

基於顏色和深度先驗機率的神經網路輻射場自由修復

劉豪剛*
國立台灣大學
台灣

沈奕超*
東京大學
日本

陳炳宇
國立台灣大學
台灣

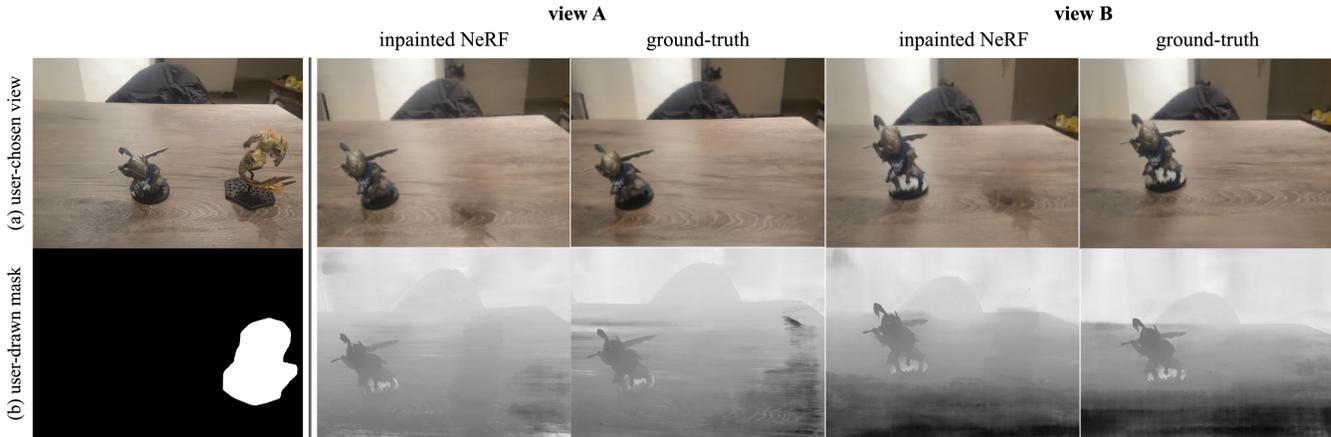


Figure 1: Given a pre-trained NeRF model, the user can (a) choose a view and (b) draw a mask to specify the unwanted object in the 3D scene. Our framework optimizes the NeRF model based on a user-provided mask and removes the unwanted object in the mask region. The optimized NeRF synthesizes inpainted results that match ground truth results from different views. (Noted that the user-chosen view is different from view A and view B.)

ABSTRACT

近年來，神經輻射場(NeRF)在電腦視覺和電腦圖學上展現了強大的能力，其利用神經網路的學習能夠生成對應場景的各視角的3D影像。而後續的相關研究也展現了神經輻射場能夠處理許多不同的任務，包括位置預測、動態場景生成、光線處理等。然而如何將2D的資訊傳播在神經輻射場這樣的3D表示法中依舊是個問題。由於我們無法得知哪幾個權重會對神經輻射場的顏色和幾何形狀產生改變，這無疑是個艱難的挑戰。在這篇論文中，我們將會介紹第一個能夠讓使用者隨意去除場景物體並基於神經輻射場的架構。在一開始的情境中，使用者可以針對一個角度下的照片去進行編輯，移除其中不想要的物體或區塊。隨後架構會將這個使用者提供的資訊(遮罩)轉移到其他角度渲染的照片並產生初步的顏色/深度參考圖。緊接著我們會根據參考圖去優化原本的神經輻射場，修改內部有關聯的權重並達到移除物體的效果。我們將我們的架構實驗在不同的場景中並發現其能夠展現視覺完整且各角度一致的渲染效果，同時也不需要人力的介入或者重新訓練整個神經輻射場。

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

*Both authors contributed equally to the paper.

Authors' addresses: 劉豪剛, 國立台灣大學, 台灣; 沈奕超, 東京大學, 日本; 陳炳宇, 國立台灣大學, 台灣.

2021. 0730-0301/2021/7-ART39 \$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

KEYWORDS

Wireless sensor networks, media access control, multi-channel, radio interference, time synchronization

ACM Reference Format:

劉豪剛, 沈奕超, and 陳炳宇. 2021. 基於顏色和深度先驗機率的神經網路輻射場自由修復. *ACM Trans. Graph.* 38, 4, Article 39 (July 2021), 9 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Recent advancements in neural rendering, such as Neural Radiance Fields (NeRF) [Mildenhall et al. 2020] has emerged as a powerful representation for the task of novel view synthesis, where the goal is to render unseen viewpoints of a scene from a given set of input images. NeRF encodes the volumetric density and color of a scene within the weights of a coordinate-based multi-layer perceptron. Several follow-up works extend original NeRF to handle different tasks, such as pose estimation [Lin et al. 2021; Wang et al. 2021b], 3D-aware image synthesis [Chan et al. 2021; Niemeyer and Geiger 2021; Schwarz et al. 2020], deformable 3D reconstruction [Liu et al. 2021a; Park et al. 2021; Pumarola et al. 2020], and modeling dynamic scenes [Gao et al. 2021; Guo et al. 2021; Xian et al. 2021].

Though NeRF achieves great performance on photo-realistic scene reconstruction and novel view synthesis, there remain enormous challenges in editing the geometries and appearances of a scene represented by a pre-trained NeRF model. Unlike traditional image editing, a user needs to transfer his/her edits on a rendered view to the NeRF model to edit the whole scene, thus introducing multiple challenges. First, it is unclear where the edited regions

appear on other rendered views. Second, because millions of parameters are used in a pre-trained NeRF model, it is unclear which parameters control the different aspects of the rendered shape and how to change the parameters according to the sparse local user input. Previous works [Liu et al. 2021b] enable users to perform color and shape editing on a category-level NeRF. However, these methods require additional category-specific training and data to support the desired editings.

In this paper, we focus on the NeRF inpainting problem, i.e., removing unwanted objects in a 3D scene represented by a pre-trained NeRF. Although we can ask a user to provide a mask and the inpainted image for each rendered view and use these images to train a new NeRF, there are several disadvantages. First, it is labor-intensive to provide masks for many rendered views. Second, there will be visual inconsistency across different inpainted views introduced by separate inpainting.

