

Removing Objects From Neural Radiance Fields

Silvan Weder^{1,2} Guillermo Garcia-Hernando¹ Áron Monszpart¹ Marc Pollefeys²
Gabriel Brostow^{1,3} Michael Firman¹ Sara Vicente¹
¹Niantic ²ETH Zurich ³University College London

Abstract

Neural Radiance Fields (NeRFs) are emerging as a ubiquitous scene representation that allows for novel view synthesis. Increasingly, NeRFs will be shareable with other people. Before sharing a NeRF, though, it might be desirable to remove personal information or unsightly objects. Such removal is not easily achieved with the current NeRF editing frameworks. We propose a framework to remove objects from a NeRF representation created from an RGB-D sequence. Our NeRF inpainting method leverages recent work in 2D image inpainting and is guided by a user-provided mask. Our algorithm is underpinned by a confidence based view selection procedure. It chooses which of the individual 2D inpainted images to use in the creation of the NeRF, so that the resulting inpainted NeRF is 3D consistent. We show that our method for NeRF editing is effective for synthesizing plausible inpaintings in a multi-view coherent manner. We validate our approach using a new and still-challenging dataset for the task of NeRF inpainting.

1. Introduction

Since the initial publication of Neural Radiance Fields (NeRFs) [41], there has been an explosion of extensions to the original framework, *e.g.*, [3, 4, 8, 12, 25, 34, 38, 41]. NeRFs are being used beyond the initial task of novel view synthesis. It is already appealing to get them into the hands of non-expert users for novel applications, *e.g.*, for NeRF editing [78] or live capture and training [45]. More casual use-cases are driving interesting new technical issues.

One of those issues is how to seamlessly remove parts of the rendered scene. Removing parts of the scene can be desirable for a variety of reasons. For example, a house scan being shared on a property selling website may need unappealing or personally identifiable objects to be removed [66]. Similarly, objects could be removed so they can be replaced in an augmented reality application, *e.g.*, removing a chair from a scan to see how a new chair fits in the environment [49]. Removing objects might also be desirable when a NeRF is part of a traditional computer vi-




Figure 1. **Removal of unsightly objects.** Our method allows for objects to be plausibly removed from NeRF reconstructions, inpainting missing regions whilst preserving multi-view coherence.

sion pipeline, *e.g.*, removing parked cars from scans that are going to be used for relocalization [42].

Some editing of NeRFs has already been explored. For example, object-centric representations disentangle labeled objects from the background, which allows editing of the trained scene with user-guided transformations [72, 75], while semantic decomposition allows selective editing and transparency for certain semantic parts of the scene [26]. However, these previous approaches only augment information from the input scan, limiting their generative capabilities, *i.e.*, the hallucination of elements that haven't been observed from any view.

With this work, we tackle the problem of removing objects from scenes, while realistically filling the resulting holes, as shown in Fig. 1. Solving this problem requires: a)

exploiting multi-view information when parts of the scene are observed in some frames but occluded in others and, b) leveraging a generative process to fill areas that are never observed. To this end, we pair the multi-view consistency of NeRFs with the generative power of 2D inpainting models [67] that are trained on large scale 2D image datasets. Such 2D inpaintings are not multi-view consistent by construction, and may contain severe artefacts. This causes corrupted reconstructions, so we design a new confidence-based view-selection scheme that iteratively removes inconsistent inpaintings from the optimization. We validate our approach on a new dataset and show that we outperform existing approaches on standard metrics of image quality, as well as producing multi-view consistent results.

In summary, we make the following contributions:

1) We propose the first approach focusing on inpainting NeRFs by leveraging the power of single image inpainting. 2) A novel view-selection mechanism that automatically removes inconsistent views from the optimization. 3) We present a new dataset for evaluating object removal and inpainting in indoor and outdoor scenes.

2. Related work

Image inpainting. Image inpainting tackles the problem of plausibly filling in missing regions in a single image. A typical approach is to use an image-to-image network with an adversarial loss, *e.g.*, [22, 51, 77, 80], or with a diffusion model [39]. Different ways have been proposed to encode the input image, *e.g.*, using masked [32] or Fourier convolutions [67]. Image inpainting was extended to also inpaint depth images by [14]. However, these methods do not give temporal consistency between video frames, nor the ability to synthesize novel views.

Removing moving objects from videos. While video inpainting is a well studied problem in computer vision [73, 79], most works focus on removing moving objects. This is typically achieved with guidance from nearby frames, *e.g.*, via estimating flow [15, 21, 73], sometimes using depth [6, 31]. Perhaps counter-intuitively, moving objects make the task easier, since their movement disoccludes the background, making most parts of the scene visible in at least one of the frames. In contrast, to inpaint videos of a static scene a generative approach is required, since large parts of the scene have never been observed.

