predicted semantic labels. This workflow is not object-based modeling since the process of learning geometry is unaware of semantics. Therefore, the geometry and semantics are not strongly associated, which results in inaccurate object representation when the prediction of either geometry or semantics is bad. Yang *et al.* present an object-compositional NeRF [34], which is a unified rendering model for the scene but respecting individual object placement in the scene. The network consists of two branches: The scene branch encodes the scene geometry and appearance, and the object branch encodes each standalone object by conditioning the output only for a specific object with everything else removed. However, as proved in recent works [39,33], object supervision suffers from 3D space ambiguity in a clustered scene. It thus requires aids from extra components such as scene guidance and 3D guard masks, which are used to distill the scene information and protect the occluded object regions.

Inspired by these works, we suggest modeling the object-level geometry directly to learn the geometry and semantics simultaneously so that the representation captures “what” and “where” things are in the scene. The inherent challenge is how to get the supervision for the object-level geometry from RGB images and 2D instance semantic. Unlike the semantic label for a 3D position that is well constrained by multiple 2D semantic maps using multi-view consistency, finding a direct connection between object-level geometry and the 2D semantic labels is non-trivial. In this paper, we propose a novel method called *ObjectSDF* for object-compositional scene representations, aiming at more accurate geometry learning in highly composite scenes and more effective extraction of individual objects to facilitate 3D scene manipulation. First, ObjectSDF represents the scene at the level of objects using a multi-layer perceptron (MLP)

that outputs the Signed Distance Function (SDF) of each object at any 3D position. Note that NeRF learns a volume density field, which has difficulty in extracting a high-quality surface [40,35,39,33,38]. In contrast, the SDF can more accurately define surfaces and the composition of all object SDFs via the minimum operation that gives the SDF of the scene. Moreover, a density distribution can be induced by the scene SDF, which allows us to apply the volume rendering to learn an object-compositional neural implicit representation with robust network training. Second, ObjectSDF builds an explicit connection between the desired semantic field and the level set prediction, which braces the insight that the geometry of each object is strongly associated with semantic guidance. Specifically, we define the semantic distribution in 3D space as a function of each object’s SDF, which allows effective semantic guidance in learning the geometry of objects. As a result, ObjectSDF provides a unified, compact, and simple framework that can supervise the training by the input RGB and instance segmentation guidance naturally, and learn the neural implicit representation of the scene as a composition of object SDFs effectively. This is further demonstrated in our experiments.

In summary, the paper has the following contributions: **1)** We propose a novel neural implicit surface representation using the signed distance functions *in an object-compositional manner*. **2)** To grasp the strong associations between object geometry and instance segmentation, we propose a simple yet effective design to incorporate the segmentation guidance organically by updating each object’s SDF. **3)** We conduct experiments that demonstrate the effectiveness of the proposed method in representing individual objects and compositional scene.

## 2 Related Work

**Neural Implicit Representation.** Occupancy Networks [17] and DeepSDF [24] are among those pioneers who introduced the idea of encoding objects or scenes implicitly using a neural network. Such a representation can be considered as a mapping function from a 3D position to the occupancy density or SDF of the input points, which is continuous and can achieve high spatial resolution. While these works require 3D ground-truth models, Scene Representation Networks (SRN) [31] and Neural Radiance Field (NeRF) [18] demonstrate that both geometry and appearance can be jointly learned only from multiple RGB images using multi-view consistency. Such an implicit representation idea is further used to predict the semantic segmentation label [41,13], deformation field [25,27], high-fidelity specular reflections [32].

This learning-by-rendering paradigm of NeRF has attracted broad interest. They also lay a foundation for many follow-up works including ours. Instead of rendering a neural radiance field, several works [36,22,14,35,33,5] demonstrate that rendering neural implicit surfaces, where gradients are concentrated around surface regions, is able to produce a high-quality 3D reconstruction. Particularly, a recent work, VolSDF [35], combines neural implicit surface with volume rendering and produces high fidelity reconstructed surfaces. Due to its superior

modeling performance, our network is built upon VolSDF. The key difference is that VolSDF only has one SDF to model the entire scene while our work models the scene SDF as a composition of multiple object SDFs.

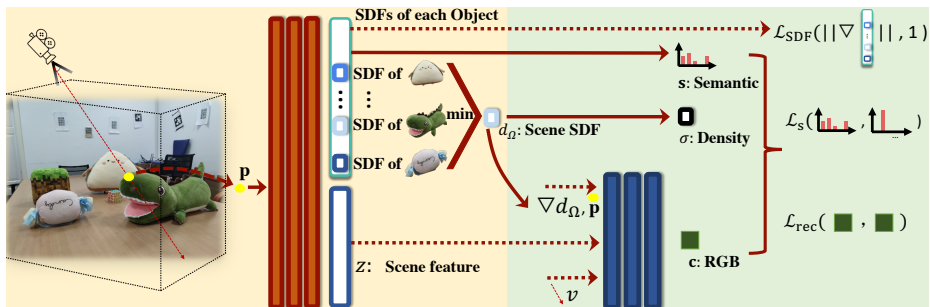
**Object-Compositional Implicit Representation.** Decomposing a holistic NeRF into several parts or object-centric representations could benefit efficient rendering of radiance fields and other applications like content generation [28,29,2]. Several attempts are made to model the scene via a composition of object representations, which can be roughly categorized as category-specific [19,7,23,21] and scene-specific [34,10,41] methods.