To address these issues, we propose a framework to help users easily remove unwanted objects by updating a pre-trained NeRF model. Given a pre-trained NeRF, the user first draws a mask over a rendered view. Given the user-drawn mask, our framework first rendered a couple of views sampled from a pre-set trajectory. Next, we transfer the user-drawn masks to these sampled views using existing video object segmentation method [Cheng et al. 2021]. Our framework then generates (i) guiding color image regions using [Cao and Fu 2021] and (ii) guiding depth images using Bilateral Solver [Barron and Poole 2016] within these masked regions. Noted that our framework can use any existing methods for generating guiding color and depth images. Finally, we formulate an optimization problem that jointly inpaints the image content within the all transferred masked regions with respect to the guiding color and depth images.

We demonstrate our framework on several scenes represented by pre-trained NeRFs in LLFF dataset. We show that our framework generates visually plausible and consistent results. Furthermore, we also demonstrate our experiments on the custom dataset to show the correctness between inpainted results and ground truth results.

2 RELATED WORK

2.1 Novel view synthesis

Constructing novel views of a scene captured by multiple images is a long standing problem in computer graphics and computer vision. Traditional methods use structure-from-motion [Hartley and Zisserman 2003] and bundle adjustment [Triggs et al. 1999] to reconstruct explicit point cloud structure and recover camera parameters. Other methods synthesize novel views by interpolating within a 4D light fields [Gortler et al. 1996; Levoy and Hanrahan 1996] and by compositing the warped layers in the multiplane image representations (MPI) [Tucker and Snavely 2020; Zhou et al. 2018]. Recently, the coordinate-based neural representations have shown significant promise as an alternative to discrete, grid representations for scene representations. Neural Radiance Fields (NeRF) [Mildenhall et al. 2020] use a multi-layer perceptron (MLP) and positional encoding to model a radiance field at an unprecedented level of fidelity. However, NeRF works on a large number of input images and requires lengthy training time. Many works attempt to reduce the number of images NeRF requires by introducing depth-supervised loss [Deng

et al. 2021; Roessle et al. 2021] and category-specific priors [Yu et al. 2021b]. Meanwhile, previous works try to reduce the training time by optimizing voxel grids of features [Sun et al. 2021; Yu et al. 2021a] and factoring the radiance field [Chen et al. 2022].

These recent advances greatly improve the practical use of NeRF. However, it is still unintuitive how a user can edit a pre-trained NeRF model. The main reason is because the neural network of a NeRF model has millions of parameters. Which parameters control the different aspects of the rendered shape and how to change the parameters to achieve desired editis are still unknown. Previous works enable users to select certain object [Ren et al. 2022], edit a NeRF model using strokes [Liu et al. 2021b], natural language [Wang et al. 2021a], and by manipulating 3D model [Yuan et al. 2022] directly. However, these methods require to learn additional category-level conditional radiance fields or segmentation network to facilitate such edits. Unlike these methods, our framework did not requires the user to prepare any additional category-specific training data and training procedure for removing unwanted objects in a pre-trained NeRF model.

2.2 Image inpainting

In recent years, two broad approaches to image inpainting exist. Patch-based method [Barnes et al. 2009; Efros and Leung 1999; Simakov et al. 2008] fill the holes by searching for patches with similar low-level image features such as rgb values. The search space can be the non-hole region of the input image or from other reference images. The inpainted results are obtained by a global optimization after the relevant patches are retrieved. These methods often fail to handle large holes where the color and texture variance is high. Meanwhile, these methods often cannot make semantically aware patch selections. Deep learning-based methods often predict the pixel values inside masks directly in a semantic-aware fashion. Thus they can synthesize more visually plausible contents especially for images like faces [Li et al. 2017; Yeh et al. 2017], objects [Pathak et al. 2016] and natural scenes [Iizuka et al. 2017]. However, these methods often focus on regular masks only. To handle irregular masks, partial convolution [Liu et al. 2018] is proposed where the convolution is masked and re-normalized to utilize valid pixels only. Yu *et al.* [Yu et al. 2018] uses GAN mechanism to maintain local and global consistency in the final results. Nazeri *et al.* [Nazeri et al. 2019] focus on improving the image structure in the inpainting results by conditioning their image inpainting network on edges in the masked regions. MST inpainting [Cao and Fu 2021] and ZITS [Dong et al. 2022] further consider both edge and line structure to synthesize more reasonable results. In this work, we use MST inpainting network [Cao and Fu 2021] to obtain the guided inpainted result because of its superior performance on inpainting images while preserving structures. Our framework can replace MST inpainting with other inpainting methods since we only used the inpainted results as a guiding signal for our optimization problem.

3 METHOD

In this section, we first summarize the mechanism of NeRF [Mildenhall et al. 2020] and formulate our problem setting.

