

# A General Implicit Framework for Fast NeRF Composition and Rendering

Xinyu Gao, Ziyi Yang, Yunlu Zhao, Yuxiang Sun, Xiaogang Jin, Changqing Zou

State Key Lab of CAD&CG, Zhejiang University

## Abstract

A variety of Neural Radiance Fields (NeRF) methods have recently achieved remarkable success in high render speed. However, current accelerating methods are specialized and incompatible with various implicit methods, preventing real-time composition over various types of NeRF works. Because NeRF relies on sampling along rays, it is possible to provide general guidance for acceleration. To that end, we propose a general implicit pipeline for composing NeRF objects quickly. Our method enables the casting of dynamic shadows within or between objects using analytical light sources while allowing multiple NeRF objects to be seamlessly placed and rendered together with any arbitrary rigid transformations. Mainly, our work introduces a new surface representation known as Neural Depth Fields (NeDF) that quickly determines the spatial relationship between objects by allowing direct intersection computation between rays and implicit surfaces. It leverages an intersection neural network to query NeRF for acceleration instead of depending on an explicit spatial structure. Our proposed method is the first to enable both the progressive and interactive composition of NeRF objects. Additionally, it also serves as a previewing plugin for a range of existing NeRF works.

## Introduction

The development of virtual worlds has revealed a strong preference for neural techniques. Neural Radiance Fields (NeRF), an image-based rendering method (Mildenhall et al. 2020), has shown its proficiency in rendering with high fidelity while also offering the added advantage of compression capabilities, resulting in a low storage footprint. As a result, NeRF has been extensively studied in human face editing (Wang et al. 2022a; Jiang et al. 2022; Sun et al. 2022b), autonomous driving (Tancik et al. 2022; Kundu et al. 2022), stylization (Gu et al. 2022; Huang et al. 2022b), mesh reconstruction (Oechsle, Peng, and Geiger 2021; Wang et al. 2021; Sun et al. 2022a), inverse rendering (Boss et al. 2021a; Knodt et al. 2021; Boss et al. 2021b, 2022; Zhang et al. 2021; Srinivasan et al. 2021), and so on.

Although a number of NeRF composition methods (Ma et al. 2023; Yang et al. 2021; Wu et al. 2022; Shuai et al.

2022; Ost et al. 2021; Guo et al. 2020) have achieved impressive results and improved the flexibility by modeling a scene as multiple semantic parts rather than as whole, compositing neural radiance fields in real-time still remains under-explored. Existing acceleration methods, such as implicit methods (Kurz et al. 2022; Neff et al. 2021) that rely on a limited scope of camera activity, hybrid methods (Reiser et al. 2021; Liu et al. 2020; Martel et al. 2021; Müller et al. 2022) that require a complex spatial structure for querying, or pre-computing methods (Yu et al. 2021; Hedman et al. 2021; Chen et al. 2022; Garbin et al. 2021) that incur significant storage costs, are far from straightforward in this task due to various challenges, let alone forming a general pipeline. First, it is difficult to satisfy low memory-storage cost and real-time speed while maintaining arbitrary affine transformations for implicit objects. Second, as the number of objects to be composed grows (e.g., a simple duplicate of 1k balls), rendering the entire scene (especially considering shadows among implicit objects) takes long, and the cost of hybrid or pre-computing methods becomes prohibitively expensive.

To address the aforementioned issues, we present a general implicit framework that can quickly perform compositing and shadow casting of NeRF objects without querying any complex structure, thereby avoiding the excessive memory-storage costs. Meanwhile, we propose Neural Depth Fields (NeDF), a novel implicit representation for implicit surface that aims for fast and precise depth estimation under a flexible camera activity scope and arbitrary affine transformations of objects, overcoming critical issues associated with previous implicit acceleration methods while retaining the benefits of implicit methods with a low storage cost. When working with a single object in a composed scene, manipulation speeds can even reach real-time in our CUDA implementation. Notably, we use the classic concept of shadow mapping based on NeDF for efficient rendering of dynamic shadows cast by either point light or parallel light sources, which represents a significant advancement over previous NeRF works aimed at acceleration.

Our work is able to provide users a fast NeRF scene composition tool allowing a quick preview of the composited results, as users may prefer interactive-level speed over perfect quality when manipulating each NeRF object. The key technical contributions are summarized below:

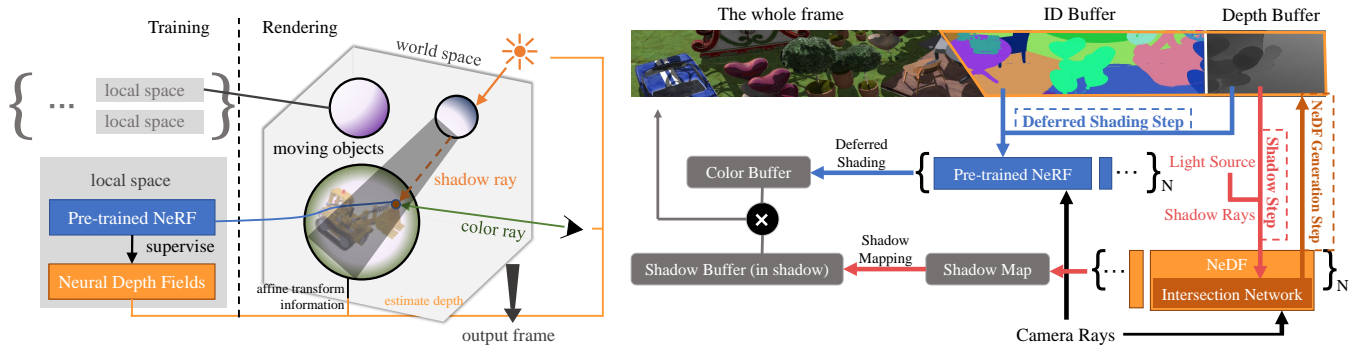


Figure 1: An overview of our framework (left) and data flow (right). In terms of framework, we use Neural Depth Fields (NeDF) as the representation of the implicit surface within each object’s local space during training, under the supervision of pre-trained NeRF. There are three main steps in render time: 1) Step1 is ”NeDFs generation step”, depth buffer and ID buffer are created by a collection of NeDFs; 2) Step2 is ”Deferred shading step”, those 2 buffers are then used by a collection of NeRFs to fill a color buffer with no shadows; 3) Step3 is ”Shadow step”, shadow rays are generated based on the depth buffer and NeDFs are used again to provide a shadow map. Finally, the color buffer and shadow buffer are multiplied to produce the results.

- A novel framework that enables interactive previewing of the NeRF composition process at a low storage cost, thereby providing a new interactive way to create novel realistic scenes from real-world scenes.
- A quick and efficient method for generating intra and inter-object dynamic shadows cast by either point light or parallel light sources.
- A new geometry representation for NeRF objects that aids intersecting/depth estimation during arbitrary manipulation.