The category-specific methods learn the object representation of a limited number of object categories using a large amount of training data in those categories. They have difficulty in generalizing to objects in other unseen categories. For example, Guo *et al.* [7] propose a bottom-up method to learn one scattering field per object, which enables rendering scenes with moving objects and lights. Ost *et al.* [23] use a neural scene graph to represent dynamic scenes and particularly decompose objects in a street view dataset. Niemeyer and Geiger [21] propose GIRAFFE that conditions latent codes to get object-centric NeRFs and thus represents scenes as compositional generative neural feature fields.

The scene-specific methods directly learn a unified neural implicit representation for the whole scene, which also respects the object placement as in the scene [41,34]. Particularly, SemanticNeRF [41] augments NeRF to estimate the semantic label for any given 3D position. A semantic head is added into the network, which is trained by comparing the rendered and real semantic maps. Although SemanticNeRF is able to predict semantic labels, it does not explicitly model each semantic entity’s geometry. The work closest to ours is ObjectNeRF [34], which uses a two-pathway architecture to capture the scene and object neural radiance fields. However, the design of ObjectNeRF requires a series of additional voxel feature embedding, object activation encoding, and separate modeling of the scene and object neural radiance fields to deal with occlusion issues and improve the rendering quality. In contrast, our approach is a simple and intuitive framework that uses SDF-based neural implicit surface representation and models scene and object geometry in one unified branch.

### 3 Method

Given a set of  $N$  posed images  $\mathcal{A} = \{x_1, x_2, \dots, x_N\}$  and the corresponding instance semantic segmentation masks  $\mathcal{S} = \{s_1, s_2, \dots, s_N\}$ , our goal is to learn an *object-compositional implicit 3D representation* that captures the 3D shapes and appearances of not only the whole scene  $\Omega$  but also individual *objects*  $\mathcal{O}$  within the scene. Different from the conventional 3D scene modeling which typically models the scene as a whole without distinguishing individual objects within it, we consider the 3D scene as a composition of individual objects and the background. A unified simple yet effective framework is proposed for 3D scene and object modeling, which offers a better 3D modeling and understanding via inherent scene decomposition and recomposition.



**Fig. 1. Overview of our proposed *ObjectSDF* framework**, consisting of two parts: an object-SDF part (left, yellow region) and a scene-SDF part (right, green region). The former predicts the SDF of each object, while the latter composites all object SDFs to predict the scene-level geometry and appearance.

Fig. 1 shows the proposed *ObjectSDF* framework of learning *object compositional neural implicit surfaces*. It consists of an object-SDF part that is responsible for modeling all instances including background (Fig. 1, yellow part) and a *scene*-SDF part that recomposes the decomposed objects in the scene (Fig. 1, green part). Note that here we use Signed Distance Function (SDF) based neural implicit surface representation to model the geometry of the scene and objects, instead of using the popular Neural Radiance Fields (NeRF). This is mainly because NeRF aims at high-quality view synthesis, not for accurate surface reconstruction, while the SDF-based neural surface representation is better for geometry modeling and SDF is also easier for the 3D composition of objects.

In the following, we first give the background of volume rendering and its combination with SDF-based neural implicit surface representation in Section 3.1. Then, we describe how to represent a scene as a composition of multiple objects within it under a unified neural implicit surface representation in Section 3.2, and emphasize our novel idea of leveraging semantic labels to supervise the modeling of individual object SDFs in Section 3.3, followed by a summary of the overall training loss in Section 3.4.

### 3.1 Background

**Volume Rendering** essentially takes the information from a radiance field. Considering a ray  $\mathbf{r}(v) = \mathbf{o} + v\mathbf{d}$  emanated from a camera position  $\mathbf{o}$  in the direction of  $\mathbf{d}$ , the color of the ray can be computed as an integral of the transparency  $T(v)$ , the density  $\sigma(v)$  and the radiance  $\mathbf{c}(v)$  over samples taken along near and far bounds  $v_n$  and  $v_f$ ,

$$\hat{\mathbf{C}}(\mathbf{r}) = \int_{v_n}^{v_f} T(v)\sigma(\mathbf{r}(v))\mathbf{c}(\mathbf{r}(v))dv. \quad (1)$$

This integral is approximated using a numerical quadrature [15]. The transparency function  $T(v)$  represents how much light is transmitted along a ray  $\mathbf{r}(v)$

and can be computed as  $T(v) = \exp(-\int_{v_n}^v \sigma(\mathbf{r}(u))du)$ , where the volume density  $\sigma(\mathbf{p})$  is the rate that light is occluded at a point  $\mathbf{p}$ . Sometimes the radiance  $\mathbf{c}$  may not be the function only of a ray  $r(v)$ , such as in [18,36]. We refer readers to [9] for more details about volume rendering.