Removing static objects from videos. Where occluded pixels are visible in other frames in the sequence, these can be used to inpaint regions [28, 43, 44]. For example, [36, 37] remove static foreground distractors from videos, *e.g.*, fences and raindrops. However, there are typically still pixels which cannot be seen in other views, for which some other method is required to fill them in. For example, [18] propagates patches from visible pixels into the region to in-

paint, and [43, 68] inpaint missing pixels via PatchMatch. Kim et al. [24] rely on a pre-computed mesh of each scene for object removal. Our key difference to these methods is that our inpaintings can be extrapolated to novel viewpoints.

Inpainting in novel view synthesis. Inpainting is often used as a *component* of novel view synthesis, to estimate textures for regions unobserved in the inputs [64, 70], *e.g.*, for panorama generation [20, 27]. Philip et al. [52] enable object removal from image-based-rendering, but with an assumption that background regions are locally planar.

2.1. Novel view synthesis and NeRFs

NeRF [41] is a highly popular image-based rendering method which uses a differentiable volume-rendering formulation to represent a scene; a multi-layer perceptron (MLP) is trained to regress the color and opacity given a 3D coordinate and ray viewing direction. This combined works on implicit 3D scene representations [9, 33, 40, 50, 56–58], with light-field rendering [11] and novel view synthesis [17, 35, 74]. Extensions include work that reduces aliasing artifacts [3], can cope with unbounded scenes [4], reconstructs a scene from only sparse views [8, 25, 38, 47] or makes NeRFs more efficient, *e.g.*, by using octrees [69, 76] or other sparse structures [59].

Depth-aware Neural Radiance fields. To overcome some of the limitations of NeRFs, particularly the requirement for dense views and the limits in the quality of the reconstructed geometry, depths can be used in training [12, 55]. These can be sparse depths from structure-from-motion [61, 62], or depth from sensors [53, 65].

Object-centric and semantic NeRFs for editing. One direction of progress in NeRFs is the decomposition of the scene into its constituent objects [48, 72, 75]. This is done based on motion for dynamic scenes [48], or instance segmentation for static scenes [75]. Both lines of work also model the background of the scene as a separate model. However, similar to video inpainting, dynamic scenes allow a better modelling of the background since more of it is visible. In contrast, visual artefacts can be seen in the background representation of [72, 75], which model static scenes. Methods that decompose the scene based on semantics [26, 82] can also be used to remove objects. However, they do not try to complete the scene when a semantic part is removed and, for example, [26] discusses how “the background behind the deleted objects can be noisy or have a hole because it lacks observation”.

Generative models for novel view synthesis. 3D aware generative models can be used to synthesize views of an object or scene from different viewpoints, in a 3D consistent manner [13, 46, 54, 63]. In contrast with NeRF models, which only have a test time component and “overfit” to a specific scene, generative models can be used to hallucinate views of novel objects by sampling in the latent variable




Figure 2. **Per-frame inpainting** can give plausible results for each frame, but they are not consistent between viewpoints and sometimes contain severe artefacts corrupting the optimization.

space. There has also been some interest in 3D generative models that work for full indoor scenes [13, 29, 54]. However, their capacity to fit the source views (or memorization) can be limited, as shown in the qualitative results of [13]. To train the generative model, [13, 29, 54] require a large dataset of indoor scenes with RGB and camera poses and in some cases depth [13, 29]. In contrast, our use of a 2D pretrained inpainting network, that can be trained on any image, is less dependent on the existence of training data and less constrained to indoor scenarios.

3. Method

We assume we are given an RGB-D sequence with camera poses and intrinsics. Depths and poses could be acquired, for example, using a dense structure-from-motion pipeline [61, 62]. For most of our experiments we capture posed RGB-D sequences directly using Apple’s ARKit framework [2], but we also show that we can relax this requirement through use of a multi-view stereo method to estimate depth for an RGB sequence. Along the way, we also assume access to a per-frame mask of the object to be removed. The goal is to learn a NeRF model from this input, which can be used to synthesize consistent novel views, where the per-frame masked region should be plausibly inpainted. An overview of our method is shown in Fig. 3.

3.1. RGB and depth inpainting network

Our method relies on a 2D single image inpainting method to inpaint each RGB image individually. Furthermore, we also require a depth inpainting network. We use both networks as black-boxes and our approach is agnostic to the method chosen. Future improvements in single image inpainting can be directly translated to improvements to our method. Given an image I_n and corresponding mask M_n , the per-image inpainting algorithm produces a new image

\tilde{I}_n . Similarly, the depth inpainting algorithm produces an inpainted depth map \tilde{D}_n . We show some results for the 2D inpainting network in Figure 2.

3.2. Background on NeRFs

Following the original NeRF paper [41], we represent the scene as an MLP F_Θ that predicts color $\mathbf{c} = [r, g, b]$ and density σ , for a 5 dimensional input containing x, y, z position and two viewing directions. The predicted color for pixel \mathbf{r} , $\hat{I}_n(\mathbf{r})$, is obtained by volume rendering along its associated ray, so

$$\hat{I}_n(\mathbf{r}) = \sum_{i=1}^K \underbrace{T_i(1 - \exp(-\sigma_i \delta_i))}_{w_i} \mathbf{c}_i, \quad T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right), \quad (1)$$

where K is the number of samples along the ray, t_i is a sample location, $\delta_i = t_{i+1} - t_i$ is the distance between adjacent samples, and w_i is the alpha compositing weight which, by construction, sum to less than or equal to 1.

