

Seal-3D: Interactive Pixel-Level Editing for Neural Radiance Fields

Xiangyu Wang^{1*} Jingsen Zhu^{2*} Qi Ye^{1†} Yuchi Huo^{3,2} Yunlong Ran¹
Zhihua Zhong² Jiming Chen¹

¹Zhejiang University, Key Lab of CS&AUS of Zhejiang Province

²State Key Lab of CAD&CG, Zhejiang University ³Zhejiang Lab

{xy-wong, zhujingsen, qi.ye}@zju.edu.cn huo.yuchi.sc@gmail.com

{yunlong-ran, zhongzhihua, cjm}@zju.edu.cn




Figure 1: Seal-3D: The first interactive pixel level NeRF editing tool. We design an interactive user editing method and system *Seal-3D*, which achieves instant ($\approx 1s$) preview (left) by our novel pretraining strategy. High-quality editing results can be further obtained by a short period (in 1 or 2 minutes) of finetuning. The editing results of our implemented editing tools (right) are view-consistent with rich shading details (e.g. shadows) on the original surface (left).

Abstract

With the popularity of implicit neural representations, or neural radiance fields (NeRF), there is a pressing need for editing methods to interact with the implicit 3D models for tasks like post-processing reconstructed scenes and 3D content creation. While previous works have explored NeRF editing from various perspectives, they are restricted in editing flexibility, quality, and speed, failing to offer direct editing response and instant preview. The key challenge is to conceive a locally editable neural representation that can directly reflect the editing instructions and update instantly. To bridge the gap, we propose a new interactive editing method and system for implicit representations, called *Seal-3D*¹, which allows users to edit NeRF models in a pixel-level and free manner with a wide range of NeRF-

like backbone and preview the editing effects instantly. To achieve the effects, the challenges are addressed by our proposed proxy function mapping the editing instructions to the original space of NeRF models in the teacher model and a two-stage training strategy for the student model with local pretraining and global finetuning. A NeRF editing system is built to showcase various editing types. Our system can achieve compelling editing effects with an interactive speed of about 1 second.

1. Introduction

Implicit neural representations, e.g. neural radiance fields (NeRF) [24], have gained increasing attention as novel 3D representations with neural networks to model a 3D scene. Benefiting from the high reconstruction accuracy and rendering quality with relatively low memory consumption, NeRF and its variations [50, 3, 33, 26, 4, 45, 41] have demonstrated great potential in many 3D applications like 3D reconstruction, novel view synthesis, and Vir-

*Equal contribution.

†Corresponding author.

Project page: <https://windingwind.github.io/seal-3d/>

¹“Seal” derived from the name of rubber stamp in Adobe Photoshop.

tual/Augmented Reality.

With the popularity of the new implicit representations and an increasing number of implicit 3D models, there is a pressing demand for human-friendly editing tools to interact with these 3D models. Editing with implicit neural representations is a fundamental technique required to fully empower the representation. Objects reconstructed from the real world are likely to contain artifacts due to the noise of captured data and the limitations of the reconstruction algorithms. In a typical 3D scanning pipeline, manual correction and refinement to remove artifacts are common stages. On the other hand, in 3D content creation applications like 3D games, animations, and filming, artists usually need to create new content based on existing 3D models.

Prior works have made attempts to edit 3D scenes represented by NeRF, including object segmentation [20, 44], object removal [19], appearance editing [14, 27, 22], and object blending [7], *etc.* These existing NeRF editing methods mainly focus on coarse-grained object-level editing and the convergence speed can not meet the demands of interactive editing. Some recent methods [48, 5] transform the editing of NeRF into mesh editing by introducing a mesh as an edit proxy. This requires the user to operate on an additional meshing tool, which limits interactivity and user-friendliness. To the best of our knowledge, there are no existing methods that are able to support interactive pixel-level editing of neural radiance fields with fast converging speed, which is mainly due to the challenges discussed below.

Unlike existing explicit 3D representations *e.g.* point cloud, textured mesh, and occupancy volume, which store the explicit geometry structure of objects and scenes, implicit representations use neural networks to query features of a 3D scene including geometry and color. Existing 3D editing methods, taking the mesh-based representation as an example, can change object geometry by displacing vertices corresponding to target object surface areas and object textures. Without explicit explainable correspondence between the visual effects and the underlying representations, editing the implicit 3D models is indirect and challenging. Further, it is difficult to locate implicit network parameters in local areas of the scene, meaning that adaptations of the network parameters may lead to undesired global changes. This results in more challenges for fine-grained editing.

To bridge the gap, in this paper, we propose an interactive pixel-level editing method and system for implicit neural representations for 3D scenes, dubbed Seal-3D. The name is borrowed from the popular 2D image editing software Adobe PhotoShop [1], as its seal tool provides similar editing operations. As shown in Fig. 1, the editing system consists of five types of editing as examples: 1) Bounding box tool. It transforms and scales things inside a bounding box, like a copy-paste operation. 2) Brushing tool. It paints

specified color on the selected zone and can increase or decrease the surface height, like an oil paint brush or graver. 3) Anchor tool. It allows the user to freely move a control point and affect its neighbor space according to the user input. 4) Color tool. It edits the color of the object surfaces.

To achieve the interactive NeRF editing effects, we address the challenges of implicit representations discussed above. First, to establish the correspondence between the explicit editing instructions to the update of implicit network parameters, we propose a proxy function that maps the target 3D space (determined by the user edit instructions from an interactive GUI) to the original 3D scene space, and a teacher-student distillation strategy to update the parameters with the corresponding content supervision acquired by the proxy function from the original scenes. Second, to enable local editing, *i.e.* mitigating the influence of the local editing effect on the global 3D scenes under the non-local implicit representations, we propose a two-stage training process: a pretraining stage of updating only the positional embedding grids with local losses for editing areas while freezing the subsequent MLP decoder to prevent global degeneration, and a finetuning stage of updating both the embedding grids and the MLP decoder with global photometric losses. With this design, the pretraining stage updates local editing features and the finetuning stage blends the local editing areas with global structures and colors of unedited space to achieve view consistency. This design has the benefit of an instant preview of the editing: the pretraining can converge very fast and presents local editing effects within approximately 1 second only.