**SDF-based Neural Implicit Surface.** SDF directly characterizes the geometry at the surface. Specifically, given a scene  $\Omega \subset \mathbb{R}^3$ , and  $\mathcal{M} = \partial\Omega$  is the boundary surface. The Signed Distance Function  $d_\Omega$  is defined as the distance from point  $\mathbf{p}$  to the boundary  $\mathcal{M}$ :

$$d_\Omega(\mathbf{p}) = (-1)^{\mathbb{1}_\Omega(\mathbf{p})} \min_{\mathbf{y} \in \mathcal{M}} \|\mathbf{p} - \mathbf{y}\|_2, \quad (2)$$

where  $\mathbb{1}_\Omega(\mathbf{p})$  is the indicator denoting whether  $\mathbf{p}$  belongs to the scene  $\Omega$  or not. If the point is outside the scene,  $\mathbb{1}_\Omega(\mathbf{p})$  returns 0; otherwise returns 1. Typically, the standard  $l_2$ -norm is used to compute the distance.

The latest neural implicit surface works [33,35] combine SDF with neural implicit function and volume rendering for better geometry modeling, by replacing the NeRF volume density output  $\sigma(\mathbf{p})$  with the SDF value  $d_\Omega(\mathbf{p})$ , which can be directly transferred into the density. Following [35], here we model the density  $\sigma(\mathbf{p})$  using a specific tractable transformation:

$$\sigma(\mathbf{p}) = \alpha \Psi(d_\Omega(\mathbf{p})) = \begin{cases} \frac{1}{2\beta} \exp\left(\frac{d_\Omega(\mathbf{p})}{\beta}\right) & \text{if } d_\Omega(\mathbf{p}) \leq 0 \\ \frac{1}{\beta} - \frac{1}{2\beta} \exp\left(-\frac{d_\Omega(\mathbf{p})}{\beta}\right) & \text{if } d_\Omega(\mathbf{p}) > 0 \end{cases} \quad (3)$$

where  $\beta$  is a learnable parameter in our implementation.

### 3.2 The scene as object composition

Unlike the existing SDF-based neural implicit surface modeling works [33,35], which either focus on a single object or treat the entire scene as one object, we consider the scene as a composition of multiple objects and aim to model their geometries and appearances jointly. Specifically, given a static scene  $\Omega$ , it can be naturally represented by the spatial composition of  $k$  different objects

$\{\mathcal{O}_i \subset \mathbb{R}^3 | i = 1, \dots, k\}$ , i.e.,  $\Omega = \bigcup_{i=1}^k \mathcal{O}_i$  (including background, as an individ-


ual object). Using the SDF representation, we denote the scene geometry by *scene*-SDF  $d_\Omega(\mathbf{p})$  and the object geometry as object-SDF  $d_{\mathcal{O}_i}(\mathbf{p})$ , and their relationship can be derived as: for any point  $\mathbf{p} \in \mathbb{R}^3$ ,  $d_\Omega(\mathbf{p}) = \min_{i=1\dots k} d_{\mathcal{O}_i}(\mathbf{p})$ . This is fundamentally different from [33,35] that directly predict the SDF of the holistic scene  $\Omega$ , while our neural implicit function outputs  $k$  distinct SDFs corresponding to different objects (see Fig. 1). The *scene*-SDF is just a minimum of the  $k$  object-SDFs, which can be implemented as a particular type of pooling.

Considering that we do not have any explicit supervision for the SDF values in any 3D position, we adopt the implicit geometric regulation loss [6] to regularize

each object SDF  $d_{\mathcal{O}_i}$  as:

$$\mathcal{L}_{SDF} = \sum_{i=1}^k \mathbb{E}_{d_{\mathcal{O}_i}} (\|\nabla d_{\mathcal{O}_i}(\mathbf{p})\| - 1)^2. \quad (4)$$

This will also constrain the scene SDF  $d_{\Omega}$ . Once we obtain the scene SDF  $d_{\Omega}$ , we use Eq. (3) to obtain the density in the holistic scene.



**Fig. 2. Semantic as a function of object-SDF.** Left: the desired 3D semantic field that should satisfy the requirement, *i.e.*, when the ray is crossing an object (the toy), the corresponding 3D semantic label should change rapidly. Thus, we propose to use the function (6) to approximate the 3D semantic field given object-SDF. Right: The plot of function (6) versus SDF.

### 3.3 Leveraging semantics for learning object-SDFs

Although our idea of treating scene-SDF as a composition of multiple object-SDFs is simple and intuitive, it is extremely challenging to learn meaningful and accurate object-SDFs since there is no explicit SDF supervision. The only object information we have is the given 2D semantic masks  $\mathcal{S} = \{s_1, s_2, \dots, s_N\}$ . So, the critical issue we need to address here is: *How to leverage 2D instance semantic masks to guide the learning of object-SDFs?*

The only existing solution we can find is the SemanticNeRF [41], which adds an additional head to predict a 3D “semantic field”  $\mathbf{s}$  in the same way as predicting the radiance field  $\mathbf{c}$ . Then, similar to Eq. (1), 2D semantic segmentation can be regarded as a volume rendering result from the 3D “semantic field” as:

$$\hat{S}(\mathbf{r}) = \int_{v_n}^{v_f} T(v)\sigma(\mathbf{r}(v))\mathbf{s}(\mathbf{r}(v))dv. \quad (5)$$

However, in our framework,  $\sigma$  is transformed from *scene*-SDF  $d_{\Omega}$ , which is further obtained from object-SDFs. The supervision on the segmentation prediction  $\hat{S}$  cannot ensure the object-SDFs to be meaningful.