The NeRF loss operates on training images as

$$\mathcal{L}_{\text{RGB}} = \sum_{n=1}^N \sum_{\mathbf{r} \in \Omega_n} \left\| I_n(\mathbf{r}) - \hat{I}_n(\mathbf{r}) \right\|^2, \quad (2)$$

where $I_n(\mathbf{r})$ is the input RGB value for pixel \mathbf{r} , and $\hat{I}_n(\mathbf{r})$ is its predicted color. Ω_n indicates the 2D domain of image n . The parameters of the MLP, Θ , are optimized to minimize this loss. Similarly to [55], if input depth is available, then an additional loss can be added,

$$\mathcal{L}_{\text{depth}} = \sum_{n=1}^N \sum_{\mathbf{r} \in \Omega_n} \left| D_n(\mathbf{r}) - \hat{D}_n(\mathbf{r}) \right|, \quad \text{with } \hat{D}_n(\mathbf{r}) = \sum_{i=1}^K w_i t_i, \quad (3)$$

where $D_n(\mathbf{r})$ is the input depth for pixel \mathbf{r} , and $\hat{D}_n(\mathbf{r})$ is the corresponding predicted depth.

Finally, a distortion regularizer loss was introduced in [4] to better constrain the NeRF optimization and remove “floaters”. It encourages the non-zero compositing weights w_i to be concentrated in a small region along the ray, so for each pixel \mathbf{r} ,

$$l_{\text{dist}}(\mathbf{r}) = \sum_{i,j} w_i w_j \left| \frac{t_i + t_{i+1}}{2} - \frac{t_j + t_{j+1}}{2} \right| + \frac{1}{3} \sum_i w_i^2 \delta_i. \quad (4)$$

3.3. Confidence-based view selection

Despite most of the individual inpainted RGB images \tilde{I}_n looking realistic, they still suffer from two issues: 1) some of the inpaintings are incorrect, and 2) despite individual plausibility, they are not multi-view consistent, *i.e.*, the same area observed in multiple views is not necessarily completed in a consistent way (Figure 2). For this reason,

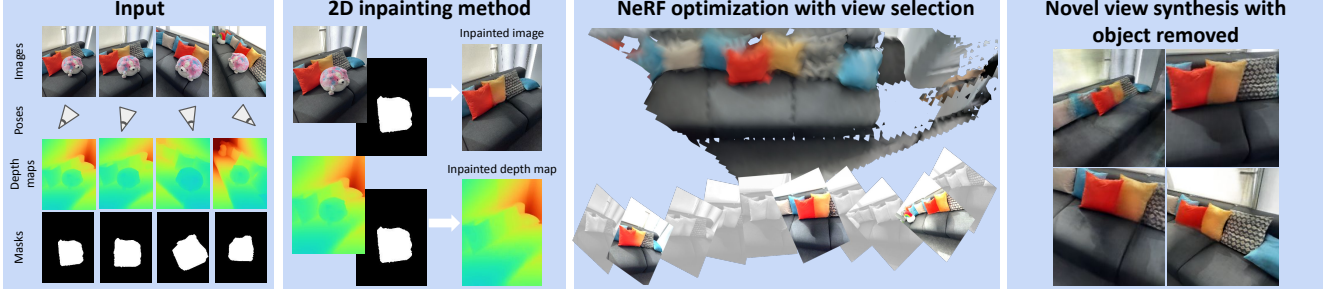


Figure 3. **An overview of our method** We take a sequence of posed RGB-D images together with corresponding 2D masks as input. The 2D frames are inpainted using [67] and then used to optimize a neural radiance field. During optimization, our confidence-based view selection automatically removes inconsistent views from the optimization preventing unwanted artefacts in the final result. Finally, novel views can be rendered from the scene, where the object has been removed.

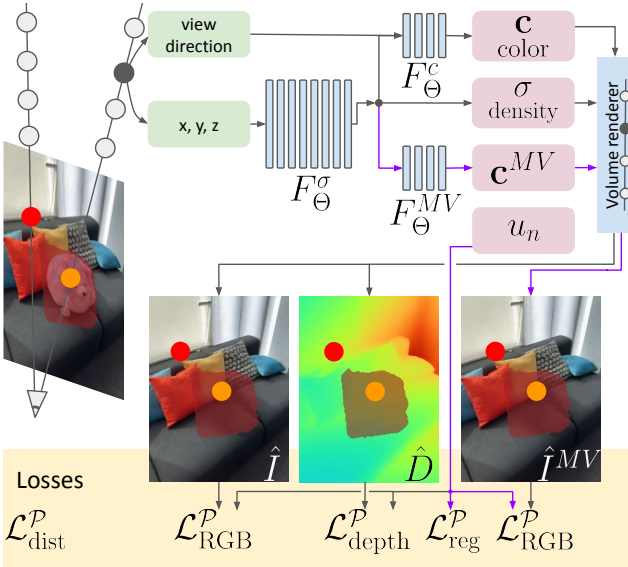


Figure 4. **Our architecture.** Our NeRF formulation contains two color heads F_{Θ}^C and F_{Θ}^{MV} , where F_{Θ}^{MV} does not take the view direction as input. This is important to encourage multi-view consistency (see text for details). It also includes uncertainty variables u_n that are used to jointly model the scene and the uncertainty of the 2D inpaintings for our automated view selection.