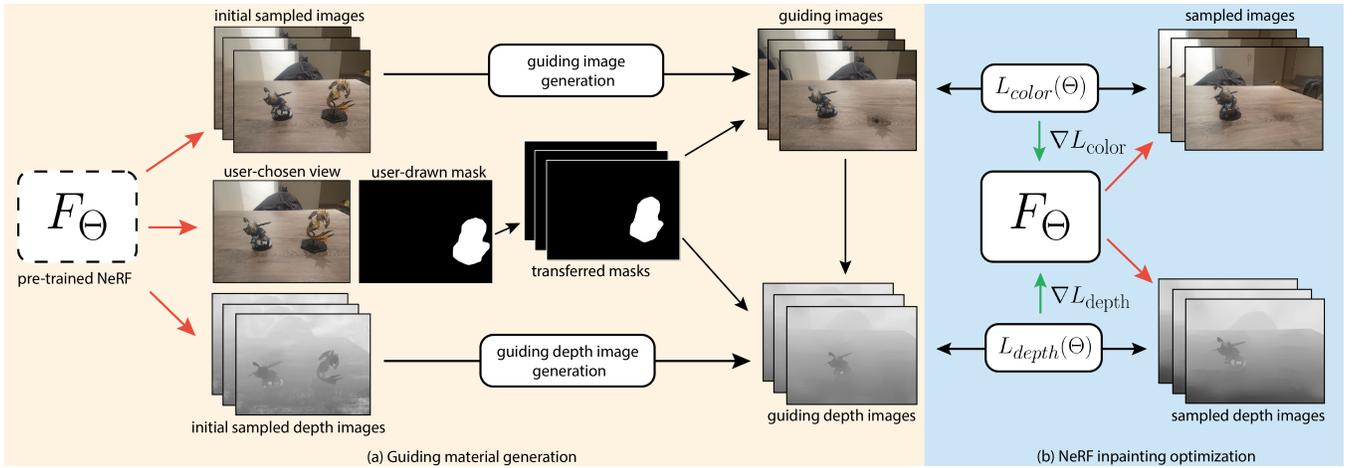


Figure 2: (a) Given a pre-trained NeRF F_Θ , an user specifies the unwanted region on an user-chosen view with a user-drawn mask. Our framework sampled initial images and initial depth images and generate both guiding images and guiding depth images. (b) Our framework update Θ by optimizing both color-guiding loss (L_{color}) and depth-guiding loss (L_{depth}). (\rightarrow denotes render a view from a NeRF model and \leftarrow denotes updating Θ by optimizing losses.)

3.1 Preliminaries: NeRF

NeRF is a continuous volumetric radiance field $F_\Theta : (\mathbf{x}, \mathbf{d}) \rightarrow (\mathbf{c}, \sigma)$ represented by a MLP network with Θ as its weights. $F_\Theta : (\mathbf{x}, \mathbf{d}) \rightarrow (\mathbf{c}, \sigma)$ takes a 3D position $\mathbf{x} = \{x, y, z\}$ and 2D viewing direction $\mathbf{d} = \{\theta, \phi\}$ as input and outputs volume density σ and directional emitted color \mathbf{c} . NeRF renders the volume of each camera rays passing through the scene by computing the volume rendering intergral using numerical quadrature. The expected color $C(\mathbf{r})$ of camera ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ is defined as:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T(t_i)(1 - \exp(-\sigma(t_i)\delta_i))c(t_i), \quad (1)$$

$$\text{where } T(t_i) = \exp\left(-\sum_{j=1}^{i-1} \sigma(t_j)\delta_j\right), \quad (2)$$

where N denotes the total quadrature points sampled between near plane t_n and far plane t_f of the camera, and $\delta_i = t_{i+1} - t_i$ is the distance between two adjacent points. We denote the color and desity at point t_i produced by NeRF model F_Θ as $c(t_i)$ and $\sigma(t_i)$.

Using the above differentiable rendering equation, we can propagte the errors and update Θ through mean square error:

$$\mathcal{L}_{mse} = \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{C}^c(\mathbf{r}) - C(\mathbf{r})\|_2^2 + \|\hat{C}^f(\mathbf{r}) - C_i(\mathbf{r})\|_2^2, \quad (3)$$

where \mathcal{R} is a ray batch, $C(\mathbf{r})$, $\hat{C}^c(\mathbf{r})$, $\hat{C}^f(\mathbf{r})$ are the ground truth, coarse volume predicted, and fine volume predicted RGB colors for ray \mathbf{r} respectively. For simplicity, we further define $F_\Theta^{\text{image}} : \mathbf{o} \rightarrow I$ as a function that takes a camera position (\mathbf{o}) as input, and outputs the rendered image of a pre-trained NeRF model F_Θ .

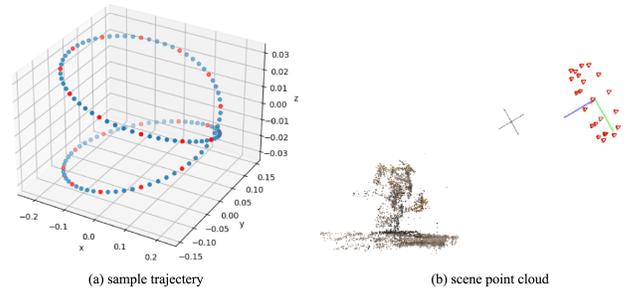


Figure 3: Our sampling strategy follows the trajectory in (a). Noted that the target scene faces toward the “+y” directon as shown in (b). Each blue dot represents a view we can sample, while the red dot represents the sample view used in the optimization framework. Also, we use sample images to construct the point cloud for better understanding. We conduct all the experiments based on this setting.

3.2 Overview

Given a pre-trained NeRF model: F_Θ , a user can specify the unwanted region by drawing a mask M_u over a user-chosen rendered view $I_u = F_\Theta^{\text{image}}(\mathbf{o}_u)$, where $M_u = 1$ for pixel outside the masked region, and \mathbf{o}_u is the user-chosen camera position. Our goal is to obtain an updated NeRF model $F_{\hat{\Theta}}$ such that the unwanted region masked by M_u is removed in every rendered views. As shown in Figure 2, our method first sample K camera positions $\mathbf{O} = \{\mathbf{o}_s | s = 1 \dots K\}$ along the test-set camera trajectory used in LLFF [Mildenhall et al. 2019] (Figure 3). For each camera position, we rendered a rgb image I_s and depth image D_s using F_Θ and obtained all rendered views $\mathbf{I} = \{I_s | s = 1 \dots K\}$ and their depth images $\mathbf{D} = \{D_s | s = 1 \dots K\}$. We will use I_s and $I(\mathbf{o}_s)$ to represent the image rendered from camera position \mathbf{o}_s interchangeably throughout the

paper. One can potentially remove the unwanted object specified by M_u using the following naive method. First, remove the content within the transferred masked region on each sampled rendered view. Then, update Θ using only the image content outside all transferred masked region by optimizing the “masked mse (mmse)” function:

$$L_{\text{mmse}} = \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{C}^c(\mathbf{r}) - C(\mathbf{r}) \odot M(\mathbf{r})\|_2^2 + \|\hat{C}^f(\mathbf{r}) - C_i(\mathbf{r}) \odot M(\mathbf{r})\|_2^2, \quad (4)$$

where M is the transferred mask on the same view as the sample ray \mathbf{r} . However, because there is no explicit guidance on what image content and structure should be in the masked region, the unwanted object will remain in the result of optimizing Equation (5). To provide explicit guidance, our method takes the user-drawn mask M_u , sampled rendered views \mathbf{I} , and sampled depth images \mathbf{D} as input, and outputs

- guiding user-chosen image and guiding depth image: I_u^G and D_u^G .
- transferred masks: $\mathbf{M} = \{M_s | s = 1 \dots K\}$.
- guiding sampled images and their guiding depth images: $\mathbf{I}^G = \{I_s^G | s = 1 \dots K\}$ and $\mathbf{D}^G = \{D_s^G | s = 1 \dots K\}$.