Therefore, we turn to a new solution that represents the 3D semantic prediction  $\mathbf{s}$  as a function of object-SDF. Our *key insight* is that the semantic information is strongly associated with the object geometry. Specifically, we analyze the property of a desired 3D “semantic field”. Considering we have  $k$  objects  $\{\mathcal{O}_i \subset \mathbb{R}^3 | i = 1, \dots, k\}$  inside the scene including background, we expect that *a desired 3D semantic label should maintain consistency inside one object while changing rapidly when crossing the boundary from one class to another*. Thus, we investigate the derivative of  $\mathbf{s}(\mathbf{p})$  at a 3D position  $\mathbf{p}$ . Particularly, we inspect the norm of  $\frac{\partial \mathbf{s}_i}{\partial \mathbf{p}}$ :

$$\begin{aligned} \left\| \frac{\partial \mathbf{s}_i}{\partial \mathbf{p}} \right\| &= \left\| \frac{\partial \mathbf{s}_i}{\partial d_{\mathcal{O}_i}} \cdot \frac{\partial d_{\mathcal{O}_i}}{\partial \mathbf{p}} \right\| && \text{(chain rule)} \\ &\leq \left\| \frac{\partial \mathbf{s}_i}{\partial d_{\mathcal{O}_i}} \right\| \cdot \left\| \frac{\partial d_{\mathcal{O}_i}}{\partial \mathbf{p}} \right\| && \text{(norm inequality)} \\ &= \left\| \frac{\partial \mathbf{s}_i}{\partial d_{\mathcal{O}_i}} \right\| \cdot \|\nabla d_{\mathcal{O}_i}(\mathbf{p})\|. \end{aligned}$$

As we adopt the implicit geometric regularization loss in Eq. (4),  $\|\nabla d_{\mathcal{O}_i}(\mathbf{p})\|$  should be close to 1 after training. Therefore, the norm of  $\partial \mathbf{s}_i / \partial \mathbf{p}$  should be bounded by the norm of  $\partial \mathbf{s}_i / \partial d_{\mathcal{O}_i}$ . In this way, we can convert the desired property of the 3D “semantic field” to  $\partial \mathbf{s}_i / \partial d_{\mathcal{O}_i}$ . Considering that crossing from one class to another class means  $d_{\mathcal{O}_i}$  is passing zero-level set (see Fig. 2 left), we come up with a simple but effective function to satisfied the property. Concretely, we use the function:

$$\mathbf{s}_i = \gamma / (1 + \exp(\gamma d_{\mathcal{O}_i})), \quad (6)$$

which is a scaled sigmoid function and  $\gamma$  is a hyper-parameter to control the smoothness of the function. The absolute value of  $\partial \mathbf{s}_i / \partial d_{\mathcal{O}_i}$  is  $\gamma^2 \exp(\gamma d_{\mathcal{O}_i}) / (1 + \exp(\gamma d_{\mathcal{O}_i}))^2$ , which meets the requirement of desired 3D semantic field, *i.e.*, smooth inside the object but a rapid change at the boundary (see Fig. 2 right).

This is fundamentally different from [41]. Here we directly transform the object-SDF prediction  $d_{\mathcal{O}_i}$  to a semantic label in 3D space. Thanks to this design, we can conduct volume rendering to convert the transformed SDF into the 2D semantic prediction using Eq. (5). With the corresponding semantic segmentation mask, we minimize the cross-entropy loss  $\mathcal{L}_s$ :

$$\mathcal{L}_s = \mathbb{E}_{\mathbf{r} \sim S} [-\log \hat{S}(\mathbf{r})]. \quad (7)$$

### 3.4 Model Training

Following [18, 35], we first minimize the reconstruction error between the predicted color  $\hat{C}(\mathbf{r})$  and the ground-truth color  $C(\mathbf{r})$  with:

$$\mathcal{L}_{rec} = \mathbb{E}_{\mathbf{r}} \|\hat{C}(\mathbf{r}) - C(\mathbf{r})\|_1. \quad (8)$$

Furthermore, we use the implicit geometric loss to regularize the SDF of each object as in Eq. (4). Moreover, the cross-entropy loss between the rendered



semantic and ground-truth semantic is applied to guide the learning of object-SDFs as in Eq. (7). Overall, we train our model with the following three losses  $\mathcal{L}_{total} = \mathcal{L}_{rec} + \lambda_1 \mathcal{L}_s + \lambda_2 \mathcal{L}_{SDF}$ , where  $\lambda_1$  and  $\lambda_2$  are two trade-off hyperparameters. We set  $\lambda_1 = 0.04$  and  $\lambda_2 = 0.1$  empirically

## 4 Experiments

The main purpose of our proposed method is to build an object-compositional neural implicit representation for scene rendering and object modeling. Therefore, we evaluate our approach in two real-world datasets from two aspects. Firstly, we quantitatively compare our scene representation ability with the state-of-the-art methods on standard scene rendering and modeling aspects. Then, we investigate the object representation ability of our method and compare it with NeRF-based object representation method [34]. Finally, we perform a model design ablation study to inspect the effectiveness of our framework.