we propose a confidence-based view selection scheme, that automatically chooses which views are used in the NeRF optimization. We associate to each image n a non-negative uncertainty measure u_n . The corresponding per-image confidence, e^{-u_n} , is used to re-weight the NeRF losses. This confidence value can be seen as a loss attenuation term, similar to the aleatoric uncertainty prediction term in [23].

The RGB loss for our model is then set out as

$$\mathcal{L}_{\text{RGB}}^{\mathcal{P}}(\hat{I}) = \sum_{n=1}^N \sum_{\mathbf{r} \in \Omega_n \setminus M_n} \left\| I_n(\mathbf{r}) - \hat{I}_n(\mathbf{r}) \right\|^2 + \sum_{n \in \mathcal{P}} e^{-u_n} \sum_{\mathbf{r} \in M_n} \left\| \hat{I}_n(\mathbf{r}) - \hat{I}_n(\mathbf{r}) \right\|^2 \quad (5)$$

where the color for pixels \mathbf{r} is supervised by the inpainted image for pixels inside the mask, and by the input RGB image for pixels outside the mask. Note that the second

term of this loss is only computed over a restricted set of images \mathcal{P} , where $\mathcal{P} \subseteq \{1, \dots, N\}$. This is denoted by the superscript \mathcal{P} in $\mathcal{L}_{\text{RGB}}^{\mathcal{P}}$. In practice, that means that only some inpainted regions are used in the NeRF optimization. We discuss below how we choose the set \mathcal{P} .

We use a similar split into pixels inside and outside the mask for the depth loss, so

$$\mathcal{L}_{\text{depth}}^{\mathcal{P}} = \sum_{n=1}^N \sum_{\mathbf{r} \in \Omega_n \setminus M_n} \left| D_n(\mathbf{r}) - \hat{D}_n(\mathbf{r}) \right| + \sum_{n \in \mathcal{P}} e^{-u_n} \sum_{\mathbf{r} \in M_n} \left| \hat{D}_n(\mathbf{r}) - \hat{D}_n(\mathbf{r}) \right|. \quad (6)$$

Finally, we include two regularizers. One is on the uncertainty weights $\mathcal{L}_{\text{reg}}^{\mathcal{P}} = \sum_{n \in \mathcal{P}} u_n$, to prevent a trivial solution where e^{-u_n} is 0. The other is a distortion regularizer, based on [3], around the loss detailed in Equation (4), so

$$\mathcal{L}_{\text{dist}}^{\mathcal{P}} = \sum_{n=1}^N \sum_{\mathbf{r} \in \Omega_n \setminus M_n} l_{\text{dist}}(\mathbf{r}) + \sum_{n \in \mathcal{P}} \sum_{\mathbf{r} \in M_n} l_{\text{dist}}(\mathbf{r}). \quad (7)$$

View direction and multi-view consistency. When optimizing the NeRF, we made three observations: a) the multi-view inconsistencies in the inpaintings are modelled by the network using the viewing direction; b) we can enforce multi-view consistency by removing the viewing direction from the input; c) the inconsistencies introduce cloud-like artefacts in the density when not using the viewing direction as input. To prevent a) and c) and correctly optimize the variables u_n that capture the uncertainty about the inpaintings \hat{I}_n , we propose: 1) adding an auxiliary network head, F_{Θ}^{MV} , to the NeRF that does not take the viewing direction as input and, 2) stopping the gradient from the color inpainting and F_{Θ}^{MV} to the density, leaving the uncertainty variable u_n as the only view-dependent input. This design forces the model to encode the inconsistencies between inpaintings into the uncertainty prediction while keeping the model consistent across views. F_{Θ}^{MV} has a loss term based on Eq. 5: $\mathcal{L}_{\text{RGB}}^{\mathcal{P}}(\hat{I}^{MV})$. See Figure 4 for an illustration of our architecture.

Algorithm 1: Iterative refinement using confidence based view selection.