In summary, our contributions are as follows:

- We propose the first interactive pixel-level editing method and system for neural radiance fields, which exemplifies fine-grained multiple types of editing tools, including geometry (bounding box tool, brush tool, and anchor tool) and color edits;
- A proxy function is proposed to establish the correspondence between the explicit editing instructions and the update of implicit network parameters and a teacher-student distillation strategy is proposed to update the parameters;
- A two-stage training strategy is proposed to enable instant preview of local fine-grained editing without contaminating the global 3D scenes.

2. Related Work

Novel view synthesis. Given a set of posed image captures of a scene, the task of novel view synthesis is to generate photo-realistic images from arbitrary novel views. Recently, neural network have been introduced into the

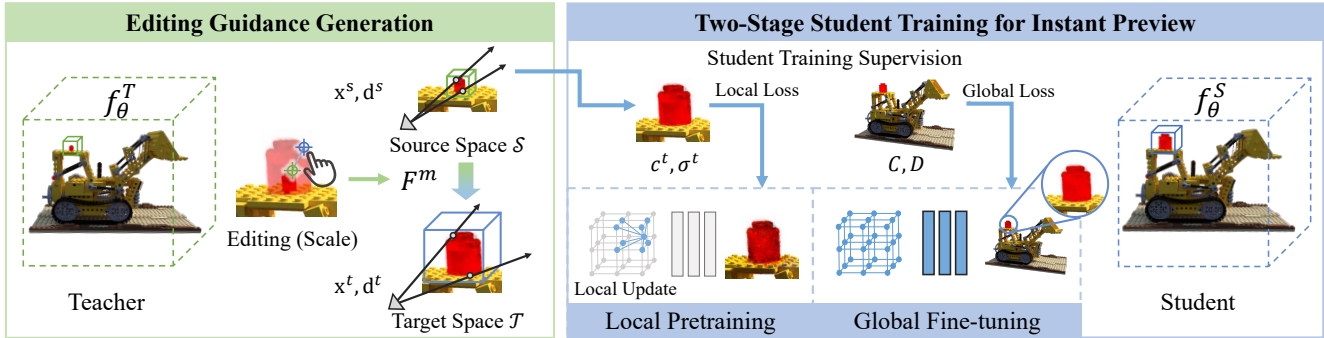


Figure 2: Illustration of the editing framework. Left: a 3D point and view direction from the target space after user editing is mapped to the original source space to get guidance c_t, σ_t from the teacher model f_{θ}^T for the student training. Right: the student training consists of two stages: fast pretraining to provide instant preview by updating partial parameters of the network with local losses and finetuning with global losses.

rendering pipeline and leveraged for multiple representations, such as voxels [21, 35], point clouds [2, 6], multi-plane images (MPIs) [17, 23, 52], and implicit representations [36, 24]. Typically, Neural radiance field (NeRF) [24] uses a single MLP to implicitly encode a scene into a volumetric field of density and color, and takes advantage of volume rendering to achieve impressive rendering results with view-dependent effects, which inspires a lot of follow-up works on human [30, 42], deformable objects [28, 29], pose estimations [18], autonomous system [32, 49], surface reconstruction [45, 41], indoor scenes [47], city [38, 43], *etc.* NeRF’s MLP representation can be enhanced and accelerated by hybrid representations, including voxels [33, 37], hashgrids [26] and tensorial decomposition [4, 40]. In this paper, our interactive editing framework is developed based on Instant-NGP [26], which achieve real-time rendering speed for NeRF inference and state-of-the-art quality of novel view synthesis.

Neural scene editing. Scene editing has been a widely researched problem in computer vision and graphics. Early method focus on editing a single static view by inserting [15, 54], relighting [16], composition [31], object moving [12, 34], *etc.* With the development of neural rendering, recent works attempt to perform editing at different levels of the 3D scene, which can be categorized as scene-level, object-level, and pixel-level editing. Scene-level editing methods focus on changing in global appearances of a scene, such as lighting [8] and global palette [14]. Intrinsic decomposition [51, 27, 9, 46, 53, 10] disentangles material and lighting field and enables texture or lighting editing. However, scene-level methods are only able to modify global attributes and are unable to apply to specified objects. Object-level editing methods use different strategies to manipulate the implicitly represented object. Object-NeRF [44] exploit per-object latent code to decompose neu-

ral radiance field into objects, enabling object moving, removal, or duplicating. Liu *et al.* [20] design a conditional radiance field model which is partially optimized according to the editing instructions to modify semantic-level color or geometry. NeRF-editing [48] and NeuMesh [5] introduce a deformable mesh reconstructed by NeRF, as an editing proxy to guide object editings. However, these methods are restricted to object-level rigid transformation or are not generalizable to arbitrary out-of-distribution editing categories. In contrast, pixel-level editing aims to provide fine-grained editing guidance precisely selected by pixels, instead of restricted by object entities. To the best of our knowledge, NeuMesh [5] is the only existing method that achieves editing at this level. However, it depends on the mesh scaffold, which limits the editing categories, *e.g.* cannot create out-of-mesh geometry structures. In contrast, our editing framework does not require any proxy geometry structures, allowing it to be more direct and extensive.

Besides, optimizing the performance of neural editing method remains an open problem. Existing methods require minutes or even hours of optimization and inference. Our method is the first pixel-level neural editing framework to achieve instant interactive (*i.e.* second-level) performance.