### 4.1 Experimental Setting

**Implementation Details.** Our systems consists of two Multi-Layer Perceptrons (MLP). (i) The first MLP  $f_\phi$  estimates each object SDF as well as a scene feature  $z$  of dimension 256 for further rendering branch, i.e.,  $f_\phi(\mathbf{p}) = [d_{\mathcal{O}_1}(\mathbf{p}), \dots, d_{\mathcal{O}_K}(\mathbf{p}), z(\mathbf{p})] \in \mathbb{R}^{K+256}$ .  $f_\phi$  consists of 6 layers with 256 channels. (ii) The second MLP  $f_\theta$  is used to estimate the scene radiance field, which takes point position  $\mathbf{p}$ , point normal  $\mathbf{n}$ , view direction  $\mathbf{d}$  and scene feature  $z$  as inputs and outputs the RGB color  $\mathbf{c}$ , i.e.,  $f_\theta(\mathbf{p}, \mathbf{n}, \mathbf{d}, z) = \mathbf{c}$ .  $f_\theta$  consists of 4 layers with 256 channels. We use the geometric network initialization technique [1,36] for both MLPs to initial the network weights to facilitate the learning of signed distance functions. We adopt the error-bounded sampling algorithm proposed by [35] to decide which points will be used in calculating volume rendering results. We also incorporate the positional encoding [18] with 6 levels for position  $\mathbf{p}$  and 4 levels for view direction  $\mathbf{d}$  to help the model capture high frequency information of the geometry and radiance field. Our model can be trained in a single GTX 2080Ti GPU with a batch size of 1024 rays. We set  $\beta = 0.1$  in Eq. 3 in the initial stage of training.

**Datasets.** Following [34] and [41], we use two real datasets for comparisons. - **ToyDesk** [34] contains scenes of a desk by placing several toys with two different layouts and capturing images in  $360^\circ$  by looking at the desk center. It also contains 2D instance segmentation for target objects as well as the camera pose for each image and a reconstructed mesh for each scene. - **ScanNet** dataset [3] contains RGB-D indoor scene scans as well as 3D segmentation annotations and projected 2D segmentation masks. In our experiments, we use the 2D segmentation masks provided in the ScanNet dataset for training, and the provided 3D meshes for 3D reconstruction evaluation.

**Comparison Baselines.** We compare our method with the recent representative works in the realm of object-compositional neural implicit representation

**Table 1. The quantitative results on scene representation.** We compare our method against recent SOTA methods [41,34], ablation designs and Ground Truth

Methods	#params	ToyDesk [34]			ScanNet [3]		
		PSNR $\uparrow$	mIOU $\uparrow$	CD $\downarrow$	PSNR $\uparrow$	mIOU $\uparrow$	CD $\downarrow$
Ground Truth (GT)	-	N/A	1.00	0.00	N/A	1.00	0.00
SemanticNeRF [41]	$\sim$ 1.26M	19.57	0.79	1.15	23.59	<b>0.57</b>	0.64
ObjectNeRF [34]	1.78M(+19.20M)	21.61	0.75	1.06	24.01	0.31	0.61
VolSDF [35]	<b>0.802M</b>	22.00	-	0.30	25.31	-	0.32
VolSDF w/ Semantic	$\sim$ 0.805M	22.00	<b>0.89</b>	0.34	<b>25.41</b>	0.56	0.28
<b>Ours</b>	<b><math>\sim</math>0.804M</b>	<b>22.00</b>	<b>0.88</b>	<b>0.19</b>	<b>25.23</b>	0.53	<b>0.22</b>

for the single static scene: **ObjectNeRF** [34] and **SemanticNeRF** [26]. ObjectNeRF uses a two-path architecture to represent object-compositional neural radiance, where one branch is used for individual object modeling while the other is for scene representation. To broaden the ability of the network to capture accurate scene information, ObjectNeRF utilizes voxel features for both scene and object branches training, as in [12], which significantly increases the model complexity. SemanticNeRF [26] is a NeRF-based framework that jointly predicts semantics and geometry in a single model for semantic labeling. The key design in this framework is an additional semantic prediction head extended from the NeRF backbone. Although this method does not directly represent objects, it can still extract an object by using semantic prediction.

**Metric.** We employ the following metrics for evaluation: **1) PSNR** to evaluate the quality of rendering; **2) mIOU** to evaluate the semantic segmentation; and **3) Chamfer Distance (CD)** to measure the quality of reconstructed 3D geometry. Besides these metrics, we also provide the number of neural network parameters (**#params**) of each method for comparing the model complexity.

**Comparison Settings.** We follow the comparison settings introduced by ObjectNeRF [34] and SemanticNeRF [41]. We use the same scene data used in [34] from ToyDesk and ScanNet for a fair comparison. To be consistent with SemanticNeRF [41], we predict the category semantic label rather than the instance semantic label for the quantitative evaluation on the ScanNet benchmark. Note that we are unable to train SemanticNeRF in the original resolution in the official codebase due to memory overflow. Therefore, we downscale the images of ScanNet and train all methods with the same data. We also noticed that the ground truth mesh may lack points in some regions, for which we apply the same crop setting for all methods to evaluate the 3D region of interest.