Data: Input images I_n , Inpainted images \tilde{I}_n , Depth maps D_n , Inpainted depth maps \tilde{D}_n , Masks M_n

Result: Trained NeRF model F_Θ with object removed

/* Set of images used for training NeRF is initialized with all images. */

$\mathcal{P} \leftarrow \{1, \dots, N\}$ and $u_n \leftarrow 0, n \in \mathcal{P}$

for $i \leftarrow 0$ **to** K_{outer} **do**

$\Theta \leftarrow$ randomly initialized

for $j \leftarrow 0$ **to** K_{grad} **do**

 /* Gradient iterations of NeRF training. */

$\Theta \leftarrow \Theta - \nabla_{\Theta} \mathcal{L}^{\mathcal{P}}$

$\mathbf{U}^{\mathcal{P}} \leftarrow \mathbf{U}^{\mathcal{P}} - \nabla_{\mathbf{U}^{\mathcal{P}}} \mathcal{L}^{\mathcal{P}}$

 /* Calculate median of confidence values. */

$m = \text{median}(\{e^{-u_n}, n \in \mathcal{P}\})$

 /* Update \mathcal{P} by removing images with small confidence. */

for $n \in \mathcal{P}$ **do**

if $e^{-u_n} < m$ **then**

$\mathcal{P} \leftarrow \mathcal{P} \setminus n$

Our final loss is then $\mathcal{L}^{\mathcal{P}} = \lambda_{\text{RGB}} \mathcal{L}_{\text{RGB}}^{\mathcal{P}}(\hat{I}) + \lambda_{\text{RGB}} \mathcal{L}_{\text{RGB}}^{\mathcal{P}}(\hat{I}^{MV}) + \lambda_{\text{depth}} \mathcal{L}_{\text{depth}}^{\mathcal{P}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}^{\mathcal{P}} + \lambda_{\text{dist}} \mathcal{L}_{\text{dist}}^{\mathcal{P}}$, which is optimized over the MLP parameters $\Theta = \{\Theta^\sigma, \Theta^c, \Theta^{MV}\}$ and the uncertainty weights $\mathbf{U}^{\mathcal{P}} = \{u_n, n \in \mathcal{P}\}$. The confidence of all images is initialized to 1, *i.e.*, $u_n := 0$.

Iterative refinement. We use the predicted per-image uncertainty, u_n , in an iterative scheme that progressively removes non-confident images from the NeRF optimization, *i.e.*, we iteratively update the set \mathcal{P} of images that contribute to the loss in masked regions. After K_{grad} steps of optimizing $\mathcal{L}^{\mathcal{P}}$, we find the median estimated confidence value m . We then remove from the training set all 2D *inpainted regions* which have associated confidence scores less than m . We then retrain the NeRF with the updated training set, and repeat these steps K_{outer} times. Note that images excluded from \mathcal{P} still participate in the optimization, but only for rays in the unmasked regions as they contain valuable information about the scene. This is summarized in Algorithm 1.

3.4. Implementation details

Masking the object to be removed. Similarly to other inpainting methods, our method requires per-frame masks as input. Manually annotating each frame with a 2D mask, as done in other inpainting methods [30, 67, 73], is time consuming. Instead, we propose to manually annotate a 3D box that contains the object, which only has to be done once per scene. We reconstruct a 3D point cloud of the scene using the camera poses and the input depth maps. This point cloud can then be imported into a 3D visualisation and editing software, such as MeshLab [10], where we specify a 3D bounding box enclosing the object to be removed. Alter-




Figure 5. **Mask refinement** Our mask refinement leads to smaller masks and therefore higher quality inpainting.

natively, we could have relied on 2D object segmentation methods, *e.g.*, [16], or 3D object bounding box detectors, *e.g.*, one of the baselines in [1], to mask the object.

Mask refinement. In practice, we observe that masks obtained from the annotated 3D bounding boxes can be quite coarse and include large parts of the background. Since large masks have a negative effect on the inpainting quality, we propose a mask refinement step to obtain masks which are tighter around the object. This step is not required for tight input masks. Intuitively, this mask refinement step removes parts of the 3D bounding box that are empty space. We start by taking all points in the reconstructed 3D point cloud that are inside the 3D bounding box. The refined mask is then obtained by rendering these points into each image and performing a simple comparison with the depth map to check occlusions in the current image. The resulting mask is cleaned up by dilating any pixel leaks caused by sensor noise using binary dilation and erosion. The effect of our mask refinement step can be seen in Figure 5.

Inpainting network. We used [67] for inpainting both RGB and depth. The inpainting of RGB images and depth maps is done independently and we use the reference network provided by the authors of [67] for both. Our depth maps are preprocessed by clipping to 5 m and linearly mapping depths in [0 m, 5 m] to pixel values of [0, 255]. This matrix is duplicated to make a $H \times W \times 3$ tensor for input to [67]. We observed empirically that this approach provided good results, but an inpainting method specific for depth maps could improve over this baseline.

NeRF estimation. The implementation of our method is built upon [47] and [3]. The shared MLP F_Θ^σ consists of eight 256-wide layers while the branches F_Θ^c and F_Θ^{MV} consist of four 128-wide layers each. The final output is activated using *softplus* for density, ReLU for the uncertainty, and sigmoid for the color channel. We weight the terms in the loss function with $\lambda_{\text{RGB}} = \lambda_{\text{depth}} = \lambda_{\text{dist}} = 1$, and $\lambda_{\text{reg}} = 0.005$. We optimize using Adam with an initial learning rate of $l_r = 0.0005$, and use the scheduler from [47]. We do a filtering step, where we remove low confidence images every $K_{\text{grad}} = 50,000$ steps, resulting in $K_{\text{outer}} = 4$ filtering steps. More details concerning implementation are in the supplementary material.