Finally, our method obtains updated parameters $\tilde{\Theta}$ by optimizing our nerf inpainting formulation: $\Phi(M_u, I_u^G, \mathbf{M}, \mathbf{D}_s^G)$.

3.3 Guiding material generation

For each sampled rendered view I_s , our goal is to generate a mask M_s that covers the same object as the user-drawn mask M_u . We use a video object segmentation method (STCN) [Cheng et al. 2021] to generate M_s . With the transferred masks M_s , we need to generate the guiding images and guiding depth images. The guiding image generation can be describe as

$$I_s^G = \rho(I_s, M_s), \quad (5)$$

where I_s^G is the guiding image, and ρ is a single image inpainting method (we used MST inpainting network [Cao and Fu 2021]). After obtaining I_s^G , we can obtain the guiding depth image using

$$D_s^G = \tau(D_s, M_s, I_s^G), \quad (6)$$

where D_s^G is the guiding depth image, and τ is a depth image completion method (we used Fast Bilateral solver [Barron and Poole 2016]). Noted that our framework can replace ρ to any other single image inpainting method and τ to any other single depth image completion method.

3.4 NeRF inpainting optimization

We obtain the updated parameters $\tilde{\Theta}$ that removes the unwanted object in the 3D scene by optimizing:

$$\tilde{\Theta} := \arg \min_{\Theta} L_{\text{color}}(\Theta) + L_{\text{depth}}(\Theta) \quad (7)$$

where L_{color} is the color-guiding loss and L_{depth} is the depth-guiding loss.

3.4.1 Color-guiding loss. The color-guiding loss used is defined as

$$L_{\text{color}}(\Theta) = L_{\text{color}}^{\text{all}}(\Theta) + L_{\text{color}}^{\text{out}}(\Theta), \quad (8)$$

where $L_{\text{color}}^{\text{all}}$ is defined on views \mathbf{O}^{all} , $L_{\text{color}}^{\text{out}}$ is defined on views \mathbf{O}^{out} , and $\mathbf{O}^{\text{all}} \cup \mathbf{O}^{\text{out}} = \{\mathbf{O}, \mathbf{o}_u\}$. $L_{\text{color}}^{\text{all}}$ is used to measure the color difference of the entire image (inside and outside of the mask) on the rendered view and is defined as:

$$L_{\text{color}}^{\text{all}}(\Theta) = \sum_{\mathbf{o} \in \mathbf{O}^{\text{all}}} F_{\Theta}^{\text{image}}(\mathbf{o}) - I_{\mathbf{o}}^G. \quad (9)$$

$L_{\text{color}}^{\text{out}}$ is used to measure the color difference outside the mask on the rendered view and is defined as:

$$L_{\text{color}}^{\text{out}}(\Theta) = \sum_{\mathbf{o} \in \mathbf{O}^{\text{out}}} (F_{\Theta}^{\text{image}}(\mathbf{o}) - I_{\mathbf{o}}^G) \odot M_{\mathbf{o}}. \quad (10)$$

In our framework, we set $\mathbf{O}^{\text{all}} = \mathbf{o}_u$ and $\mathbf{O}^{\text{out}} = \mathbf{O}$.

3.4.2 Depth-guiding loss. While we can obtain visual plausible inpainted color results using the color-guided loss, it often generates incorrect depth, which might cause incorrect geometry and keep some unwanted objects in the scene. To fix these incorrect geometries, we introduce a depth-guided loss, which is defined as:

$$L_{\text{depth}}(\Theta) = \sum_{\mathbf{o}_s \in \mathbf{O}} \|D^f(\mathbf{o}_s) - D^G(\mathbf{o}_s)\|_2^2 + \|D^c(\mathbf{o}_s) - D^G(\mathbf{o}_s)\|_2^2, \quad (11)$$

where $D^f(\mathbf{o}_s)$ is the fine volume predicted depth image, $D^c(\mathbf{o}_s)$ is the coarse volume predicted depth image, and we rendered both depth image from a sampled camera position \mathbf{o}_s using F_{Θ} . We compute the depth $D^f(\mathbf{o}_s)$ through computing the accumulation of σ from ray batches.

4 EXPERIMENTS AND EVALUATIONS

In this section, we show qualitative results on LLFF [Mildenhall et al. 2019] dataset and our custom dataset, followed by ablation studies.

4.1 Implementation detail

We implement our framework in PyTorch [Paszke et al. 2019] and Python 3.9. We test our framework on a machine with Intel i7-7800X and a GTX-1080 graphics card to train our models. For each scene, we first train a model initialized to random weights and optimize it for 200,000 steps with a batch size of 4,096 using Adam [Kingma and Ba 2014], which takes about 18 to 20 hours. The sample points used in fine and coarse model are 128 and 64, respectively. To inpaint each scene, we optimize Eq. 7 for 50,000 steps which takes about five hours in total.