It is worth noting that both our method and SemanticNeRF are able to produce the semantic label in the output. However, ObjectNeRF does not explicitly predict the semantic label in their framework. Therefore, we calculate the depth of each object which is computed from the volume density predicted from each object branch in ObjectNeRF [4]. Then, we use the object with the nearest

depth as the pixel semantic prediction of ObjectNeRF for calculating the mIOU metric. For scene rendering and the 3D reconstruction ability of ObjectNeRF, we adopt the result from the scene branch for evaluation. More details can be found in the supplementary.

## 4.2 Scene-level Representation Ability

To evaluate the scene-level representation ability, we first compare the scene rendering, object segmentation, and 3D reconstruction results. As shown in Tab. 1, our framework outperforms other methods on the Toydesk benchmark and is comparable or even better than the SOTA methods on the ScanNet dataset. The qualitative results shown in Fig. 3 demonstrate that both our method and SemanticNeRF are able to produce fairly accurate segmentation masks. ObjectNeRF, on the other hand, renders noisy semantic masks as shown in the third row of Fig. 3. We believe the volume density predicted by the object branch is susceptible to noisy semantic prediction for points that are further from the object surface. Therefore, when calculating the depth of each object, it results in artifacts and leads to noisy rendering.

In terms of 3D structure reconstruction, thanks to the accurate SDF in capturing surface information, our framework can recover much more accurate geometry compared with other methods. We also calculate the number of model parameters of each method. Due to the feature volume used in ObjectNeRF, their model needs additional 19.20M parameters. In contrast, the number of parameters of our model is about 0.804M, which is about 36% and 54% reductions from ObjectNeRF and SemanticNeRF, respectively. This demonstrates the compactness and efficiency of our proposed method.

## 4.3 Object-level Representation Ability

Besides the scene-level representation ability, our framework can naturally represent each object by selecting the specific output channel of object-SDFs for volume rendering. ObjectNeRF [34] can also isolate an object in a scene by computing the volume density and color of the object using the object branch network with a specific object activation code. We evaluate the object-level representation ability based on the quality of rendering and reconstruction of each object. Particularly, we compare our method against ObjectNeRF on Toydesk02 which contains five toys in the scene as shown in Fig. 4. We show the rendered opacity and RGB images of each toy from the same camera pose. It can be seen that our proposed method can render the objects more precisely with accurate opacity to describe each object. In contrast, ObjectNeRF often renders noisy images despite utilizing the opacity loss and 3D guided mask to stop gradient during training. Moreover, the accurate rendering of the occluded cubes (the last two columns in Fig. 4) demonstrates that our method handles occlusions much better than ObjectNeRF. We also compare the geometry reconstructions of all the five objects on the left of Fig. 4.

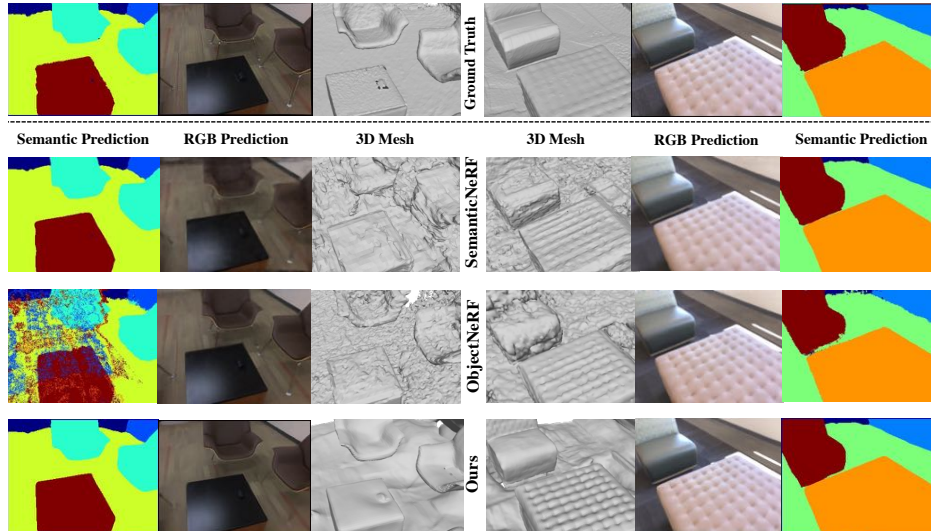


Fig. 3. Qualitative Comparison with SemanticNeRF [41] and ObjectNeRF [34] on scene-level representation ability. We show the reconstructed meshes, predicted RGB images, and semantic masks of each method together with the ground truth results from two scenes in ScanNet.





Fig. 4. Instance Results of ObjectNeRF [34] and Ours. We show the reconstructed mesh, rendered opacity, and RGB images of different objects.