3. Method

We introduce Seal-3D, an interactive pixel-level editing method for neural radiance fields. The overall pipeline is illustrated in Fig. 2, which consists of a pixel-level proxy mapping function, a teacher-student training framework, and a two-stage training strategy for the student NeRF network under the framework. Our editing workflow starts with the proxy function which maps the query points and ray directions according to user-specified editing rules. Then a NeRF-to-NeRF teacher-student distillation framework follows, where a teacher model with editing mapping rules of geometry and color supervises the train-

ing of a student model (Sec. 3.2). The key to interactive fine-grained editing is the two-stage training for the student model (Sec. 3.3). In an extra pretraining stage, the points, ray directions, and inferred ground truth inside edit space from the teacher model are sampled, computed, and cached previously; only parameters with locality are updated and the parameters causing global changes are frozen. After the pretraining stage, the student model is finetuned with a global training stage.

3.1. Overview of NeRF-based Editing Problem

We first make a brief introduction to neural radiance fields and then analyze the challenges of NeRF-based editing problems and the limitations of existing solutions.

3.1.1 NeRF Preliminaries

Neural radiance fields (NeRFs) provide implicit representations for a 3D scene as a 5D function: $f : (x, y, z, \theta, \phi) \mapsto (c, \sigma)$, where $\mathbf{x} = (x, y, z)$ is a 3D location and $\mathbf{d} = (\theta, \phi)$ is the view direction, while c and σ denote color and volume density, respectively. The 5D function is typically parameterized as an MLP f_θ .

To render an image pixel, a ray \mathbf{r} with direction \mathbf{d} is shot from the camera position \mathbf{o} through the pixel center according to the intrinsics and extrinsics of the camera. K points $\mathbf{x}_i = \mathbf{o} + t_i \mathbf{d}, i = 1, 2, \dots, K$ are sampled along the ray, and the network f_θ is queried for their corresponding color and density:

$$(c_i, \sigma_i) = f_\theta(\mathbf{x}_i, \mathbf{d}) \quad (1)$$

Subsequently, the predicted pixel color $\hat{C}(\mathbf{r})$ and depth value $\hat{D}(\mathbf{r})$ are computed by volume rendering:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^K T_i \alpha_i c_i, \quad \hat{D}(\mathbf{r}) = \sum_{i=1}^K T_i \alpha_i t_i \quad (2)$$

$$T_i = \prod_{j<i} (1 - \alpha_j), \quad \alpha_i = 1 - \exp(-\sigma_i \delta_i) \quad (3)$$

where α_i is the alpha value for blending, T_i is the accumulated transmittance, and $\delta_i = t_{i+1} - t_i$ is the distance between adjacent points. NeRF is trained by minimizing the photometric loss between the predicted and ground truth color of pixels.

In this paper, we build our interactive NeRF editing system upon Instant-NGP [26], which achieves nearly real-time rendering performance for NeRF. Although our implementation of instant interactive editing relies on hybrid representations for NeRF to achieve the best speed performance, our proposed editing framework does not rely on a specific NeRF backbone and can be transplanted to other frameworks as long as they follow the aforementioned volume rendering pipeline.

3.1.2 Challenges of NeRF-based Editing

NeRF-like methods achieve the state-of-the-art quality of scene reconstruction. However, the 3D scene is implicitly represented by network parameters, which lacks interpretability and can hardly be manipulated. In terms of scene editing, it is difficult to find a mapping between the *explicit* editing instructions and the *implicit* update of network parameters. Previous works attempt to tackle this by means of several restricted approaches:

NeRF-Editing [48] and NeuMesh [5] introduce a mesh scaffold as a geometry proxy to assist the editing, which simplifies the NeRF editing task into mesh modification. Although conforming with existing mesh-based editing, the editing process requires extracting an additional mesh, which is cumbersome. In addition, the edited geometry is highly dependent on the mesh proxy structure, making it difficult to edit spaces that are not easy or able to be represented by meshes while representing these spaces is one key feature of the implicit representations. Liu *et al.* [20] designs additional color and shape losses to supervise the editing. However, their designed losses are only in 2D photometric space, which limits the editing capability of a 3D NeRF model. Furthermore, their method only supports editing of semantic-continuous geometry in simple objects, instead of arbitrary pixel-level complex editing.

Moreover, to the best of our knowledge, existing methods have not realized interactive editing performance considering both quality and speed. Liu *et al.* [20] is the only existing method that completes optimization within a minute (37.4s according to their paper), but their method only supports extremely simple objects and does not support fine-grained local edits (see Fig. 10 for details). Other editing methods (*e.g.* NeuMesh [5]) usually require hours of network optimization to obtain edit results.

In this paper, we implement an interactive pixel-level editing system, which can be extended to new editing types easily using similar editing strategies as the traditional explicit 3D representation editing. Our method does not require any explicit proxy structure (instead, a proxy function, see Sec. 3.2) and can define various pixel-level editing effects without an explicit geometry proxy. It also enables *instant preview* ($\approx 1s$) (see Sec. 3.3). Tab. 1 compares the edit capabilities between our method and previous methods.

Method	w/o Explicit Proxy	Pixel-Level	Interactive	Time
Ours	✓	✓	✓	seconds
NeuMesh [5]	✗	(partial)	✗	hours
NeRF-Editing [48]	✗	✗	✗	hours

Table 1: **Comparison with recent methods in edit capabilities.** Our method supports arbitrary editing, does not require any explicit geometry proxy, and achieves interactive editing in seconds.