4.2 Evaluation

4.2.1 Datasets. To verify our framework’s performance, we create a custom dataset that contains three custom scenes: **figyua**, **desuku**, and **terebe**. The purpose is to obtain the ground truth results of NeRF inpainting. For each custom scene, we collect a pair of photo set, i.e., (**original** and **removed**). For **original** set, we keep all the objects in the scene and take photos from 24 camera positions. For **removed** set, we remove one object in the scene manually and take

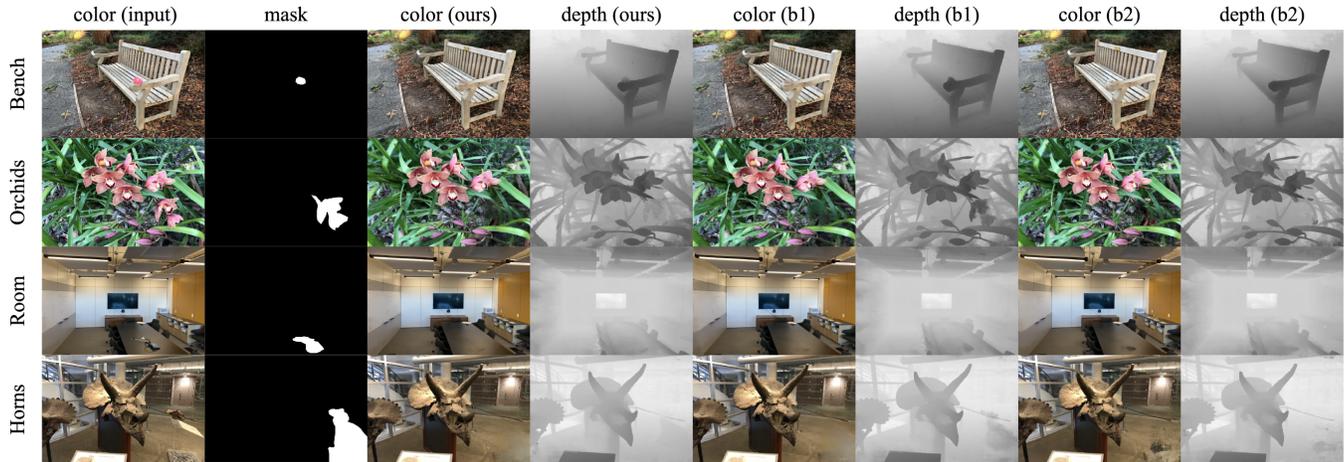


Figure 4: Qualitative comparison - LLFF dataset. For each scene, we show the user-chosen view image and the user-provided mask on the left. We then show the color image and depth image generated by different methods: our method (ours), *baseline1* (b1), and *baseline2* (b2). The depth map of b1 still keep depth of the unwanted object. Meanwhile, the color of b2 might cause noise or shadow on the scene (shown in horns). Our method, compared to these two baselines, have better color and correct geometry on final results.

photos from the same 24 camera positions. The **original** sets is used as the input of our framework, and the **removed** set is used as the ground truth of the results after the inpainting optimization.

4.2.2 Experiment setup. As we are the first to free-form inpainting on NeRF, we propose two baseline methods for comparisons:

baseline1: per-view color updating We update the pre-trained NeRF model F_{Θ} with all guiding images I^G by optimizing

$$\tilde{\Theta} := \arg \min_{\Theta} \sum_{I_s^G \in I^G} (F_{\Theta}^{\text{image}}(\mathbf{o}_s) - I_s^G) \quad (12)$$

baseline2: masked mse retraining We re-train a new NeRF model using all guiding images I^G by optimizing:

$$\tilde{\Theta} := \arg \min_{\Theta} \sum_{I_s^G \in I^G} (F_{\Theta}^{\text{image}}(\mathbf{o}_s) - I_s^G) \odot M_s, \quad (13)$$

where Θ is randomly initialized.

Both *baseline1* and *baseline2* did not consider depth information during updating the pre-trained NeRF model or re-train a new NeRF model.

We compared the inpainted results of our framework to inpainted results of two baseline methods on LLFF dataset and our custom dataset described in Section 4.2.1. For LLFF dataset, we perform qualitative evaluation by applying three methods on each pre-trained model. For our custom dataset, we perform both qualitative and quantitative evaluations. For each scene, we trained two separate NeRF models for the **original** set and the **removed** set. We apply three methods to the pre-trained NeRF model using the **original** set. We then compared the inpainted result with the image generated by the pre-trained NeRF model using the **removed** set.

4.2.3 Results and discussions.

LLFF dataset We show the qualitative comparison between our method and two baseline methods using LLFF dataset in Figure 4

and Figure 6. In Figure 4, we observed that the depth maps of the inpainted NeRF generated by *baseline1* did not match the inpainted image content. In Figure 6, we showed that there are obvious visual inconsistencies between views in the results generated by *baseline1*. To avoid these visual inconsistency, we choose to provide color guidance using only the user-chosen view and let the NeRF model maintain the view consistency by itself. *Baseline2* recovers visual satisfactory image content without any color guidance inside the masked region. However, *baseline2* still generate results that losses fine structures or synthesize some unnatural patches at complicated regions, which can be observed at Horns and Orchids.

Custom dataset We show the qualitative comparison between our method and two baseline methods using our custom dataset in Figure 1 and Figure 5. In Figure 5, we can observe that although *baseline1* can synthesize rgb content closer to ground truth, it still fail to generate correct depth map. On the other hand, *baseline2* recovers the content in the masked region guided by the content from different views but still creates noisy and blurry content. Our framework generates closer color and depth images to the ground truth rendered results compared to the two baseline methods.

Overall, our inpainting optimization updates a pre-trained NeRF model to obtain correct geometry and preserve visual consistency across views.

4.3 Ablation study

4.3.1 How important is the depth-guiding loss? Introducing the depth-guiding loss (L_{depth}) is one of the major contribution of our framework. We validate its effectiveness by comparing with the optimization results using color-guiding loss (L_{color}) only, depth-guiding loss (L_{depth}) only, and both losses.

We showed the results in Figure 7. We observe that optimizing using L_{depth} only already leads to correct geometries inside the masked region but introduces color noises outside the masked