	Synthetic objects - Masked					Real objects - Masked				
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Depth L_1 \downarrow	Depth L_2 \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Depth L_1 \downarrow	Depth L_2 \downarrow
No novel view synthesis										
LaMa [67] \dagger	27.999	0.898	0.060	<u>0.070</u>	0.075	-	-	-	-	-
E ² FGVI [30] $\dagger\dagger$	24.568	0.874	0.102	-	-	-	-	-	-	-
NeRF based ablations										
Masked NeRF	26.126	0.882	0.093	0.084	0.105	21.644	0.815	0.142	<u>0.096</u>	<u>0.054</u>
Inpainted NeRF	<u>28.760</u>	<u>0.905</u>	0.086	0.278	0.400	23.705	0.848	0.134	0.145	0.121
Inpainted NeRF + inpainted Depth	27.568	0.898	0.094	0.231	0.318	23.652	0.844	0.136	0.313	0.387
3D scene completion										
PixelSynth [54] \ddagger	25.481	0.887	0.116	-	-	25.438	0.851	0.152	-	-
CompNVS [29]	17.389	0.823	0.171	1.697	3.641	-	-	-	-	-
Object compositional NeRF [75]*	-	-	-	-	-	21.757	0.836	0.134	0.312	0.341
Our method										
Ours	29.437	0.916	<u>0.078</u>	0.069	<u>0.096</u>	<u>25.271</u>	0.859	0.125	0.071	0.044
Ours – depth predicted using [60]	26.540	0.895	0.087	0.112	0.118	25.010	<u>0.856</u>	<u>0.128</u>	0.142	0.140

Table 1. **Comparison with baselines and state of the art methods.** Our method is either **best** or **second-best** compared to other novel-view synthesis baselines in inpainting the missing regions of the scene, by propagating multi-view information and leveraging 2D inpainting information. Notes: \dagger These methods can’t be evaluated on the proposed real dataset as they do not synthesize novel views. \ddagger These methods do not produce depth maps. * [75] requires the actual object therefore it can’t be evaluated on the proposed synthetic dataset.

4. Experiments

4.1. Datasets

While previous approaches have tackled static object removal from videos, no standard dataset/metrics to evaluate these systems has been proposed, to our knowledge. This work introduces an RGB-D dataset of real scenes, designed to evaluate the quality of object removal. Our dataset will be made public and has two variants, which are used differently when benchmarking:

Ours — Real objects. This dataset comprises 17 scenes, which are a mixture of indoors/outdoors with landscape/portrait orientation, and focus on a small area with at least one object, with one of them being the labeled object of interest. They vary in difficulty in terms of background texture, size of the object, and complexity of scene geometry. Please see supplementary material for a visualization of the full dataset. For each scene, we collected two sequences, one with and the other without the object that we want to remove. The sequences are collected using ARKit [2] on an iPhone 12 Pro with Lidar, and contain RGB-D images and poses. The RGB images are resized to 192×256 , the same resolution as the depth maps. Both sequences for the same scene use the same camera coordinate system. The sequences vary in length, from 194 to 693 frames.

As previously described, the masks are obtained by annotating a 3D bounding box of the object of interest and are refined for all of the scenes. For each scene, we use the sequence with the object and corresponding masks for training the NeRF model, and the sequence without the object for testing. The use of real objects makes it easier to evaluate how the systems deal with real shadows and reflections, as well as novel view synthesis.

Ours — Synthetic objects. Most video and image inpaint-

ing methods, *e.g.*, [30, 67], do not perform novel view synthesis, meaning such methods cannot be fairly evaluated on our ‘Real objects’ dataset. We therefore introduce a separate synthetically-augmented variant of our dataset. This uses the same scenes as the real objects dataset, but we only use the sequence that *doesn’t* contain the object. We then manually position a 3D object mesh from ShapeNet [7] in each scene. The object is placed so that it has a plausible location and size, *e.g.*, a laptop on a table. The masks are obtained by projecting the mesh into the input images, which is the only use we make of the 3D object mesh. Please see supplementary material for a visualisation of the full synthetic dataset. For this synthetic dataset, following *e.g.*, [4], we use every 8th frame for testing and the rest of the frames for training the NeRF model.

ARKitScenes. We further validate our approach qualitatively on ARKitScenes [5]. This is an RGB-D dataset of 1,661 scenes, where depth was captured via iPhone Lidar.