## Related Work

**Compositing Neural Radiance Fields.** Modeling a scene as a single NeRF remains generally limited in functionality. In many scenarios, the ability to isolate individual objects and manipulate them under controlled constraints is necessary (e.g., driving scenes (Krishnan et al. 2023)). In order to enhance the controllability and editability of the NeRF-based virtual scenes, one straightforward idea is to model scenes as multiple semantic parts rather than a whole at the beginning. This can be achieved by separately compressing background and foreground objects into two branches represented by neural voxels (Yang et al. 2021) or a set of neural SDFs (Wu et al. 2022). In order to enable a better perception for foreground objects, clues such as 3D motion (Shuai et al. 2022) or segmentation mask (Kundu et al. 2022; Tancik et al. 2022) in 2D image space may be required as a preprocessing step. Another idea is compositing pre-trained NeRF models together to form a more complex and controllable scene. To achieve this, explicit spatial structures like multi-granularity voxels (Liu et al. 2020) and 3D bounding box (Ost et al. 2021; Shuai et al. 2022) are employed to indicate the sampling range for each nerf model in global space. By manipulating these voxels/boxes, NeRF objects are transformed from their local space coordinates to the global world space. Additionally, techniques that pre-compute radiance in explicit structures have shown promise for scene compositing. Despite the benefits of uti-

lizing multiple neural fields, doing so may result in significant computational requirements, leading to notable storage overhead in the pre-compute methods or reduced rendering speeds in implicit methods, which can negatively affect the ease with which users can manipulate objects in an interactive preview and loading speed. This limitation contradicts the goal of enabling real-time scene editing and rendering. Unlike these methods, we attempt to composite NeRF objects interactively and progressively for a quick preview.

**Accelerating Rendering of NeRF-modeled Scenes** Real-time rendering has been a long-standing goal in the field of computer graphics, and this objective extends to the realm of NeRF. The explicit (pre-computed) methods achieved real-time by pre-computing radiance color as spherical Harmonics coefficients (Yu et al. 2021) or feature vectors (Hedman et al. 2021; Chen et al. 2022), and then storing them in spatial structures like octree for quick inference, thereby reducing the computational overhead at render time. Among these explicit methods, the state-of-the-art method (Yariv et al. 2023) converts NeRF scenes into a combination of meshes and spherical Gaussians, resulting in fast rendering speed on modern mobile graphics pipelines. However, the explicit method can result in high storage costs for even a single scene. In contrast, implicit (MLP-based) methods that mainly focus on importance sampling (Piala and Clark 2021; Neff et al. 2021; Kurz et al. 2022) significantly reduce storage overhead and computational burden. But these methods can limit the activity scope of the camera to a specific view cell. The hybrid method (Liu et al. 2020) incorporates a neural component within the explicit grid to achieve a cost-effective solution. All of these methods are highly specialized, designed to address their own accelerating path. Our goal, on the other hand, is to investigate a general pipeline that enables rapid composition across various NeRF works.

**Neural Light Fields (NeLF)** The representation of Light Fields has been widely utilized in computer graphics to model a scene as a 4D representation, enabling fast image-

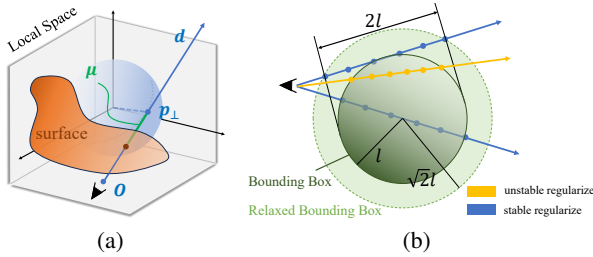


Figure 2: (a) Given a ray traveling through an object’s local space, the first intersection on the surface can be obtained by the distance between the intersection and the tangency point  $\mathbf{p}_\perp$  (The sphere is in the center of local space). Since  $\mathbf{p}_\perp$  can be quickly and uniquely determined, this representation leads to a direct intersection computation. (b) To improve the stability of NeDF, we regularize the main body of each ray, approximated by a tuple of points sampled along the ray, within the relaxed bounding box before feeding it into the intersection network. The exact bounding box is relaxed because it can cause unstable range when a ray does not travel through the center.

based rendering. The emergence of deep learning in recent years has led to the investigation of Neural Light Fields (NeLF) as an extension of this concept (Meng et al. 2019; Feng and Varshney 2021; Wizadwongsa et al. 2021; Attal et al. 2022; Sitzmann et al. 2021; Suhail et al. 2022). Recently, R2L (Cao et al. 2022; Wang et al. 2022b) successfully obtained NeLF with a comparable quality to NeRF, which provides a straightforward insight into the capability of light fields’ representation to generate high-quality results under  $360^\circ$  cases. In comparison to NeRF, NeLF has the distinct advantage of requiring only a single evaluation per ray, thereby eliminating the need for time-consuming ray marching. However, the composition task is unsuitable for NeLF due to its lack of precise depth information, though several past approaches (Dansereau and Bruton 2004; Tomic and Berkner 2014; Kim et al. 2013) have tried to leverage light fields to extract geometry information for depth estimation. Nevertheless, NeLF has inspired us to efficiently composite NeRF in a geometry-aware manner.

## Methodology

Given a set of objects represented by NeRF, our approach provides a quick preview of the resulting scene in which all of the NeRF objects are manipulated with arbitrary affine transformations and rendered together. The overview and data flow of our render framework is shown in Fig. 1, which is also described by the algorithm pseudo-codes in appendix.

### Neural Depth Fields (NeDF)

Neural Depth Field is defined on a 5-dimensional continuous function  $\mathcal{F}_\Phi$  that represents a scene. This function maps an arbitrary ray (represented by a starting point and a direction) to the distance/depth between the starting point of the ray  $\mathbf{r}$  to the ray-object intersection point in a 3D virtual world space as shown in Fig. 2a. Assuming that each object used

for scene composition is opaque, this output distance can enable the computation of the surface point color of composed NeRF, as well as the relative spatial occlusion relationship of each composed object, which are essential for scene composition and rendering.

The continuous function  $\mathcal{F}_\Phi$  is implemented with a Multilayer Perceptron network, termed as *intersection network*.  $\mathcal{F}_\Phi$  does not directly infer the depth between the starting point of the ray  $\mathbf{r}$  and the ray-object intersection point. It instead infers the distance between the ray-object intersection point and the tangency point at which the ray  $\mathbf{r}$  is tangent to a sphere centered at the origin in the local space of the object (i.e.,  $\mathcal{F}_\Phi$  is used to compute  $\mu$ , and  $\mu$  is the distance between the tangency point  $\mathbf{p}_\perp$  of the ray  $\mathbf{r}$  to the ray-object intersection point, as shown in Fig. 2a). We choose to learn  $\mu$  rather than  $D(\mathbf{r})$  is because it’s impossible to learn an infinite space.  $\mu$  is defined in the local space and has a limited range of values.