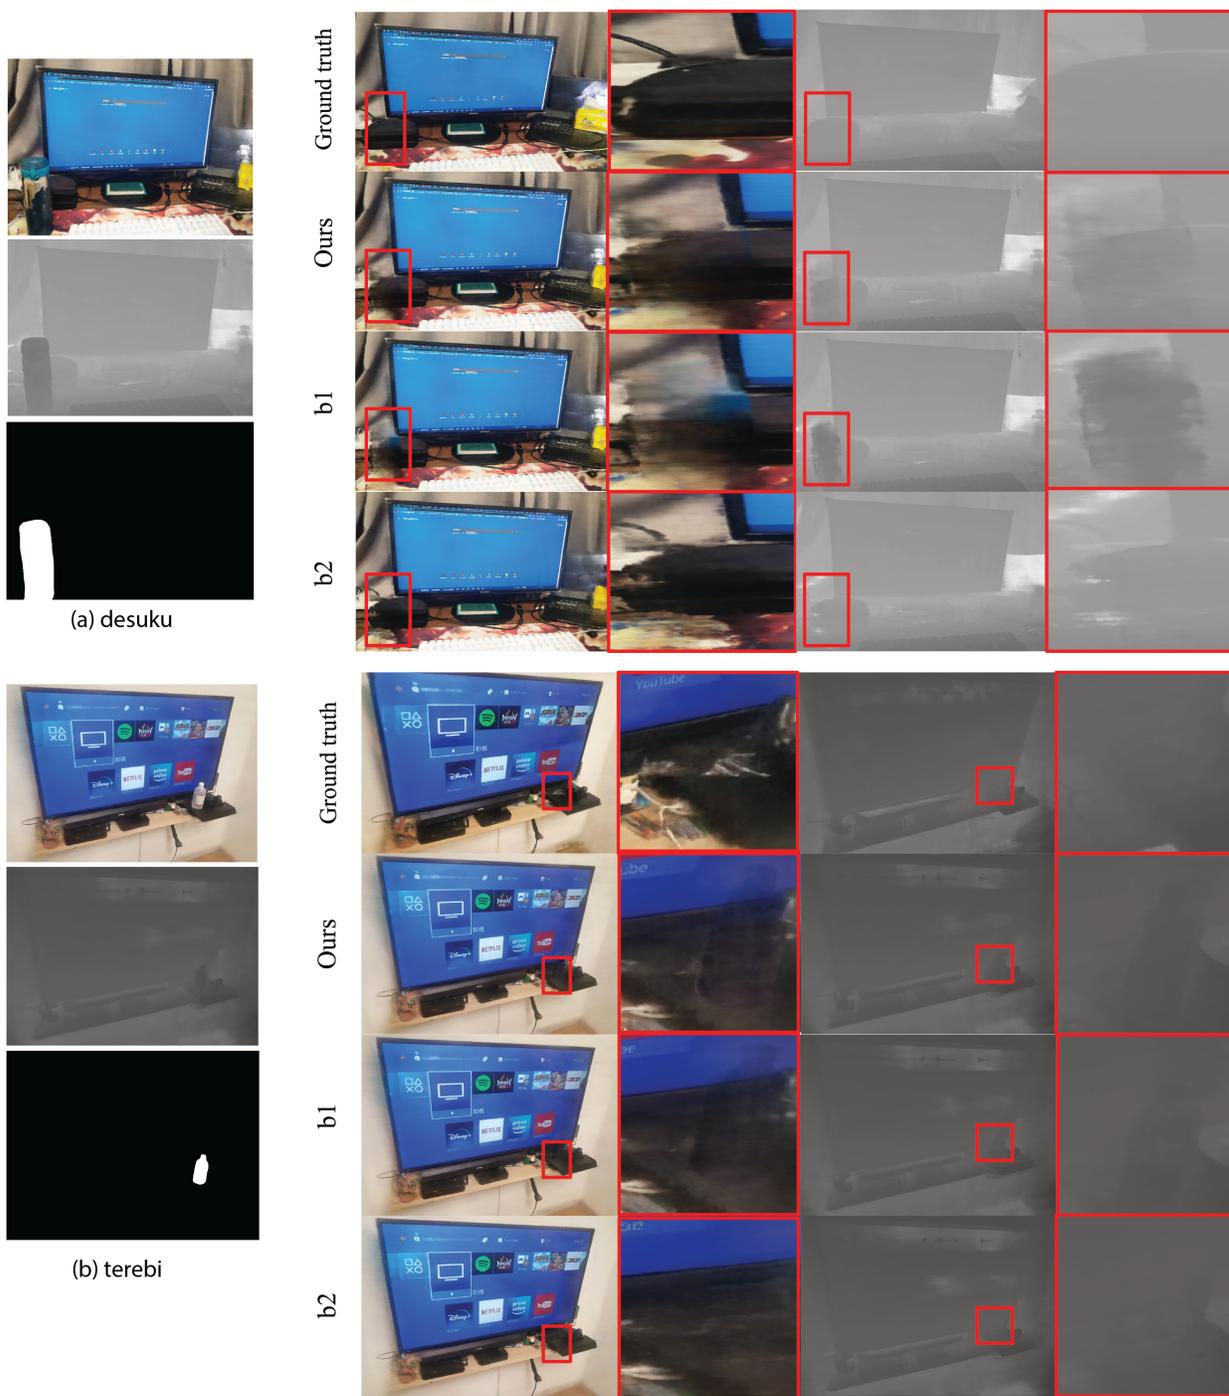


Figure 5: Qualitative comparison - custom dataset. For each custom scene, we demonstrate the ground truth rendered image, results generated by our framework, *baseline1* (b1), and *baseline2* (b2). Our framework generates more accurate depth maps and synthesizes more fine structures compared to *baseline1*. Compared to *baseline2*, our framework synthesizes more realistic and shape results.

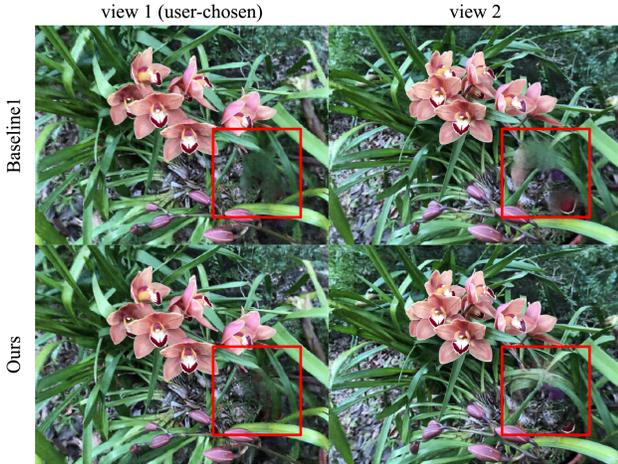


Figure 6: Qualitative comparison - visual consistency. The rendered views generated by *baseline1* have severe visual inconsistency across different views (within the red box region). Meanwhile, our method synthesize visual consistent results across different views.

region. Our method optimizes both losses and generate correct geometries without color noises. In Figure 8, we can also observed that the unwanted object in the region with high depth variations can be removed by using L_{depth} only (red box). However, using L_{depth} only loses color information in the flat region (blue box). Our method combines these two losses and remove the unwanted object without losing color information in the flat region.