4.2. Metrics

To evaluate the object removal and inpainting quality, we compare our system’s output image against the ground truth image, for each test image in the dataset. All metrics in the paper are only computed inside the masked region. Metrics for the full image are provided in the supplementary material. We use the three standard metrics for NeRF evaluation [41]: PSNR [19], SSIM [71] and LPIPS [81]. To evaluate the geometric completion, we compute the L_1 and L_2 error between the rendered and the ground-truth depth maps inside masked regions. The metrics are averaged over all frames of a sequence, and then averaged over all sequences.



Figure 6. **Results on ARKitScenes [5].** We can successfully remove objects from casually captured sequences in indoor scenes.

4.3. Comparison with baselines

In Table 1, we compare our approach with alternative methods for object removal, with a focus on methods that use an underlying NeRF representation.

Image and video inpainting baselines. We compare with two state-of-the-art methods for image *LaMa* [67] and video inpainting *E²FGVI* [30]. In both cases, we use their reference implementation and provided trained network. Neither of these methods allows novel view synthesis, so they are only evaluated on the synthetic objects dataset.

NeRF-based ablations. We compare different ways of training a baseline NeRF model. The *Masked NeRF* baseline corresponds to training a NeRF using the full input RGB-D data, but pixels and depths in the masked regions are ignored in the NeRF losses. *Inpainted Images* is a NeRF trained with all inpainted images, but not using the inpainted depth maps, while *Inpainted Images + Inpainted Depth* uses all the inpainted images and inpainted depth maps. This baseline corresponds to *All views* in Table 2.

3D scene completion. We compare with several published works for 3D scene completion. Note that none of these baselines specifically targets inpainting, so results should be viewed with this in mind. For all of them, we use their publicly available implementation. *PixelSynth* [54] and *CompNVS* [29] were both proposed for scene *outpainting*. Given one or a few frames of a scene their goal is to complete the scene to enable novel view synthesis. Both of these methods rely on a generative model of indoor scenes and neither requires test-time optimization. Both methods are adapted to use our masks as input. We show qualitative results for *CompNVS* in the supplementary materials. *Object compositional NeRF* [75] enables objects in NeRFs to be edited via user-editing of pose transformations. We adapt

Method	# of views	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	$L_1\downarrow$	$L_2\downarrow$
All views	82 - 382	27.568	0.898	0.094	0.231	0.318
1/10th	8 - 38	27.098	0.900	0.079	0.202	0.291
1/50th	1 - 7	26.718	0.893	0.087	0.229	0.309
Single view	1	26.232	0.892	0.079	0.133	0.198
Ours	10 - 185	29.437	0.916	0.078	0.069	0.096

Table 2. **Ablation on view selection methods.** We validate our view selection formulation by comparing to alternative approaches. Ours consistently produces better performing models.

their code for object *removal* by setting a transformation that moves the object outside the camera’s field of view.

Our method. We compare two different versions of our method. *Ours* is our proposed approach, which uses the method described in Section 3. *Ours – depth from [60]* is our method, where the depth maps are obtained using a state-of-the-art multi-view depth prediction method [60] to show that we do not necessarily rely on sensor depth.

As shown in Table 1, our method is superior to other novel-view synthesis baselines across most appearance and depth metrics. Moreover, as opposed to the single image inpainting *LaMa* [67], our method is close to multi-view consistent, significantly reducing inter-frame flickering. To scrutinize the flickering, we refer to the supplementary video.

Qualitative comparison. In Figure 7, we show that the proposed method successfully removes the selected object compared to the baselines. While *Masked NeRF* fails to complete large holes and *Inpainted NeRF* suffers from bad inpaintings in the training set, our method can leverage the 2D inpaintings, while avoiding integrating artefacts by removing those input frames. In contrast to *Object compositional NeRF* [75], leveraging the inpaintings also helps to mitigate the appearance of artefacts below the object’s surface. Compared to [75] and [54], our method is better able to generate plausible scene completions. We also show results from the ARKitScenes dataset in Figure 6, and please also see our supplementary video.

4.4. View selection ablation

Here we validate that our view selection procedure from Section 3.3 contributes to improved results. We compare our method with different view selection strategies in Table 2 on our synthetic object dataset. *All views* uses all the inpainting views when training the NeRF model. The other baselines use a subset of views to train the NeRF model, spaced at regular intervals: every 10th frame for *1/10th*; every 50th frame for *1/50th* and a single middle frame for *Single view*. The number of views used for each sequence varies depending on the length of the sequence. We outperform these baselines, suggesting that our proposed strategy for view selection is effective in choosing a good set of views to include in the NeRF optimization.




Figure 7. **Qualitative comparisons with baseline.** Our method significantly improves over 3D scene completion baselines for the task of removing objects from a scene using 3D inpainting. Further, it mitigates artefacts and stabilizes convergence compared to the inpainting baseline without automatic view selection. †: We also compare with two NeRF based baselines: *Masked NeRF* and *Inpainted NeRF*.

4.5. Limitations

Our method is upper bounded by the performance of the 2D inpainting method. When the masks are too large along the entire trajectory, the 2D inpainting fails and no realistic views can be selected. Our renderings sometimes suffer from blurring, caused by flickering of high-frequency textures in the 2D inpaintings. Furthermore, cast shadows or reflections of the object are not handled well. We leave tackling these challenges for future work.