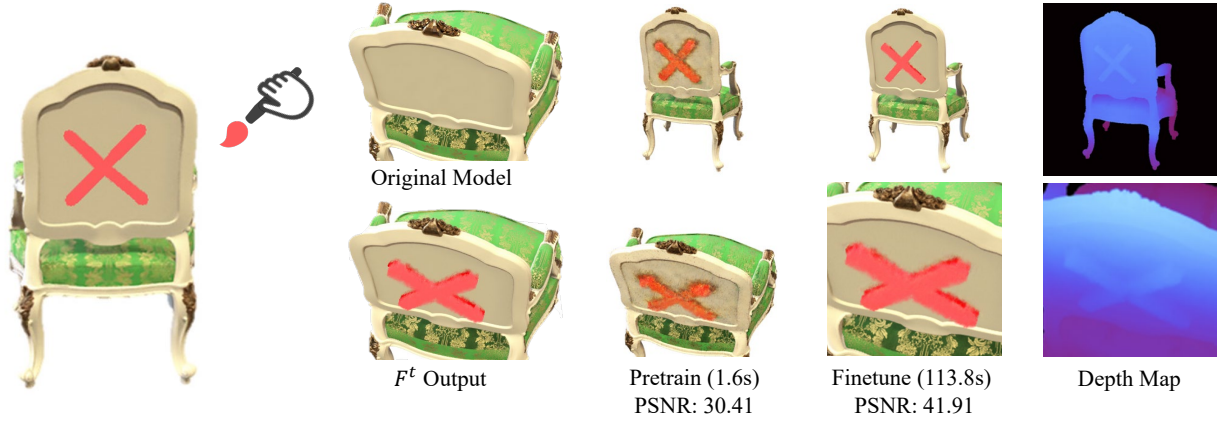


Figure 3: Example of brush editing: 3D painting with color and thickness.

3.2. Editing Guidance Generation

Our design implements NeRF editing as a process of knowledge distillation. Given a pretrained NeRF network fitting a particular scene that serves as a teacher network, we initialize an extra NeRF network with the pretrained weights as a student network. The teacher network f_{θ}^T generates editing guidance from the editing instructions input by the user, while the student network f_{θ}^S is optimized by distilling editing knowledge from the editing guidance output by the teacher network. In the subsection, editing guidance generation for the student model supervision is introduced and illustrated on the left of Fig. 2.

Firstly, the user edit instructions are read from the interactive NeRF editor as pixel-level information. The source space $\mathcal{S} \subset \mathbb{R}^3$ is the 3D space for the original NeRF model and the target space $\mathcal{T} \subset \mathbb{R}^3$ is the 3D space for the NeRF model after editing. The target space \mathcal{T} is warped to the original space \mathcal{S} by $F^m: \mathcal{T} \mapsto \mathcal{S}$. F^m transforms points within the target space and their associated directions according to editing rules which are exemplified below. With the function, the “pseudo” desired edited effects c^T, σ^T for each 3D point and view direction in the target space can be acquired by querying the teacher NeRF model f_{θ}^T : the transformed points and directions (in source space) are fed into the teacher network get the color and density. The process can be expressed as

$$\mathbf{x}^s, \mathbf{d}^s = F^m(\mathbf{x}^t, \mathbf{d}^t), \mathbf{x}^s \in \mathcal{S}, \mathbf{x}^t \in \mathcal{T}, \quad (4)$$

$$c^T, \sigma^T = f_{\theta}^T(\mathbf{x}^s, \mathbf{d}^s) \quad (5)$$

Where $\mathbf{x}^s, \mathbf{d}^s$ denotes source space point position and direction and $\mathbf{x}^t, \mathbf{d}^t$ denotes target space point position and direction.

For brevity, we define the entire process as *teacher inference process* $F^t := f_{\theta}^T \circ F^m: (\mathbf{x}^t, \mathbf{d}^t) \mapsto (c^T, \sigma^T)$. The inference result c^T, σ^T mimics the edited scene and acts as the teacher label, the information of which is then distilled

by the student network in the network optimization stage.

The mapping rules of F^m can be designed according to arbitrary editing targets. In particular, we implement 4 types of editing as examples.

- Bounding shape tool, which supports common features in traditional 3D editing software including copy-paste, rotation, and resizing. The user provides a bounding shape to indicate the original space \mathcal{S} to be edited and rotates, translates, and scales the bounding box to indicate the target effects. The target space \mathcal{T} and mapping function F^m are then parsed by our interface

$$\mathbf{x}^s = S^{-1} \cdot R^T \cdot (\mathbf{x}^t - \mathbf{c}^t) + \mathbf{c}^s,$$

$$\mathbf{d}^s = R^T \cdot \mathbf{d}^t$$

$$F^m := (\mathbf{x}^t, \mathbf{d}^t) \mapsto \begin{cases} (\mathbf{x}^s, \mathbf{d}^s) & , \text{if } \mathbf{x}^t \in \mathcal{T} \\ (\mathbf{x}^t, \mathbf{d}^t) & , \text{otherwise} \end{cases}$$

where R is rotation, S is scale, and $\mathbf{c}^s, \mathbf{c}^t$ are the center of \mathcal{S}, \mathcal{T} , respectively.

With this tool, we even support cross-scene object transfer, which can be implemented by introducing the NeRF of the transferred object as an additional teacher network in charge of part of the teacher inference process within the target area. We give a result in Fig. 7.