Specifically, for a given ray  $\mathbf{r}$ , we mathematically define  $\mathcal{F}_\Phi$  as:

$$\mathcal{F}_\Phi : \mathbf{r} \mapsto (\mu, \hat{\alpha}), \quad (1)$$

where  $\hat{\alpha}$  is a binary indicator that signifies whether the ray intersects the object’s surface.

Once the  $\mu$  is predicted by the intersection network, the depth  $D(\mathbf{r})$  between the starting point of the ray  $\mathbf{r}$  and the ray-object intersection point can be calculated as:

$$D(\mathbf{r}) = |\mathbf{o} - \mathbf{p}_\perp| - \mu, \quad (2)$$

where  $\mathbf{o}$  is the starting point of the ray and  $\mathbf{p}_\perp$  is the tangency point. We adopt a set of points sampled along rays to represent an oriented ray to improve the stable performance of the intersection network. Specifically, we regularize the sampling within a specific local range along rays and sample 16 points in total (see Fig. 2b).

### Intersection network

Directly implementing the intersection network as a regression problem of the distance  $\mu$  is not a good choice and tends to produce an inaccurate estimation because of the depth discontinuity issue (Neff et al. 2021) during composition, as illustrated in Fig. 3(a). This inaccurate estimation around edges can lead to floating artifacts during rendering as shown in Fig. 3(c).

To address this problem, we instead build the intersection network as a multi-level classifier. As shown in Fig. 3(b), we first regularize  $\mu$  in a limited range  $[-l, l]$  with symmetric center as  $\mathbf{p}_\perp$ . We then divide the limited range into two levels using a segment operator  $\mathcal{S}$ , namely the coarse level with  $N_c$  bins and the fine level with  $N_f$  bins. The intersection network then is formulated as:

$$\mathcal{F}_\Phi : \mathbf{r} \mapsto (\hat{v}_c, \hat{v}_f, \hat{\alpha}), \quad (3)$$

where  $\hat{v}_c, \hat{v}_f$  are 1D vectors representing the coarse- and fine-level target bins to which the surface belongs, respectively. The resolution of the two-level classifier is  $N_c \times N_f$  (where  $N_c$  and  $N_f$  are set to 64 and 128, respectively). Then,  $\mu$  is calculated as follows:

$$\mu = \mathcal{S}^{-1} \circ \mathcal{F}_\Phi = \lambda_1 p_c(\hat{v}_c) + \lambda_2 p_f(\hat{v}_f) - l, \quad (4)$$




Figure 3: Directly regressing  $\mu$  may result in significant errors in the discontinuity area (a) during composition, particularly during shadow casting (see the top row of c), we address this problem by replacing the regression of  $\mu$  with a multi-level classifier (b).

where  $\mathcal{S}^{-1}$  is the inverse operator of  $\mathcal{S}$ ,  $p_c(\hat{v}_c) = \text{argmax}(\hat{v}_c)/N_c$ ,  $p_f(\hat{v}_f) = \text{argmax}(\hat{v}_f)/N_f$ ,  $\lambda_1$  is  $2l$  and  $\lambda_2$  is  $2l/N_c$ . Because the target bin’s indicator operation is a non-differentiable operation, we choose to use binary cross-entropy (BCE) loss (see Eq. 6) during training. We use the  $\text{argmax}$  function to pick out the target bin during testing.

In our implementation, the intersection network is an MLP. The MLP’s head is a single linear layer; the MLP’s body is made up of 16 residual blocks, each of which is made up of two fully connected layers followed by ReLU activation; the MLP’s tail has two branches to produce  $\hat{\alpha}$ ,  $\hat{v}_c$  and  $\hat{v}_f$  (see Appendix for more details). We supervise this intersection network with the pre-trained NeRF models. Specifically, the distance from the camera to the first ray-object intersection can be measured by:

$$D(\mathbf{r}) = \sum_{i=1}^N w_i (t_i |\mathbf{d}|), t_i \in [t_n, t_f], \quad (5)$$

where  $w_i$  is the weight of  $i$ -th sample along rays in NeRF’s render equation,  $\mathbf{d}$  is the direction of ray  $\mathbf{r}$ ,  $t_n$  and  $t_f$  are the near and far clip ranges of the camera, respectively. This depth is then converted to  $\mu$  using the inverse of Eq. 2 for supervision.

The total training loss for the intersection network can be formulated as:

$$\mathcal{L}_{NeDF} = \text{BCE}(\hat{v}_c, v_c) + \text{BCE}(\hat{v}_f, v_f) + 0.1 \times \text{BCE}(\hat{\alpha}, \alpha), \quad (6)$$

where  $v_c, v_f$  are one-hot vectors of the same dimension as  $\hat{v}_c, \hat{v}_f$ . They can be obtained by applying segment operator  $\mathcal{S}$  on  $\mu$  generated by NeRF.  $\alpha$  is the ground truth mask generated by NeRF.

### Fast NeRF Composition with NeDF

Built on NeDF, we are able to achieve a fast composition and rendering framework. We illustrate the structure and data processing flow of this framework on the figure in the right of Fig. 1. It generally consists of three steps. First, the NeDF generation step takes all of the rays cast from the camera (eye position) through each pixel of the frame as input to

output a depth,  $D(\mathbf{r})$ , and object ID buffers. It saves the corresponding depth and object identification information for each element of  $D(\mathbf{r})$ . Deferred Shading step, as the second step, takes the ID and depth buffers outputted by the previous NeDF Generation step and the rays  $\mathbf{r}$  as an input. Then, it outputs a color buffer that holds the raw color (without shadow) information about the pixels. Lastly, the Shadow step uses the shadow rays,  $D(\mathbf{r})$ , and  $\mathbf{r}$  to generate the shadow information in a shadow buffer output. The final rendering result is produced by multiplying the color buffer and shadow buffer. Next, we first describe how the method handles the affine transformation of each NeRF object, as well as the necessary details for rendering acceleration.