#### 4.4 Ablation Study

Our framework is built upon VolSDF [35] to develop an object-compositional neural implicit surface representation. Instead of modeling individual object SDFs, an alternative way to achieve the same goal is to add a semantic head to

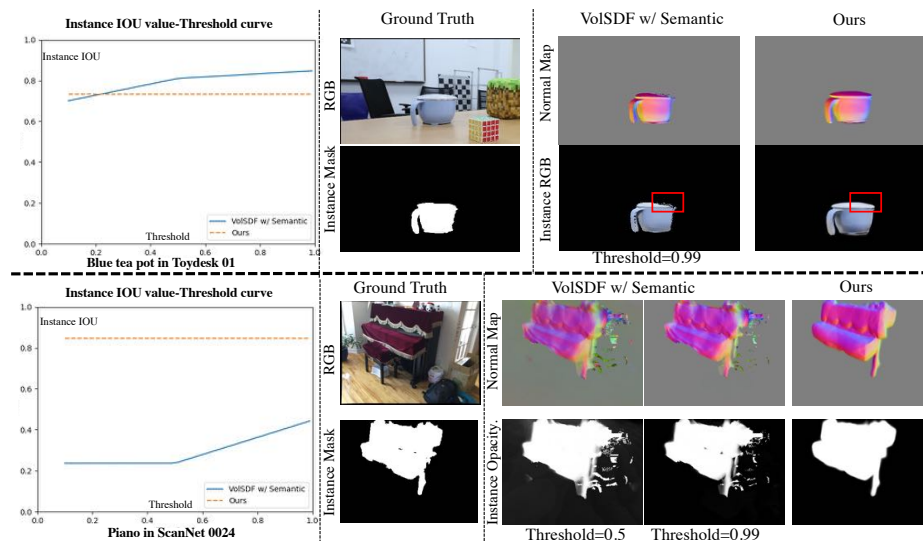


**Fig. 5. Ablation study results on scene representation ability.** We show the rendered RGB image and rendered normal map together with ground truth image.

VolSDF to predict the semantic label given each 3D location, which is similar to the approach done in [41]. We name this variant as “VolSDF w/ Semantic”.

We first evaluate the scene-level representation ability between our method and the variant “VolSDF w/ Semantic” in Tab. 1. For completeness, we also include the vanilla VolSDF, but due to the lack of semantic head, it cannot be evaluated on mIOU. While the comparing methods achieve similar performance on image rendering measured by PSNR, our method excels at geometric reconstruction. This is further demonstrated in Fig. 5, where we render the RGB image and normal map of each method. From the rendered normal maps, we can see that our method captures more accurate geometry compared with the two baselines. For example, our method can recover the geometry of the floor and the details of the sofa legs. The key difference between our method and the two variants is that we directly model each object SDF inside the scene. This indicates that our object-compositional modeling can improve the full understanding of 3D scene both semantically and geometrically.

To investigate the object representation ability of “VolSDF w/ semantic”, we obtain an implicit object representation by using the prediction of semantic labels to determine the volume density of an object. In particular, given a semantic prediction in a 3D position, we can truncate the object semantic value by a threshold to decide whether to use the density to represent this object. We evaluate the object representation ability on two instances from ToyDesk and Scannet, respectively, in Fig. 6. We choose the object, which is not occluded to extract the complete segmentation mask, and then use this mask to evaluate the semantic prediction result for each instance. Because the instance mask generated by “VolSDF w/ Semantic” is controlled by the semantic value threshold, we plot the curve of IOUs under different thresholds (blue line). This reveals an inherent challenge for “VolSDF w/ Semantic”, *i.e.*, how to find a generally suitable threshold across different instances or scenes. For instance, we notice “VolSDF w/ Semantic” could gain a high IOU value with a high threshold of 0.99, but it will miss some information of the teapot (as highlighted by the red box). While



**Fig. 6. Instance Results of Ours and “VoISDF w/ Semantic”.** We show the curve between instance IOU value and semantic value threshold (left), the ground truth instance image and mask (middle), the rendered normal map, and RGB/opacity of each instance under different threshold values (right).

using the same threshold of 0.99 on ScanNet 0024 (bottom), it fails in separating the piano. In contrast, our instance prediction is invariant to the threshold as shown in the yellow dash line. This suggests that the separate modeling of 3D structure and semantic information is undesirable to extract accurate instance representation when either prediction is inaccurate. We also observe that given a fairly rough segmentation mask during training, our framework can produce a smooth and high-fidelity object representation as shown in Fig. 6.

## 5 Conclusion and Future Work

We have presented an object-compositional neural implicit surface representation framework, namely *ObjectSDF*, which learns the signed distance functions of all objects in a scene from the guidance of 2D instance semantic segmentation masks and RGB images using a single network. Our model unifies the object and scene representations in one framework. The main idea behind it is building a strong association between semantic information and object geometry. Extensive experimental results on two datasets have demonstrated the strong ability of our framework in both 3D scene and object representation. Future work includes applying our model for various 3D scene editing applications and efficient training of neural implicit surfaces.