- Brushing tool, similar to the sculpt brush in traditional 3D editing that lifts or descends the painted surface. The user scribbles with a brush and \mathcal{S} is generated by ray casting on brushed pixels. The brush normal \mathbf{n} , and pressure value $p(\cdot) \in [0, 1]$ are defined by user, which determines the mapping:

$$\mathbf{x}^s = \mathbf{x}^t - p(\mathbf{x}^t)\mathbf{n},$$

$$F^m := (\mathbf{x}^t, \mathbf{d}^t) \mapsto (\mathbf{x}^s, \mathbf{d}^t)$$

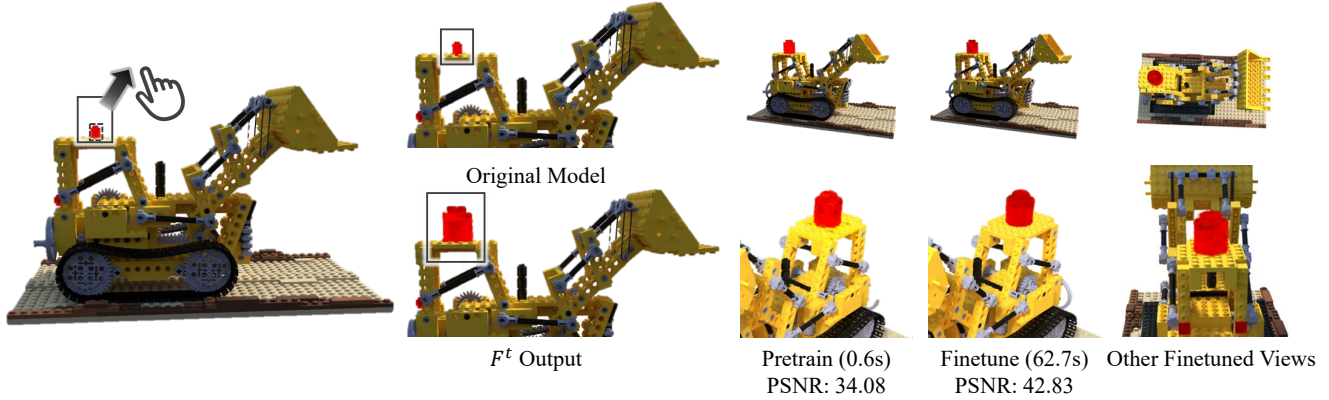


Figure 4: Example of bounding shape editing: bulb scaling.

- Anchor tool, where the user defines a control point \mathbf{x}^c and a translation vector \mathbf{t} . The region surrounding \mathbf{x}^c will be stretched by a translation function $\text{stretch}(\cdot; \mathbf{x}^c, \mathbf{t})$. Then the mapping is its inverse:

$$\begin{aligned} \mathbf{x}^s &= \text{stretch}^{-1}(\mathbf{x}^t; \mathbf{x}^c, \mathbf{t}) \\ F^m &:= (\mathbf{x}^t, \mathbf{d}^t) \rightarrow (\mathbf{x}^s, \mathbf{d}^t) \end{aligned}$$

please refer to the supplementary material for the explicit expressions of $\text{stretch}(\cdot; \mathbf{x}^c, \mathbf{t})$.

- Non-rigid transform tool, which allows for the accurate and flexible transformation of selected space.

$$\begin{aligned} \mathbf{x}^s &= R \cdot \mathbf{x}^t + t, \\ \mathbf{d}^s &= R \cdot \mathbf{d}^t, \\ F^m &:= (\mathbf{x}^t, \mathbf{d}^t) \mapsto (\mathbf{x}^s, \mathbf{d}^s) \end{aligned}$$

Where R, t are interpolated from the transform matrixes of the three closest coordinates of a pre-defined 3D blending control grid with position and transformation of each control point. The results can be found in Fig. 9.

- Color tool, which edits color via color space mapping (single color or texture). Here the spatial mapping is identical and we directly map the color output of the network to HSL space, which helps for color consistency. Our method is capable of preserving shading details (e.g. shadows) on the modified surface. We achieve this by transferring the luminance (in HSL space) offsets on the original surface color to the target surface color. Implementation details of this shading preservation strategy are presented in the supplementary.

For the training strategy of distillation, the student model f_θ^S is optimized with the supervision of pseudo ground

truths generated by the aforementioned teacher inference process F^t . The editing guidance from the teacher model is distilled into the student model by directly applying the photometric loss between pixel values \hat{C}, \hat{D} accumulated by Eq. (2) from the teacher and student inference.

However, we find that the convergence speed of this training process is slow ($\approx 30s$ or longer), which cannot meet the needs of instant preview. To tackle this problem, we design a two-stage training strategy: the first stage aims to converge instantly (within 1 second) so that a coarse editing result can be immediately presented to the user as a preview, while the second stage further finetunes the coarse preview to obtain a final refinement.

3.3. Two-stage Student Training for Instant Preview

Local pretraining for instant preview. Usually, the edit space is relatively small compared to the entire scene, so training on the global photometric loss is wasteful and leads to slow convergence. To achieve instant preview of editing, we adopt a local pretraining stage before the global training begins. The local pretraining process consists of: 1) uniformly sample a set $\mathcal{X} \subset \mathcal{T}$ of local points within the target space and a set \mathcal{D} of directions on the unit sphere, and feed them into the teacher inference process F^t to obtain teacher labels c^T, σ^T , and cache them in advance; 2) the student network is trained by local pretraining loss $\mathcal{L}_{\text{local}}$:

$$(c^T, \sigma^T) = F^t(\mathbf{x}, \mathbf{d}), (c^S, \sigma^S) = f_\theta^S(\mathbf{x}, \mathbf{d}), \quad (6)$$

$$\mathcal{L}_{\text{local}} = \sum_{\mathbf{x} \in \mathcal{X}, \mathbf{d} \in \mathcal{D}} \lambda_1 \|c^T - c^S\|_1 + \lambda_2 \|\sigma^T - \sigma^S\|_1 \quad (7)$$

where c^S, σ^S are the predicted color and density of sampled points $\mathbf{x} \in \mathcal{X}$ by the student network, and c^T, σ^T are cached teacher labels. This pretraining stage is very fast: after only about 1 second of optimization, the rendered image of the student network shows plausible color and shape consistent with the editing instructions.

However, training on only the local points in the editing area may lead to degeneration in other global areas unrelated to the editing due to the non-local implicit neural network. We observe the fact that in hybrid implicit representations (such as Instant-NGP [26]), local information is mainly stored in the positional embedding grids, while the subsequent MLP decodes global information. Therefore, in this stage, all parameters of the MLP decoder are frozen to prevent global degeneration. Experimental illustrations will be presented in Sec. 4.3 and Fig. 12.