**Manipulating Objects.** Given a 3D vector  $\mathbf{v}$  in world space and a corresponding 3D vector  $\mathbf{q}$  in canonical space, we define the rigid transformation of a NeRF object as a mapping function  $\mathcal{G} : \mathbf{q}(\mathbf{R}, \mathbf{T}, s) \mapsto \mathbf{v}$  where  $\mathbf{R}, \mathbf{T}, s$  specify the rotation, translation, and scale, respectively. Previous NeRF works focusing on deformation such as (Pumarola et al. 2021; Park et al. 2021a,b; Ost et al. 2021) transform the sampling points to canonical space to achieve the transformation. We instead directly transform the starting point and direction of rays to local canonical space and use this transformed rays to query the implicit field. As the intersection network takes rays  $\mathbf{r}$  as input, the depth  $\hat{D}(\mathbf{r})$  in global space can be computed using the inverse of the affine transformation function  $\mathcal{G}$ :

$$\hat{D}(\mathbf{r}) = |\mathbf{o} - \mathbf{p}_\perp| - s (\mathcal{S}^{-1} \circ \mathcal{F}_\Phi) (\mathcal{G}^{-1}(\mathbf{r})). \quad (7)$$

**Deferred Shading.** We term the step generating the raw color information as Deferred Shading step because each pixel only requires one  $\mathcal{F}_\Theta$  to shade in our framework. As the intersection network can accurately estimate the depth information, we are able to replace the ray marching step in previous NeRF works with a degenerated ray casting form, i.e., a pixel color can be acquired by a single evaluation of  $\mathcal{F}_\Theta$ . To further shade all objects in a lighting-aware manner, it is necessary to handle different lighting conditions in the original space from which the objects are collected, which is a non-trivial task. We simplify this processing by only shading with the primary color. Specifically, we obtain each ray’s color by:

$$c, \sigma = \mathcal{F}_\Theta (\mathcal{G}^{-1}(\mathbf{o} + \hat{D}(\mathbf{r})\mathbf{d}), \mathcal{G}^{-1}(\mathbf{d})), \quad (8)$$

where color  $c$  at the ray-object intersection point is directly used to approximate the ray’s color  $\hat{C}(\mathbf{r})$ . Density  $\sigma$  is employed to check whether the estimated  $\hat{D}$  is convincible. An extremely small  $\sigma$  indicates that  $\hat{D}$  does not reach the object surface. In this situation, we set a threshold that optionally assists in resampling these outliers using primary volume render equation in NeRF for better quality but a slight speed penalty (see our experimental results in Table 2). Although replacing the raymarching step with a single sample limits its usefulness, it is adequate for previewing. Users may prefer interactive-level speed over perfect quality when manipulating each object. In addition, once the adjustment is confirmed, the pipeline can re-render the scene and generate a set of high-quality results using only the NeRF models.

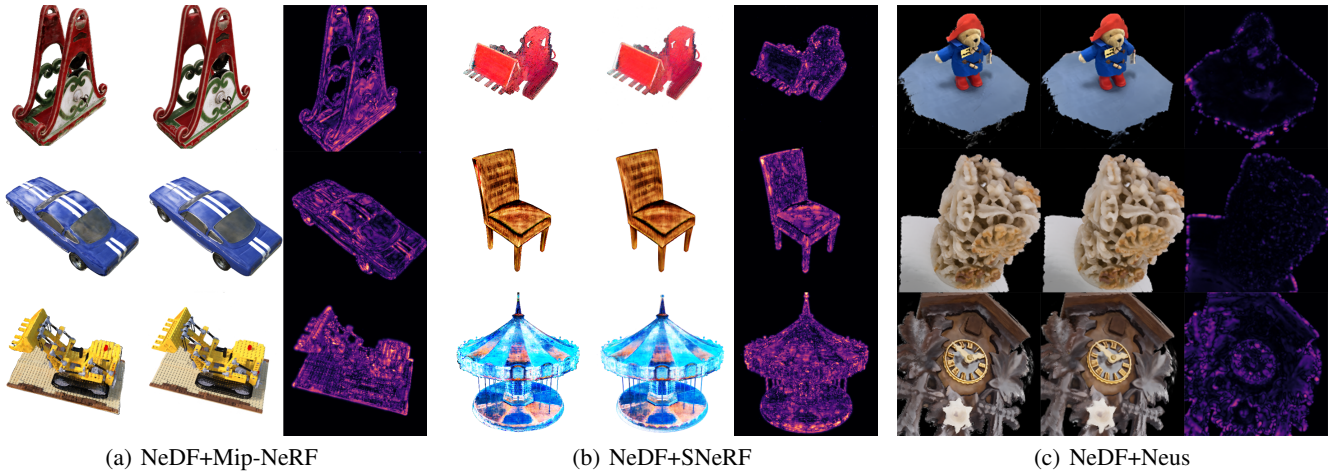


Figure 4: Visual quality illustration of the general plugin. For each method, the left column shows the previewing result, the middle column shows the naive result, and the right column visualizes the flip errors (Andersson et al. 2020).

	NeDF+Mip-NeRF vs. Mip-NeRF		NeDF+SNeRF vs. SNeRF		NeDF+Neus vs. Neus	
	time (s)	s-time(%) $\uparrow$	time (s)	s-time(%)	time (s)	s-time(%)
Outlier Processing (rs)	0.417/15.677	97.34	0.519/15.528	96.66	2.11/21.72	90.29
w/o Outlier Processing (wo rs)	0.234/15.677	98.51	0.225/15.528	98.55	0.202/21.72	99.07

Table 1: Our proposed framework serves as a plugin for three representative NeRF models: Mip-NeRF, SNeRF, and Neus. Term "s-time" refers to the percentage of time that our method saves. The optional outlier processing in Eq. 8 provides better quality but a slight speed penalty. The resolution is set to  $800 \times 800$ .

### Fast Dynamic shadows Casting

Shadows for NeRF models can be generated using the traditional volume rendering technique developed by (Max 1995). It, however, requires an additional volumetric evaluation for each sample point along a ray to determine whether the point is in shadows, which significantly increases the time consumption. Given that target objects have opaque in our cases, we generate the dynamic shadows based on the fundamental concept of shadow mapping (Williams 1978) with NeDF for quick preview. We detail this dynamic shadow casting as an additional step (Shadow Step, also named step 3) in appendix.

### Experiments

In this section, we first demonstrate that our pipeline indeed takes effect for various NeRF works. Following that, we show how our framework can assemble a large number of neural objects to form a virtual 3D world. Furthermore, we show our approach can be combined with traditional rendering techniques to create a mixed rendering pipeline. We implement our framework in two versions: Pytorch version for a convenient reproduction and CUDA version for testing performance limitation.

**Training details.** A single Nvidia A100 GPU is used to train NeDF for each scene from our  $N$ -object dataset (see section 'N-object dataset testing' below). We employ Adam as the optimizer and set the learning rate to  $5e-4$ . We use

the pre-trained NeRF models to generate 500 random views (more random views, less artifacts) for supervision for each NeDF model of a single object, and train the intersection network over 60W iterations until convergence. The batchsize of training rays is set to 4,096 during each iteration.