4.3.2 How important is color-guiding within the masked regions from sampled views? In our framework, we only use $L_{\text{color}}^{\text{all}}$ to guide the color reconstruction inside the masked regions. We validate this function design by adjusting the number of views used to compute $L_{\text{color}}^{\text{all}}$ during the optimization.

We compared the results of following three settings:

- (1) only *user-chosen* view is used to guide the color inside the masked region, i.e., $\mathbf{O}^{\text{all}} = \mathbf{o}_u$ and $\mathbf{O}^{\text{out}} = \mathbf{O}$.
- (2) *three sampled views* are used to guide the color inside the masked region, i.e., $\mathbf{O}^{\text{all}} = \{\mathbf{o}_i, \mathbf{o}_j, \mathbf{o}_k\}$ where i, j, k are randomly sampled, and $\mathbf{O}^{\text{out}} = \{\mathbf{o}_s | \mathbf{O} - \mathbf{O}^{\text{all}}\}$.
- (3) *all sampled views* (i.e., 24) are used to guide the color inside the masked region, i.e., $\mathbf{O}^{\text{all}} = \mathbf{O}$ and $\mathbf{O}^{\text{out}} = \mathbf{o}_u$.

As shown in Figure 9, we observe that more visual inconsistency will be introduced when we use more inpainted images as color guidance. Our framework obtains stable results for most of the scene using *user-chosen* view only; thus, we choose to not to include other inpainted images in the computation of $L_{\text{color}}^{\text{all}}$.

5 LIMITATIONS AND FUTURE WORK

Fuse color and depth guidance. Our framework leverages existing image inpainting and depth completion method to generate initial guidance materials. Once there are appearances or geometries artifacts in the initial guidance materials, the optimized NeRF

might output undesired or incorrect results. Meanwhile, our framework shares the same limitations as the image inpainting method we used. For example, our framework fails to inpaint the image region with high reflectance content (Figure 10) or with a thin structure. It is possible to design a fusion method to fuse color and depth guidances from multiple methods.

Update masks and guidance materials. In our current framework, we used fixed masks and guidance materials during the optimization. However, this is sub-optimal when the unwanted object is occluded in some views. In the future, we plan to extend our framework to update the masks in every optimization step using 3D volume features extracted from the NeRF model. We also plan to use a discriminator to constrain the synthesized results and improve the view consistency.

Volume feature for mask transferring Our current framework uses the existing video-based object segmentation method to transfer the user-drawn mask. It is possible to perform mask transferring by conducting 3D volume segmentation using the volume feature extracted from the pre-trained NeRF.

6 CONCLUSION

In this paper, we propose the first framework that enables users to remove unwanted objects or retouch undesired regions in a 3D scene represented by a pre-trained NeRF. Our framework requires no additional category-specific data and training. Instead, we formulated a novel optimization to inpaint the pre-trained NeRF with the generated RGB-D guidances. We demonstrated our framework handles a variety of scenes well, and we also validate our framework using a custom dataset where ground truth inpainted results are available. We believe that the custom dataset we proposed and our framework can foster future research on neural radiance field editings.

REFERENCES

- Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. 2009. Patch-Match: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.* 28, 3 (2009), 24.
- Jonathan T. Barron and Ben Poole. 2016. The Fast Bilateral Solver. *ECCV* (2016).
- Chenjie Cao and Yanwei Fu. 2021. Learning a sketch tensor space for image inpainting of man-made scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14509–14518.
- Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. 2021. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5799–5809.
- Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. 2022. TensorRF: Tensorial Radiance Fields. *arXiv preprint arXiv:2203.09517* (2022).
- Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. 2021. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *Advances in Neural Information Processing Systems* 34 (2021).
- Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. 2021. Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791* (2021).
- Qiaole Dong, Chenjie Cao, and Yanwei Fu. 2022. Incremental Transformer Structure Enhanced Image Inpainting with Masking Positional Encoding. *arXiv preprint arXiv:2203.00867* (2022).
- Alexei A Efros and Thomas K Leung. 1999. Texture synthesis by non-parametric sampling. In *Proceedings of the seventh IEEE international conference on computer vision*, Vol. 2. IEEE, 1033–1038.
- Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. 2021. Dynamic View Synthesis from Dynamic Monocular Video. *arXiv preprint arXiv:2105.06468* (2021).
- Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. 1996. The lumigraph. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. 43–54.