**Acknowledgements** This research is partially supported by Monash FIT Start-up Grant and SenseTime Gift Fund.

## References

1. Atzmon, M., Lipman, Y.: Sal: Sign agnostic learning of shapes from raw data. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
2. Chen, Y., Wu, Q., Zheng, C., Cham, T.J., Cai, J.: Sem2nerf: Converting single-view semantic masks to neural radiance fields. arXiv preprint arXiv:2203.10821 (2022)
3. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In: CVPR (2017)
4. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised nerf: Fewer views and faster training for free. arXiv preprint arXiv:2107.02791 (2021)
5. Deng, Y., Yang, J., Xiang, J., Tong, X.: Gram: Generative radiance manifolds for 3d-aware image generation. arXiv preprint arXiv:2112.08867 (2021)
6. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. arXiv preprint arXiv:2002.10099 (2020)
7. Guo, M., Fathi, A., Wu, J., Funkhouser, T.: Object-centric neural scene rendering. arXiv preprint arXiv:2012.08503 (2020)
8. Hassan, M., Choutas, V., Tzionas, D., Black, M.J.: Resolving 3d human pose ambiguities with 3d scene constraints. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 2282–2292 (2019)
9. Kajiya, J.T., Von Herzen, B.P.: Ray tracing volume densities. ACM SIGGRAPH computer graphics **18**(3), 165–174 (1984)
10. Kohli, A., Sitzmann, V., Wetzstein, G.: Semantic Implicit Neural Scene Representations with Semi-supervised Training. In: International Conference on 3D Vision (3DV) (2020)
11. Li, K., Rezatofghi, H., Reid, I.: Moltr: Multiple object localization, tracking and reconstruction from monocular rgb videos. IEEE Robotics and Automation Letters **6**(2), 3341–3348 (2021)
12. Liu, L., Gu, J., Lin, K.Z., Chua, T.S., Theobalt, C.: Neural sparse voxel fields. arXiv preprint arXiv:2007.11571 (2020)
13. Liu, X., Xu, Y., Wu, Q., Zhou, H., Wu, W., Zhou, B.: Semantic-aware implicit neural audio-driven video portrait generation. arXiv preprint arXiv:2201.07786 (2022)
14. Luan, F., Zhao, S., Bala, K., Dong, Z.: Unified shape and svbrdf recovery using differentiable monte carlo rendering. In: Computer Graphics Forum. vol. 40, pp. 101–113. Wiley Online Library (2021)
15. Max, N.: Optical models for direct volume rendering. IEEE Transactions on Visualization and Computer Graphics **1**(2), 99–108 (1995)
16. McCormac, J., Handa, A., Davison, A., Leutenegger, S.: Semanticfusion: Dense 3d semantic mapping with convolutional neural networks. In: 2017 IEEE International Conference on Robotics and automation (ICRA). pp. 4628–4635. IEEE (2017)
17. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4460–4470 (2019)
18. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European conference on computer vision. pp. 405–421. Springer (2020)



19. Nguyen-Phuoc, T.H., Richardt, C., Mai, L., Yang, Y., Mitra, N.: Blockgan: Learning 3d object-aware scene representations from unlabelled images. *Advances in Neural Information Processing Systems* **33**, 6767–6778 (2020)
20. Nie, Y., Han, X., Guo, S., Zheng, Y., Chang, J., Zhang, J.J.: Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 55–64 (2020)
21. Niemeyer, M., Geiger, A.: Giraffe: Representing scenes as compositional generative neural feature fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11453–11464 (2021)
22. Oechsle, M., Peng, S., Geiger, A.: Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5589–5599 (2021)
23. Ost, J., Mannan, F., Thurey, N., Knodt, J., Heide, F.: Neural scene graphs for dynamic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2856–2865 (2021)
24. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 165–174 (2019)
25. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5865–5874 (2021)
26. Prajwal, K., Mukhopadhyay, R., Namboodiri, V.P., Jawahar, C.: A lip sync expert is all you need for speech to lip generation in the wild. In: *Proceedings of the 28th ACM International Conference on Multimedia*. pp. 484–492 (2020)
27. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10318–10327 (2021)
28. Rebain, D., Jiang, W., Yazdani, S., Li, K., Yi, K.M., Tagliasacchi, A.: Derf: Decomposed radiance fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14153–14161 (2021)
29. Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 14335–14345 (2021)
30. Rosinol, A., Gupta, A., Abate, M., Shi, J., Carlone, L.: 3d dynamic scene graphs: Actionable spatial perception with places, objects, and humans. *arXiv preprint arXiv:2002.06289* (2020)
31. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. *arXiv preprint arXiv:1906.01618* (2019)
32. Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J.T., Srinivasan, P.P.: Ref-NeRF: Structured view-dependent appearance for neural radiance fields. *CVPR* (2022)
33. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS* (2021)
34. Yang, B., Zhang, Y., Xu, Y., Li, Y., Zhou, H., Bao, H., Zhang, G., Cui, Z.: Learning object-compositional neural radiance field for editable scene rendering. In: *International Conference on Computer Vision (ICCV)* (October 2021)



35. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. arXiv preprint arXiv:2106.12052 (2021)
36. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* **33** (2020)
37. Yu, H.X., Guibas, L., Wu, J.: Unsupervised discovery of object radiance fields. In: *International Conference on Learning Representations* (2022), <https://openreview.net/forum?id=rwE8SshAlxw>
38. Zhang, K., Luan, F., Li, Z., Snavely, N.: Iron: Inverse rendering by optimizing neural sdfs and materials from photometric images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5565–5574 (2022)
39. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492 (2020)
40. Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)* **40**(6), 1–18 (2021)
41. Zhi, S., Laidlow, T., Leutenegger, S., Davison, A.: In-place scene labelling and understanding with implicit scene representation. In: *Proceedings of the International Conference on Computer Vision (ICCV)* (2021)