### General Acceleration Plugin

To evaluate the efficacy of the proposed framework acting as a general pipeline, we conduct experiments on the pre-trained models of certain classic implicit NeRFs. Three representative implicit neural representation based NeRF methods, i.e., Mip-NeRF (Barron et al. 2021), SNeRF (Nguyen-Phuoc, Liu, and Xiao 2022), and Neus (Wang et al. 2021), are employed for this set of experiments. Mip-NeRF represents the kind of NeRF methods that focus on novel view synthesis and rendering. SNeRF is a typical NeRF method that supports appearance style transfer, while Neus is a representative for NeRF methods designed for task of geometry reconstruction. We integrate these three representative NeRF techniques on our dataset and show their corresponding quantitative performance results in Table 1 and visual results in Fig. 4. Evidently, we can see that these representative NeRF methods, when used in conjunction with our proposed framework, significantly reduce rendering time while maintaining an acceptable level of quality degradation.

	air_balloons	idr_chair	nerf_chair	car	lego4pi	carousel	carousel_hors	cherryBlossom	ferrise-base	ferrise-wheel	ferrise-bin	plant	basketball	football	wood_chair	wood_table	wood_floor	ground1	ground2
b3															1	1	1		
b10		1			1			1				1			1	1		4	
b20		1	1	1	1	1	5	1				3			2	1		4	
b50	3	2	1	1	1	1	8	2	2	2	10	3			2	1	1	4	6

Figure 5: Number of objects contained in  $b3, b10, b20, b50$  scene. The number in two ball cases is empty, which is because we use them for the capability pressure test and repeated basketballs and footballs (sharing fifty-fifty) in the 100-ball and 1k-ball scenes.

## Storage and Speed

As shown in SNeRF (the same architecture with vanilla NeRF) column in Table 1, our framework can speed up rendering more than  $30 \sim 40$  times over the vanilla NeRF. For each model, our NeDF model takes 10MB (except 3MB taken by NeRF).

**Comparison with other solutions.** Although there’s no general pipeline for NeRF composition, we force a comparison with NSVF (Liu et al. 2020), a hybrid representation based NeRF, which speed up rendering of  $10 \sim 20$  times over the vanilla NeRF but gained a higher quality than us. To manifest performance of storage cost we compare our framework with DOnERF (3.6MB) (Neff et al. 2021) and CC-NeRF (1.5MB) (Tang et al. 2022). The experiments show that NSVF can produce higher image quality than our framework, but our framework can save time by more than 2 times over NSVF. As for the the storage cost, CC-NeRF consumes 1.5MB of storage, which is smaller than our model. However, our framework has several significant advantages that the other approaches (NSVF and CC-NeRF) lack, including: 1) unlike to our framework, neither NSVF nor CC-NeRF can be used as a general acceleration plugin and 2) NSVF and CC-NeRF cannot achieve interactive-level manipulation speeds or generate shadows like our method and 3) both of CC-NeRF and NSVF incur additional memory cost when rendering repeated objects while our framework avoids this issue.

## Capability evaluation

**$N$ -object dataset testing.** First, we collected an  $N$ -object dataset consisting of 22 distinct NeRF objects: 7 objects were chosen from previous work (3 from Neus’ dataset of real-world objects, 4 from synthesis dataset (Mildenhall et al. 2020; Zhang et al. 2022; Verbin et al. 2022)), and the remaining 15 objects were newly created by Blender 3D. To evaluate the framework’s ability to composite complex virtual scenes (see Fig. 6), we conduct a set of virtual scene compositions with four-level increasing complexity:  $b3, b10, b20, b50$ . The  $b3$  and  $b10$  are two static scenes with lower complexity, and the  $b20$  and  $b50$  are dynamic scenes where some objects move according to user-defined transformation settings. We allow to insert the same NeRF ob-




Figure 6: Results on our  $N$ -object dataset. Each 4-image pair shows: ID map, depth map, results with shadows using our previewing framework, and results with only naive NeRFs (shadowless). The Rightside of last two rows is the flip error (Andersson et al. 2020) between previewing and NeRF’s naive composition results). From top to bottom:  $b3$  and  $b10$ ,  $b20$  and  $b50$ , 100-ball and 1k-ball virtual scenes. Scenes  $b20$  and  $b50$  are both dynamic, with carousels and ferris wheels moving through user control settings; scenes 100-ball and 1k-ball have 100 and 1,000 balls (half basketballs and half footballs), respectively.

ject instances repeatedly into the scene to increase the complexity of the virtual scene. For each scene, we set different camera views (such as the views in Fig. 6) and render 120 frames to evaluate the average performance. Table 2 shows the consumed time and quality of each composition. Figure 7 depicts the time consumption ratio of each rendering step (see appendix) in rendering a single frame. As complexity increases the resampling ratio increases, which results in significant differences in computational time across different scenes in the *Deferred Shading step*. To test the capability upper bound of our framework, we also conduct a pressure test on virtual scenes with 100 and 1k balls to see how stable the proposed approach is for scene composition and shadow casting (see Fig. 6). Despite taking a longer time (see Tab. 2), our pipeline delivers stable performance and does not incur additional memory cost for those repeated balls.

**CUDA implementation.** CUDA version of the proposed framework is built with Vulkan, TensorRT, and a customized CUDA kernel to enable real-time manipulation on Nvidia RTX 4090 GPU. With this, we enable the user to construct a virtual scene in a progressive and interactive manner. The real-time manipulation benefits from a buffer reuse tech-

		b3-scene		b10-scene		b20-scene		b50-scene		100-ball		1k-ball	
		time(s)	PSNR	time	PSNR	time	PSNR	time	PSNR	time	PSNR	time	PSNR
wo s	rs	0.240	25.78	0.658	23.32	1.678	20.02	1.881	18.67	1.887	26.54	16.63	25.26
	wo rs	0.209	23.97	0.443	19.92	0.873	15.26	1.209	15.43	1.741	24.19	16.34	22.58
shadow	rs	0.398	\	1.040	\	2.441	\	2.943	\	3.073	\	28.97	\
	wo rs	0.367	\	0.813	\	1.628	\	2.286	\	3.002	\	27.66	\

Table 2: Performance of the implemented Pytorch framework on an  $N$ -object dataset ('wo s' means no shadow is calculated; 'rs' stands for resampling, i.e., outlier processing in Table 1). The resolution is set to  $900 \times 600$ . We use the naive composition (right side in Fig. 6) by NeRF as ground truth for composition without shadows. We only test the average time (in seconds) in each frame for shadow composition. It should be noted that the quality of the 1K-ball pressure testing was superior to the former due to the simpler geometry of the objects. Nonetheless, due to excessive overlapping between balls, the time increases rapidly.

nique to accelerate the performance during manipulation. Specifically, when the user manipulates an object, the other objects' rendering processes are frozen, and their buffers are reused to decrease the computational load.

### Compatibility with traditional graphics pipelines

The proposed framework is also suitable for integration with traditional rendering pipelines to build a mixed pipeline that supports dynamic shadow casting across meshes and neural objects in a single scene (see Fig. 8). Specifically, we first use a shared camera & light settings across mesh and neural scene to generate rendering buffers. Subsequently, these two scenes are seamlessly integrated using depth-based composition. Finally, shadow mapping is applied to produced a mixed result.