Global Finetuning. After pretraining, we continue to finetune f_{θ}^S to refine the coarse preview to a fully converged result. This stage is similar to the standard NeRF training, except that the supervision labels are generated by the teacher inference process instead of image pixels.

$$\mathcal{L}_{\text{global}} = \sum_{\mathbf{r} \in \mathcal{R}} \lambda_3 \|\hat{C}^T - \hat{C}^S\|_2 + \lambda_4 \|\hat{D}^T - \hat{D}^S\|_1 \quad (8)$$

where \mathcal{R} denote the set of sampled rays in the minibatch and $(\hat{C}^T, \hat{D}^T), (\hat{C}^S, \hat{D}^S)$ are accumulated along ray \mathbf{r} by Eq. (2) according to $(c^T, \sigma^T), (c^S, \sigma^S)$, respectively.

It is worth mentioning that the student network is capable of generating results of better quality than the teacher network that it learns from. This is because the mapping operation in the teacher inference process may produce some view-inconsistent artifacts in the pseudo ground truths. However, during the distillation, the student network can automatically eliminate these artifacts due to the multi-view training that enforces view-consistent robustness. See Sec. 4.2 and Fig. 6 for details.

4. Experiments and Analysis

4.1. Implementation Details

Network. In order to disentangle shape and color latent information within the hashgrids, we split the single hash table in the NeRF network architecture of Instant-NGP [26] into two: a density grid \mathcal{G}^σ and a color grid \mathcal{G}^c , with the same settings as the original density grid in the open-source PyTorch implementation torch-ngp [39]. We do this to make it possible to make fine-grained edits of one to one of the color or geometry properties without affecting the other. The rest of the network architecture remains the same, including a sigma MLP f^σ and a color MLP f^c . For a spatial point \mathbf{x} with view direction \mathbf{d} , the network predicts volume density σ and color c as follows:

$$\sigma, \mathbf{z} = f^\sigma(\mathcal{G}^\sigma(\mathbf{x})) \quad (9)$$

$$c = f^c(\mathcal{G}^c(\mathbf{x}), \mathbf{z}, \text{SH}(\mathbf{d})) \quad (10)$$

where \mathbf{z} is the intermediate geometry feature, and SH is the spherical harmonics directional encoder [26]. The same as

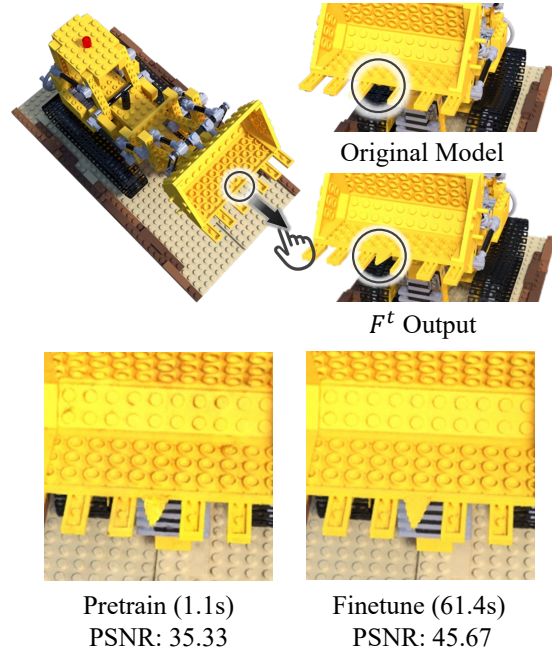


Figure 5: Example of anchor editing: fake tooth.




Figure 6: Example of editing on the real-world scene: thumbs up to thumbs down (DTU Scan 83).



Figure 7: Example of object transfer editing: from the Lego scene (NeRF Blender) to the family scene (Tanks and Temples).

Instant-NGP’s settings, f^σ has 2 layers with hidden channel 64, f^c has 3 layers with hidden channel 64, and \mathbf{z} is a 15-channel feature.

We compare our modified NeRF network with the vanilla

architecture in the Lego scene of NeRF Blender Synthetic dataset[24]. We train our network and the vanilla network on the scene for 30,000 iterations. The result is as follows:

- Ours: training time 441s, PSNR 35.08dB
- Vanilla: training time 408s, PSNR 34.44dB

We observe slightly slower runtime and higher quality for our modified architecture, indicating that this modification causes negligible changes.

Training. We select Instant-NGP [26] as the NeRF backbone of our editing framework. Our implementations are based on the open-source PyTorch implementation torch-ngp [39]. All experiments are run on a single NVIDIA RTX 3090 GPU. Note that we make a slight modification to the original network architecture. Please refer to the supplementary material for details.

During the pretraining stage, we set $\lambda_1 = \lambda_2 = 1$ and the learning rate is fixed to 0.05. During the finetuning stage, we set $\lambda_3 = \lambda_4 = 1$ with an initial learning rate of 0.01. Starting from a pretrained NeRF model, we perform 50-100 epochs of local pretraining (for about 0.5-1 seconds) and about 50 epochs of global finetuning (for about 40-60 seconds). The number of epochs and time consumption can be adjusted according to the editing type and the complexity of the scene. Note that we test our performance in the absence of tiny-cuda-nn [25] which achieves superior speed to our backbone, which indicates that our performance has room for further optimization.

Datasets. We evaluate our editing in the synthetic NeRF Blender Dataset [24], and the real-world captured Tanks and Temples [13] and DTU [11] datasets. We follow the official dataset split of the frames for the training and evaluation.