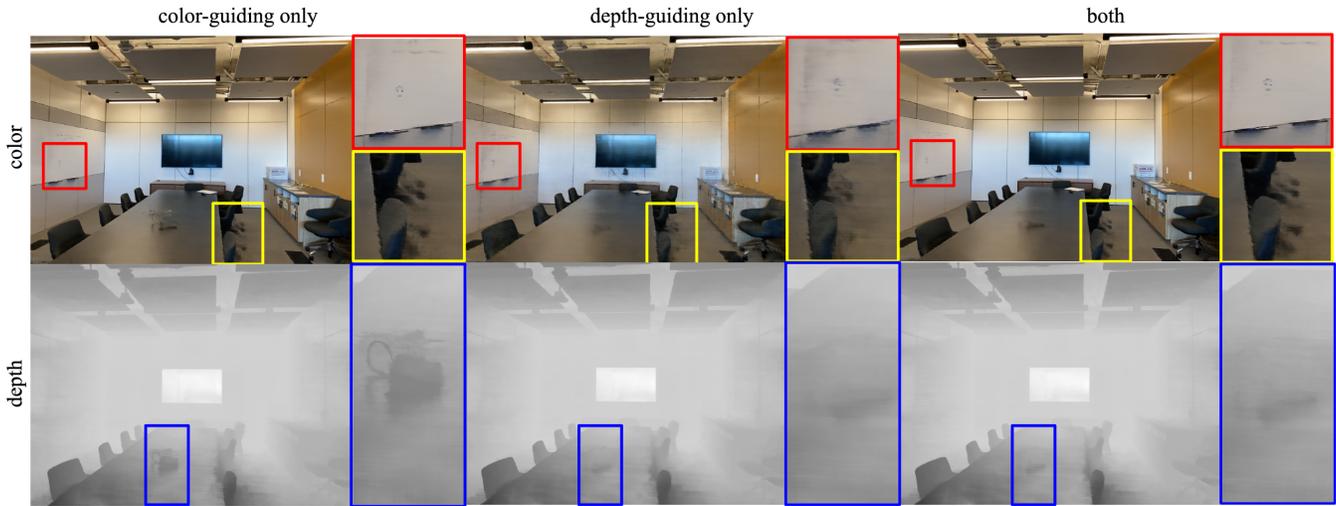


Figure 7: Depth-guiding ablation result. We show the optimization results using different guiding losses. We observed that the geometry of the unwanted object could not be removed using color-guiding term only (depth inside the blue box). On the other hand, using depth-guiding term only helps to get correct geometry but introduces color noises outside the masked region (red box and yellow box). Our method combines both terms to generate correct geometry (blue box) without introducing any color noises (red box and yellow box).



Figure 8: Depth-guiding losses discussion. The unwanted object in the region with high depth variations can be removed by using depth-guiding loss (red box). However, using depth-guiding loss only loses color information in the flat region (blue box). Combining both losses removes the unwanted object without losing any color information.

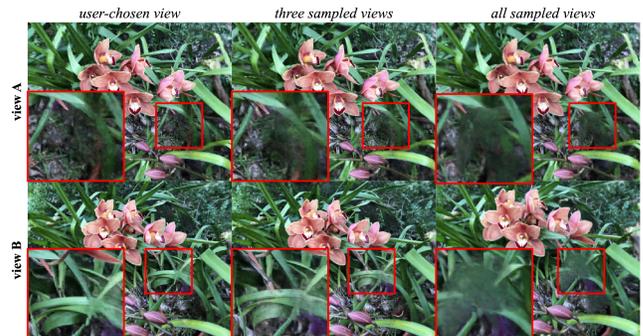


Figure 9: Color-guiding ablation result. We show the optimization results using different number of views to compute L_{color}^{all} . The visual inconsistency becomes larger in the masked region (red box) when we increase the number of views used to compute L_{color}^{all} .

Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. 2021. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5784–5794.

Richard Hartley and Andrew Zisserman. 2003. *Multiple view geometry in computer vision*. Cambridge university press.

Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. 2017. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)* 36, 4 (2017), 1–14.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

Marc Levoy and Pat Hanrahan. 1996. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. 31–42.

Yijun Li, Sifei Liu, Jimei Yang, and Ming-Hsuan Yang. 2017. Generative face completion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3911–3919.

Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. 2021. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5741–5751.

Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. 2018. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European conference on computer vision (ECCV)*. 85–100.

Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. 2021a. Neural Actor: Neural Free-view Synthesis of Human Actors with Pose Control. *ACM SIGGRAPH Asia* (2021).

Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. 2021b. Editing conditional radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5773–5783.

Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. 2019. Local Light Field Fusion: Practical View Synthesis with Prescriptive Sampling Guidelines. *ACM Transactions on Graphics (TOG)* (2019).

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.

Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Z Qureshi, and Mehran Ebrahimi. 2019. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212* (2019).

Michael Niemeyer and Andreas Geiger. 2021. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11453–11464.

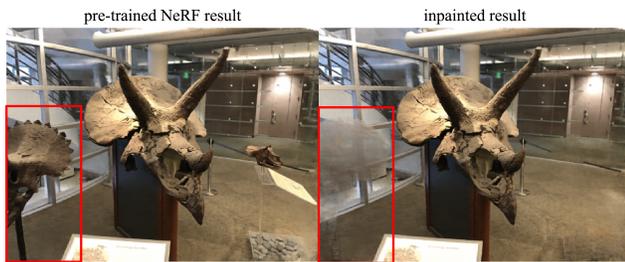


Figure 10: Our framework fails to inpaint the mask region in the left region (red box) and introduces artifacts in the optimized results.

- Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. 2021. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5865–5874.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2536–2544.
- Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. 2020. D-NeRF: Neural Radiance Fields for Dynamic Scenes. <https://arxiv.org/abs/2011.13961> (2020).
- Zhongzheng Ren, Aseem Agarwala[†], Bryan Russell[†], Alexander G. Schwing[†], and Oliver Wang[†]. 2022. Neural Volumetric Object Selection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. ([†] alphabetic ordering).
- Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. 2021. Dense Depth Priors for Neural Radiance Fields from Sparse Input Views. *arXiv preprint arXiv:2112.03288* (2021).
- Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. 2020. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems 33* (2020), 20154–20166.
- Denis Simakov, Yaron Caspi, Eli Shechtman, and Michal Irani. 2008. Summarizing visual data using bidirectional similarity. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1–8.
- Cheng Sun, Min Sun, and Hwann-Tzong Chen. 2021. Direct Voxel Grid Optimization: Super-fast Convergence for Radiance Fields Reconstruction. *arXiv preprint arXiv:2111.11215* (2021).
- Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. 1999. Bundle adjustment—a modern synthesis. In *International workshop on vision algorithms*. Springer, 298–372.
- Richard Tucker and Noah Snavely. 2020. Single-view view synthesis with multiplane images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 551–560.
- Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. 2021a. CLIP-NeRF: Text-and-Image Driven Manipulation of Neural Radiance Fields. *arXiv preprint arXiv:2112.05139* (2021).
- Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. 2021b. NeRF-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064* (2021).
- Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. 2021. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9421–9431.
- Raymond A Yeh, Chen Chen, Teck Yian Lim, Alexander G Schwing, Mark Hasegawa-Johnson, and Minh N Do. 2017. Semantic image inpainting with deep generative models. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5485–5493.
- Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. 2021a. Plenoxels: Radiance Fields without Neural Networks. *arXiv preprint arXiv:2112.05131* (2021).
- Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. 2021b. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference*

on *Computer Vision and Pattern Recognition*. 4578–4587.

Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. 2018. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5505–5514.

Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. 2022. NeRF-editing: geometry editing of neural radiance fields. *arXiv preprint arXiv:2205.04978* (2022).

Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. 2018. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817* (2018).