## Conclusions and Limitations

We present a general framework for a fast composition of NeRF objects with dynamic shadows. The core technical contribution is the use of Neural Depth Fields to allow for the direct intersection of rays and implicit surfaces to quickly determine the spatial relationship between objects as well as surface color computation of NeRF objects. To the best of our knowledge, this is the first framework that allows for the rapid creation of a virtual scene composed of a large number of NeRF objects with dynamic shadows. Our method can also be combined with traditional renderers to form a mixed pipeline, increasing its versatility and applicability.

The proposed framework is currently limited to solid surface objects, and predicting weight distribution in local space rather than a single depth may be a better option. In the future, this compositing framework can be further enhanced with more powerful editable NeRF works (Martin-Brualla et al. 2021; Tancik et al. 2022; Nguyen-Phuoc, Liu, and Xiao 2022; Zhang et al. 2023; Tojo and Umetani 2022; Zhang et al. 2021, 2022; Verbin et al. 2022) to achieve global illumination and improved visual quality.

## References

Andersson, P.; Nilsson, J.; Akenine-Möller, T.; Oskarsson, M.; Åström, K.; and Fairchild, M. D. 2020. FLIP: A difference evaluator for alternating images. *Proceedings of*




Figure 7: The time consumption ratio of our framework for the main steps. *Green* and *Red* bars represent the results with and without outliers resampling, respectively. Regarding the green bar, refer to this figure in conjunction with Fig. 6, when the camera view varies, certain objects which occupy a larger portion of the frame, e.g., the trees, have a higher resampling ratio and thus increase the processing time. Focusing attention on the red bar, the Step 2 (deferred shading step) consistently assumes a relatively minor temporal allocation across all scenarios. This is because, regardless of the number of objects in the scene, a pixel will be shaded only once with a single NeRF object.

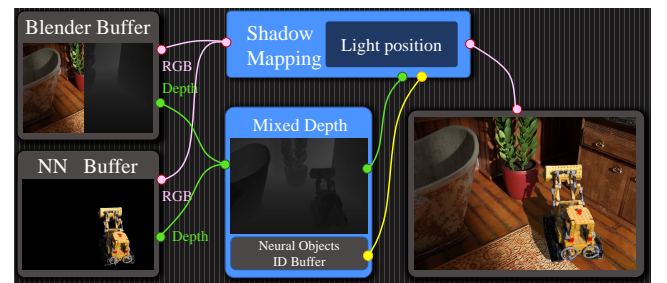


Figure 8: Given the information of camera view and light position, the buffer produced by the proposed framework (NN Buffer) and the *Cycles* engine of Blender 3D can be combined to form a mixed render result (see our video for a dynamic result). The *Bathroom* scene is from (Huang et al. 2022a).

- the ACM on Computer Graphics and Interactive Techniques*, 3(2): 15:1–15:23.
- Attal, B.; Huang, J.-B.; Zollhöfer, M.; Kopf, J.; and Kim, C. 2022. Learning neural light fields with ray-space embedding networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19819–19829.
- Barron, J. T.; Mildenhall, B.; Tancik, M.; Hedman, P.; Martin-Brualla, R.; and Srinivasan, P. P. 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5855–5864.
- Boss, M.; Braun, R.; Jampani, V.; Barron, J. T.; Liu, C.; and Lensch, H. 2021a. NerD: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12684–12694.
- Boss, M.; Engelhardt, A.; Kar, A.; Li, Y.; Sun, D.; Barron, J.; Lensch, H.; and Jampani, V. 2022. Samurai: Shape and material from unconstrained real-world arbitrary image collections. *Advances in Neural Information Processing Systems*, 35: 26389–26403.
- Boss, M.; Jampani, V.; Braun, R.; Liu, C.; Barron, J.; and Lensch, H. 2021b. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. *Advances in Neural Information Processing Systems*, 34: 10691–10704.
- Cao, J.; Wang, H.; Chemerys, P.; Shakhrai, V.; Hu, J.; Fu, Y.; Makoviichuk, D.; Tulyakov, S.; and Ren, J. 2022. Real-time neural light field on mobile devices. *arXiv preprint arXiv:2212.08057*.
- Chen, Z.; Funkhouser, T.; Hedman, P.; and Tagliasacchi, A. 2022. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. *arXiv preprint arXiv:2208.00277*.
- Dansereau, D.; and Bruton, L. 2004. Gradient-based depth estimation from 4D light fields. In *2004 IEEE International Symposium on Circuits and Systems (IEEE Cat. No.04CH37512)*, volume 3, III–549.
- Feng, B. Y.; and Varshney, A. 2021. Signet: Efficient neural representation for light fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14224–14233.
- Garbin, S. J.; Kowalski, M.; Johnson, M.; Shotton, J.; and Valentin, J. 2021. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14346–14355.
- Gu, J.; Liu, L.; Wang, P.; and Theobalt, C. 2022. StyleNeRF: A style-based 3D aware generator for high-resolution image synthesis. In *International Conference on Learning Representations*.
- Guo, M.; Fathi, A.; Wu, J.; and Funkhouser, T. 2020. Object-centric neural scene rendering. *arXiv preprint arXiv:2012.08503*.
- Hedman, P.; Srinivasan, P. P.; Mildenhall, B.; Barron, J. T.; and Debevec, P. 2021. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5875–5884.
- Huang, X.; Zhang, Q.; Feng, Y.; Li, H.; Wang, X.; and Wang, Q. 2022a. Hdr-nerf: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18398–18408.
- Huang, Y.-H.; He, Y.; Yuan, Y.-J.; Lai, Y.-K.; and Gao, L. 2022b. StylizedNeRF: Consistent 3D scene stylization as stylized NeRF via 2D-3D mutual learning. In *Computer Vision and Pattern Recognition (CVPR)*, 18342–18352.
- Jiang, K.; Chen, S.-Y.; Liu, F.-L.; Fu, H.; and Gao, L. 2022. NeRFFaceEditing: Disentangled face editing in neural radiance fields. In *SIGGRAPH Asia 2022 Conference Papers*, 1–9.
- Kim, C.; Zimmer, H.; Pritch, Y.; Sorkine-Hornung, A.; and Gross, M. 2013. Scene reconstruction from high spatio-angular resolution light fields. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 32(4): 73:1–73:12.
- Knodt, J.; Bartusek, J.; Baek, S.-H.; and Heide, F. 2021. Neural ray-tracing: Learning surfaces and reflectance for relighting and view synthesis. *arXiv preprint arXiv:2104.13562*.
- Krishnan, A.; Raj, A.; Zhang, X.; Carlson, A.; Tseng, N.; Sridhar, S.; Jaipuria, N.; and Hays, J. 2023. LANE: Lighting-aware neural fields for compositional scene synthesis. *arXiv:2304.03280*.
- Kundu, A.; Genova, K.; Yin, X.; Fathi, A.; Pantofaru, C.; Guibas, L. J.; Tagliasacchi, A.; Dellaert, F.; and Funkhouser, T. 2022. Panoptic neural fields: A semantic object-aware neural scene representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12871–12881.
- Kurz, A.; Neff, T.; Lv, Z.; Zollhöfer, M.; and Steinberger, M. 2022. AdaNeRF: Adaptive sampling for real-time rendering of neural radiance fields. In *European Conference on Computer Vision*, 254–270. Springer.
- Liu, L.; Gu, J.; Zaw Lin, K.; Chua, T.-S.; and Theobalt, C. 2020. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33: 15651–15663.
- Ma, T.; Li, B.; He, Q.; Dong, J.; and Tan, T. 2023. Semantic 3D-Aware Portrait Synthesis and Manipulation Based on Compositional Neural Radiance Field. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(2): 1878–1886.
- Martel, J. N. P.; Lindell, D. B.; Lin, C. Z.; Chan, E. R.; Monteiro, M.; and Wetzstein, G. 2021. ACORN: Adaptive coordinate networks for neural scene representation. *ACM Transactions on Graphics (SIGGRAPH)*, 40(4).
- Martin-Brualla, R.; Radwan, N.; Sajjadi, M. S.; Barron, J. T.; Dosovitskiy, A.; and Duckworth, D. 2021. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7210–7219.
- Max, N. 1995. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2): 99–108.