4.2. Experimental Results

Qualitative NeRF editing results. We provide extensive experimental results in all kinds of editing categories we design, including bounding shape (Figs. 4 and 6), brushing (Fig. 3), anchor (Fig. 5), and color (Fig. 1). Our method not only achieves a huge performance boost, supporting instant preview at the second level but also produces more visually realistic editing appearances, such as shading effects on the lifted side in Fig. 3 and shadows on the bumped surface in Fig. 8. Besides, results produced by the student network can even outperform the teacher labels, e.g. in Fig. 6 the F^t output contains floating artifacts due to view inconsistency. As analyzed in Sec. 3.3, the distillation process manages to eliminate this. We also provide an example of object transfer (Fig. 7): the bulb in the Lego scene (of Blender dataset) is transferred to the child’s head in the family scene of Tanks and Temples dataset.




Figure 8: Comparisons on texture/color painting between NeuMesh [5] and our method. Note that NeuMesh requires hours of finetuning while ours needs only seconds.




Figure 9: Comparison on qualitative and quantitative between NeuMesh [5] and our method. The PSNR is computed from the editing result and the rendering of the ground truth mesh with the same editing applied.




Figure 10: Comparison of the pixel-wise editing ability between baselines [20, 5] and ours. Note that [20] does not focus on the same task of pixel-wise editing as the other two. We are not to compete with their method.

Comparisons to baselines. Existing works have strong restrictions on editing types, which focus on either geometry editing or appearance editing, while ours is capable of doing both simultaneously. Our brushing and anchor tools can create user-guided out-of-proxy geometry structures, which no existing methods support. We make comparisons on color and texture painting supported by NeuMesh [5] and




Figure 11: Ablation studies on two-stage training strategy. Zoom in for degradation details of “w/o finetuning”.




Figure 12: Ablation study on MLP fixing.

Liu *et al.* [20].

Fig. 8 illustrates two comparisons between our method and NeuMesh [5] in scribbling and a texture painting task. Our method significantly outperforms NeuMesh, which contains noticeable color bias and artifacts in the results. In contrast, our method even succeeds in rendering the shadow effects caused by geometric bumps.

Fig. 9 illustrates the results of the same non-rigid blending applied to the Mic from NeRF Blender[24]. It clearly shows that being mesh-free, We have more details than NeuMesh[5], unlimited by mesh resolution.

Fig. 10 shows an overview of the pixel-wise editing ability of existing NeRF editing methods and ours. Liu *et al.* [20]’s method does not focus on the pixel-wise editing task and only supports textureless simple objects in their paper. Their method causes an overall color deterioration

within the edited object, which is highly unfavorable. This is because their latent code only models the global color feature of the scene instead of fine-grained local features. Our method supports fine-grained local edits due to our local-aware embedding grids.

4.3. Ablation Studies

Effect of the two-stage training strategy. To validate the effectiveness of our pretraining and finetuning strategy, we make comparisons between our full strategy (3rd row), finetuning-only (1st row) and pretraining-only (2nd row) in Fig. 11. Our pretraining can produce a coarse result in only 1 second, while photometric finetuning can hardly change the appearance in such a short period. The pretraining stage also enhances the subsequent finetuning, in 30 seconds our full strategy produces a more complete result. However, pretraining has a side effect of local overfitting and global degradation. Therefore, our two-stage strategy makes a good balance between both and produces optimal results.

MLP fixing in the pretraining stage. In Fig. 12, we validate our design of fixing all MLP parameters in the pretraining stage. The result confirms our analysis that MLP mainly contains global information so it leads to global degeneration when MLP decoders are not fixed.

5. Conclusion

We have introduced an interactive framework for pixel-level editing for neural radiance fields supporting instant preview. Specifically, we exploit the two-stage teacher-student training method to provide editing guidance and design a two-stage training strategy to achieve instant network convergence to obtain coarse results as a preview. Unlike previous works, our method does not require any explicit proxy (such as mesh), improving interactivity and user-friendliness. Our method also supports preserving shading effects on the edited surface. One limitation is that our method does not support complex view-dependent lighting effects such as specular reflections, and can not change the scene illumination, which can be improved by introducing intrinsic decomposition. Besides, our method does not handle the reconstruction failures (such as floating artifacts) of the original NeRF network.

ACKNOWLEDGEMENT

This work was supported in part by the Fundamental Research Funds for the Central Universities; NSFC under Grants (62103372, 62088101, 62233013); the Key Research and Development Program of Zhejiang Province (2021C03037); Zhejiang Lab (121005-PI2101); Information Technology Center and State Key Lab of CAD&CG, Zhejiang University.