- Meng, N.; So, H. K.-H.; Sun, X.; and Lam, E. Y. 2019. High-dimensional dense residual convolutional neural network for light field reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(3): 873–886.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 405–421.
- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4): 102:1–102:15.
- Neff, T.; Stadlbauer, P.; Parger, M.; Kurz, A.; Mueller, J. H.; Chaitanya, C. R. A.; Kaplanyan, A.; and Steinberger, M. 2021. DONeRF: Towards real-time rendering of compact neural radiance fields using depth oracle networks. *Computer Graphics Forum*, 40(4): 45–59.
- Nguyen-Phuoc, T.; Liu, F.; and Xiao, L. 2022. SNeRF: Stylized Neural Implicit Representations for 3D Scenes. *ACM Transactions on Graphics*, 41(4).
- Oechsle, M.; Peng, S.; and Geiger, A. 2021. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5589–5599.
- Ost, J.; Mannan, F.; Thurey, N.; Knodt, J.; and Heide, F. 2021. Neural scene graphs for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2856–2865.
- Park, K.; Sinha, U.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Seitz, S. M.; and Martin-Brualla, R. 2021a. Nerfies: Deformable neural radiance fields. *ICCV*, 5865–5874.
- Park, K.; Sinha, U.; Hedman, P.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Martin-Brualla, R.; and Seitz, S. M. 2021b. HyperNeRF: A higher-dimensional representation for topologically varying neural radiance fields. *ACM Transactions on Graphics*, 40(6).
- Piala, M.; and Clark, R. 2021. Terminerf: Ray termination prediction for efficient neural rendering. In *2021 International Conference on 3D Vision (3DV)*, 1106–1114. IEEE.
- Pumarola, A.; Corona, E.; Pons-Moll, G.; and Moreno-Noguer, F. 2021. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10318–10327.
- Reiser, C.; Peng, S.; Liao, Y.; and Geiger, A. 2021. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14335–14345.
- Shuai, Q.; Geng, C.; Fang, Q.; Peng, S.; Shen, W.; Zhou, X.; and Bao, H. 2022. Novel view synthesis of human interactions from sparse multi-view videos. In *SIGGRAPH Conference Proceedings*, 1–10.
- Sitzmann, V.; Rezkikov, S.; Freeman, B.; Tenenbaum, J.; and Durand, F. 2021. Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems*, 34: 19313–19325.
- Srinivasan, P. P.; Deng, B.; Zhang, X.; Tancik, M.; Mildenhall, B.; and Barron, J. T. 2021. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7495–7504.
- Suhail, M.; Esteves, C.; Sigal, L.; and Makadia, A. 2022. Light field neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8269–8279.
- Sun, J.; Chen, X.; Wang, Q.; Li, Z.; Averbuch-Elor, H.; Zhou, X.; and Snavely, N. 2022a. Neural 3D reconstruction in the wild. In *ACM SIGGRAPH 2022 Conference Proceedings*, 1–9.
- Sun, J.; Wang, X.; Zhang, Y.; Li, X.; Zhang, Q.; Liu, Y.; and Wang, J. 2022b. Fenerf: Face editing in neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7672–7682.
- Tancik, M.; Casser, V.; Yan, X.; Pradhan, S.; Mildenhall, B.; Srinivasan, P. P.; Barron, J. T.; and Kretzschmar, H. 2022. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8248–8258.
- Tang, J.; Chen, X.; Wang, J.; and Zeng, G. 2022. Compressible-composable NeRF via rank-residual decomposition. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 14798–14809. Curran Associates, Inc.
- Tojo, K.; and Umetani, N. 2022. Recolorable posterization of volumetric radiance fields using visibility-weighted palette extraction. *Computer Graphics Forum*, 41(4): 149–160.
- Tosic, I.; and Berkner, K. 2014. Light field scale-depth space transform for dense depth estimation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 441–448.
- Verbin, D.; Hedman, P.; Mildenhall, B.; Zickler, T.; Barron, J. T.; and Srinivasan, P. P. 2022. Ref-NeRF: structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5481–5490. IEEE.
- Wang, D.; Chandran, P.; Zoss, G.; Bradley, D.; and Gotardo, P. 2022a. Morf: Morphable radiance fields for multiview neural head modeling. In *ACM SIGGRAPH 2022 Conference Proceedings*, 1–9.
- Wang, H.; Ren, J.; Huang, Z.; Olszewski, K.; Chai, M.; Fu, Y.; and Tulyakov, S. 2022b. R2l: Distilling neural radiance field to neural light field for efficient novel view synthesis. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, 612–629. Springer.
- Wang, P.; Liu, L.; Liu, Y.; Theobalt, C.; Komura, T.; and Wang, W. 2021. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 27171–27183.