References

- [1] Adobe Inc. Adobe photoshop. **2**
- [2] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. 2020. **3**
- [3] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. *ICCV*, 2021. **1**
- [4] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, 2022. **1, 3**
- [5] Chong Bao and Bangbang Yang, Zeng Junyi, Bao Hujun, Zhang Yinda, Cui Zhaopeng, and Zhang Guofeng. Neumesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In *European Conference on Computer Vision (ECCV)*, 2022. **2, 3, 4, 8, 9**
- [6] Peng Dai, Yinda Zhang, Zhuwen Li, Shuaicheng Liu, and Bing Zeng. Neural point cloud rendering via multi-plane projection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7830–7839, 2020. **3**
- [7] Jianfei Guo, Zhiyuan Yang, Xi Lin, and Qingfu Zhang. Template nerf: Towards modeling dense shape correspondences from category-specific object images. *arXiv preprint arXiv:2111.04237*, 2021. **2**
- [8] Michelle Guo, Alireza Fathi, Jiajun Wu, and Thomas Funkhouser. Object-centric neural scene rendering. *arXiv preprint arXiv:2012.08503*, 2020. **3**
- [9] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, Light, and Material Decomposition from Images using Monte Carlo Rendering and Denoising. *arXiv:2206.03380*, 2022. **3**
- [10] Jinkai Hu, Chengzhong Yu, Hongli Liu, Lingqi Yan, Yiqian Wu, and Xiaogang Jin. Deep real-time volumetric rendering using multi-feature fusion. In *ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH '23*, New York, NY, USA, 2023. Association for Computing Machinery. **3**
- [11] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 406–413. IEEE, 2014. **8**
- [12] Natasha Kholgade, Tomas Simon, Alexei Efros, and Yaser Sheikh. 3d object manipulation in a single photograph using stock 3d models. *ACM Transactions on Computer Graphics*, 33(4), 2014. **3**
- [13] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017. **8**
- [14] Zhengfei Kuang, Fujun Luan, Sai Bi, Zhixin Shu, Gordon Wetzstein, and Kalyan Sunkavalli. Palettenerf: Palette-based appearance editing of neural radiance fields. *arXiv preprint arXiv:2212.10699*, 2022. **2, 3**
- [15] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2475–2484, 2020. **3**
- [16] Zhengqin Li, Jia Shi, Sai Bi, Rui Zhu, Kalyan Sunkavalli, Miloš Hašan, Zexiang Xu, Ravi Ramamoorthi, and Manmohan Chandraker. Physically-based editing of indoor scene lighting from a single image. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VI*, pages 555–572. Springer, 2022. **3**
- [17] Zhengqi Li, Wenqi Xian, Abe Davis, and Noah Snavely. Crowdsampling the plenoptic function. In *European Conference on Computer Vision*, pages 178–196. Springer, 2020. **3**
- [18] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. **3**
- [19] Hao-Kang Liu, I Shen, Bing-Yu Chen, et al. Nerf-in: Free-form nerf inpainting with rgb-d priors. *arXiv preprint arXiv:2206.04901*, 2022. **2**
- [20] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. **2, 3, 4, 8, 9**
- [21] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4):65:1–65:14, July 2019. **3**
- [22] Aryan Mikaeili, Or Perel, Mehdi Safaei, Daniel Cohen-Or, and Ali Mahdavi-Amiri. Sked: Sketch-guided text-based 3d editing, 2023. **2**
- [23] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019. **3**
- [24] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. **1, 3, 8, 9**
- [25] Thomas Müller. Tiny CUDA neural network framework, 2021. <https://github.com/nvmlabs/tiny-cuda-nn>. **8**
- [26] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. **1, 3, 4, 7, 8**
- [27] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8280–8290, 2022. **2, 3**
- [28] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. **3**

- [29] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *ACM Trans. Graph.*, 40(6), dec 2021. 3
- [30] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 3
- [31] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. In *ACM SIGGRAPH 2003 Papers*, pages 313–318. 2003. 3
- [32] Yunlong Ran, Jing Zeng, Shibo He, Jiming Chen, Lincheng Li, Yingfeng Chen, Gimhee Lee, and Qi Ye. Neurar: Neural uncertainty for autonomous 3d reconstruction with implicit neural representations. *IEEE Robotics and Automation Letters*, 8(2):1125–1132, 2023. 3
- [33] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 1, 3
- [34] Rakshith Shetty, Mario Fritz, and Bernt Schiele. Adversarial scene editing: Automatic object removal from weak supervision. In *Advances in Neural Information Processing Systems 31*, pages 7716–7726, Montréal, Canada, 2018. Curran Associates. 3
- [35] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2019. 3
- [36] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, 2019. 3
- [37] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022. 3
- [38] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul Srinivasan, Jonathan T. Barron, and Henrik Kretzschmar. Block-NeRF: Scalable large scene neural view synthesis. *arXiv*, 2022. 3
- [39] Jiayang Tang. Torch-ngp: a pytorch implementation of instant-ngp, 2022. <https://github.com/ashawkey/torch-ngp>. 7, 8
- [40] Jiayang Tang, Xiaokang Chen, Jingbo Wang, and Gang Zeng. Compressible-composable nerf via rank-residual decomposition. *arXiv preprint arXiv:2205.14870*, 2022. 3
- [41] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 1, 3
- [42] Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. HumanNeRF: Free-viewpoint rendering of moving people from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16210–16220, June 2022. 3
- [43] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In *The European Conference on Computer Vision (ECCV)*, 2022. 3
- [44] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *International Conference on Computer Vision (ICCV)*, October 2021. 2, 3
- [45] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. 1, 3
- [46] Weicai Ye, Shuo Chen, Chong Bao, Hujun Bao, Marc Pollefeys, Zhaopeng Cui, and Guofeng Zhang. Intrinsicnerf: Learning intrinsic neural radiance fields for editable novel view synthesis. 2022. 3
- [47] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 3
- [48] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: Geometry editing of neural radiance fields. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3, 4
- [49] Jing Zeng, Yanxu Li, Yunlong Ran, Shuo Li, Fei Gao, Lincheng Li, Shibo He, Jiming Chen, and Qi Ye. Efficient view path planning for autonomous implicit reconstruction. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4063–4069, 2023. 3
- [50] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020. 1
- [51] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)*, 40(6):1–18, 2021. 3
- [52] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. In *SIGGRAPH*, 2018. 3
- [53] Jingsen Zhu, Yuchi Huo, Qi Ye, Fujun Luan, Jifan Li, Dianbing Xi, Lisha Wang, Rui Tang, Wei Hua, Hujun Bao, et al. I2-sdf: Intrinsic indoor scene reconstruction and editing via raytracing in neural sdf. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12489–12498, 2023. 3
- [54] Jingsen Zhu, Fujun Luan, Yuchi Huo, Zihao Lin, Zhihua Zhong, Dianbing Xi, Rui Wang, Hujun Bao, Jiayang Zheng, and Rui Tang. Learning-based inverse rendering of complex indoor scenes with differentiable monte carlo raytracing. In *SIGGRAPH Asia 2022 Conference Papers*. ACM, 2022. 3