- Williams, L. 1978. Casting curved shadows on curved surfaces. In *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques*, 270–274.
- Wizadwongsa, S.; Phongthawee, P.; Yenphraphai, J.; and Suwajanakorn, S. 2021. Nex: Real-time view synthesis with neural basis expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8534–8543.
- Wu, Q.; Liu, X.; Chen, Y.; Li, K.; Zheng, C.; Cai, J.; and Zheng, J. 2022. Object-compositional neural implicit surfaces. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVII*, 197–213. Springer.
- Yang, B.; Zhang, Y.; Xu, Y.; Li, Y.; Zhou, H.; Bao, H.; Zhang, G.; and Cui, Z. 2021. Learning object-compositional neural radiance field for editable scene rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13779–13788.
- Yariv, L.; Hedman, P.; Reiser, C.; Verbin, D.; Srinivasan, P. P.; Szeliski, R.; Barron, J. T.; and Mildenhall, B. 2023. BakedSDF: Meshing Neural SDFs for Real-Time View Synthesis. *arXiv*.
- Yu, A.; Li, R.; Tancik, M.; Li, H.; Ng, R.; and Kanazawa, A. 2021. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5752–5761.
- Zhang, X.; Srinivasan, P. P.; Deng, B.; Debevec, P.; Freeman, W. T.; and Barron, J. T. 2021. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)*, 40(6): 1–18.
- Zhang, Y.; He, Z.; Xing, J.; Yao, X.; and Jia, J. 2023. Ref-NPR: Reference-based non-photorealistic radiance fields for controllable scene stylization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4242–4251.
- Zhang, Y.; Sun, J.; He, X.; Fu, H.; Jia, R.; and Zhou, X. 2022. Modeling indirect illumination for inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18643–18652.

# Appendix: A General Implicit Framework for Fast NeRF Composition and Rendering

Xinyu Gao, Ziyi Yang, Yunlu Zhao, Yuxiang Sun, Xiaogang Jin, Changqing Zou

State Key Lab of CAD&CG, Zhejiang University

## Intersection Network

**Architecture.** The Intersection Network is built using an MLP (Fig. 1). The MLP’s head is a single linear layer that converts the input dimension  $D_i$  to the feature dimension  $D_f$ . The MLP’s body is made up of 16 residual blocks, each of which is made up of two fully connected layers followed by ReLU activation. The MLP’s tail has two branches, each with a single linear to fit output dimension. One branch produces  $\hat{\alpha}$  and  $\hat{v}_c$ , while the other produces  $\hat{v}_f$  (Equation 3 in the main paper).

**The choice for hyper-parameters.** In our experiment,  $D_i$  is 1,008 (16 points for a ray, each with 10-level sinusoidal position encoding, as in NeRF). We set  $D_f$  to 256 and use 16 residual blocks with skip-connection, following the same idea as Wang et al. (?), where residual blocks with skip-connection play an important role in deep networks. More residual blocks result in a slight improvement in quality but a significant decrease in speed, and 16 blocks is a good choice for trade-off in most cases.

**Training details.** We employ Adam as the optimizer and set the learning rate to  $5e-4$ . We use the pre-trained NeRF models to generate 500 random views (more random views, less artifacts) for supervision for each NeDF model of a single object, and train the intersection network over 60W iterations until convergence. The batchsize of training rays is set to 4,096 during each iteration.

## Algorithm Pseudo Code

We detailed algorithm 1 and 2 in pseudo-code corresponding to the three main steps in our pipeline.

## A brief Comparison Result

We force a brief comparison with NSVF (?) in caption of Fig. 2 (the most complex and dynamic scene from our dataset). Because NSVF and ours are two totally different system, it’s hard to set a simultaneous event and we only made a comparison in terms of speed.

---

### Algorithm 1: Fast NeRF Composition

---

**Input:**  $\mathcal{M}$ , all models in the view frustum. Each model has a unique id;  
**Output:** color buffer  $I_{rgb}$  of compositing results  
**NeDF Generation Step (STEP 1);** Output: depth buffer  $I_{depth}$ , id buffer  $I_{id}$ ;  
*clear depth buffer with  $\infty$ , reset id buffer;*  
**foreach** model  $m$  of  $\mathcal{M}$  **do**  
    use  $\mathcal{F}_\Phi$  of  $m$ ;  
    **foreach** ray cast through pixel  $i$  **do**  
        render depth  $d$  and mask  $\alpha$  by  $\mathcal{F}_\Phi$ ;  
        **if**  $d < I_{depth}^i$  and  $\alpha = 1$  **then**  
            write to  $I_{depth}^i$  with  $d$ ;  
            write to  $I_{id}^i$  with  $m.id$ ;  
        **end**  
    **end**  
**end**  
**Deferred Shading Step (STEP 2);**  
*clear color buffer with clear-color value;*  
**foreach** model  $m$  of  $\mathcal{M}$  **do**  
    use  $\mathcal{F}_\Theta$  of  $m$ ;  
    **foreach** pixel  $i$  where  $I_{id}^i == m.id$  **do**  
        project pixel  $i$  to a single 3D point  $\mathbf{x}_i$  using  $I_{depth}^i$ ;  
        render color at  $\mathbf{x}_i$  by  $\mathcal{F}_\Theta$  and write to  $I_{rgb}^i$ ;  
    **end**  
**end**

---

arXiv:2308.04669v2 [cs.CV] 14 Aug 2023

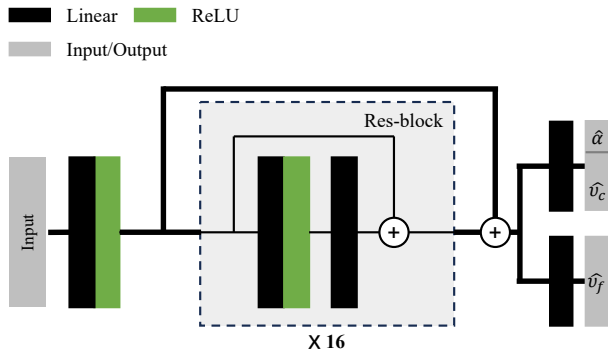


Figure 1: Architecture of Intersection Network.

---

**Algorithm 2:** Fast Dynamic Shadow Casting

---

**Input:**  $\mathcal{M}$ , all Models in view frustum. Each model has a unique id.  $L$ , position of a point light source;

**Output:** color buffer  $I_{rgb}$  with shadows

**Shadow Step (STEP 3):** generate shadow by shadow rays;

```

foreach pixel  $i$  with a valid  $I_{id}^i$  do
  project pixel  $i$  to a single 3D point  $\mathbf{x}_i$  using
     $I_{depth}^i$ ;
  cast shadow ray from  $L$  to  $\mathbf{x}_i$ ;
  depth  $D$  in shadow maps can be obtained like in
  STEP 1;
  if  $D + \epsilon < |L - \mathbf{x}_i|$  then
    multiply  $I_{rgb}^i$  with an intensity value
       $\beta(0 < \beta < 1)$ ;
  end

```

**end**

---




Figure 2: The above frame (with shadow) with a resolution of 2,000x800 was rendered in 4.48 seconds using our Pytorch implementation, whereas vanilla NeRF takes over 130 seconds and the hybrid representation based NeRF method NSVF (?) takes over 10 seconds (without the presence of a shadow). The above scene is dynamic, e.g., a carousel and a ferris wheel that both rotate slowly (see the demo video in the supplemental materials for more details).