5. Conclusion

We have presented a framework to train neural radiance fields, where objects are plausibly removed from the output renderings. Our method draws upon existing work in 2D inpainting, and introduces an automatic confidence-based view selection scheme to select single-view inpaintings with multi-view consistency. We experimentally vali-




Figure 8. **Failure cases and limitations.** Our method can not recover when the 2D inpainting method fails all the frames, for example when the mask covers a large part of the image. Further, our method keeps the shadows of the removed object, if they are not included in the object mask.

dated that our proposed method improves novel-view synthesis from 3D inpainted scenes compared to existing work. We have also introduced a dataset for evaluating this work, which sets a benchmark for other researchers in the field.

References

- [1] Adel Ahmadyan, Liangkai Zhang, Artsiom Ablavatski, Jianing Wei, and Matthias Grundmann. Objectron: A large scale dataset of object-centric videos in the wild with pose annotations. In *CVPR*, 2021. (page 5)
- [2] Apple. ARKit. Accessed: 14 October 2022. (pages 3 and 6)
- [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. In *ICCV*, 2021. (pages 1, 2, 4, and 5)
- [4] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, 2022. (pages 1, 2, 3, and 6)
- [5] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, and Elad Shulman. ARK-itscenes - a diverse real-world dataset for 3D indoor scene understanding using mobile RGB-D data. In *NeurIPS*, 2021. (pages 6 and 7)
- [6] Borna Bešić and Abhinav Valada. Dynamic object removal and spatio-temporal RGB-D inpainting via geometry-aware adversarial learning. *IEEE Transactions on Intelligent Vehicles*, 7(2), 2022. (page 2)
- [7] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. (page 6)
- [8] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. MVSNeRF: Fast generalizable radiance field reconstruction from multi-view stereo. In *ICCV*, 2021. (pages 1 and 2)
- [9] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3D shape reconstruction and completion. In *CVPR*, 2020. (page 2)
- [10] Paolo Cignoni, Massimiliano Corsini, and Guido Ranzuglia. MeshLab: an open-source 3D mesh processing system. *ERCIM News*, (73), 2008. (page 5)
- [11] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. In *Computer Graphics Forum*, volume 31, pages 305–314. Wiley Online Library, 2012. (page 2)
- [12] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised NeRF: Fewer views and faster training for free. In *CVPR*, June 2022. (pages 1 and 2)
- [13] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W Taylor, and Joshua M Susskind. Unconstrained scene generation with locally conditioned radiance fields. In *ICCV*, 2021. (pages 2 and 3)
- [14] Ryo Fujii, Ryo Hachiuma, and Hideo Saito. RGB-D image inpainting using generative adversarial network with a late fusion approach. In *International Conference on Augmented Reality, Virtual Reality and Computer Graphics*, pages 440–451. Springer, 2020. (page 2)
- [15] Chen Gao, Ayush Saraf, Jia-Bin Huang, and Johannes Kopf. Flow-edge guided video completion. In *ECCV*, 2020. (page 2)
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. (page 5)
- [17] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (TOG)*, 37(6):1–15, 2018. (page 2)
- [18] Jan Herling and Wolfgang Broll. Advanced self-contained object removal for realizing real-time diminished reality in unconstrained environments. In *ISMAR*, 2010. (page 2)
- [19] Alain Hore and Djemel Ziou. Image quality metrics: PSNR vs. SSIM. In *ICPR*. IEEE, 2010. (page 6)
- [20] Ching-Yu Hsu, Cheng Sun, and Hwann-Tzong Chen. Moving in a 360 world: Synthesizing panoramic parallaxes from a single panorama. *arXiv preprint arXiv:2106.10859*, 2021. (page 2)
- [21] Jia-Bin Huang, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. Temporally coherent completion of dynamic video. *ACM Transactions on Graphics (TOG)*, 35(6), 2016. (page 2)
- [22] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 2017. (page 2)
- [23] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *NeurIPS*, 30, 2017. (page 4)
- [24] Joohyung Kim, Janghun Hyeon, and Nakju Doh. Generative multiview inpainting for object removal in large indoor spaces. *International Journal of Advanced Robotic Systems*, 2021. (page 2)
- [25] Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *CVPR*, 2022. (pages 1 and 2)
- [26] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing NeRF for editing via feature field distillation. In *NeurIPS*, 2022. (pages 1 and 2)
- [27] Jing Yu Koh, Honglak Lee, Yinfei Yang, Jason Baldrige, and Peter Anderson. Pathdreamer: A world model for indoor navigation. In *ICCV*, 2021. (page 2)
- [28] Vincent Lepetit, Marie-Odile Berger, and LORIA-INRIA Lorraine. An intuitive tool for outlining objects in video sequences: Applications to augmented and diminished reality. In *ISMAR*, pages 159–160, 2001. (page 2)
- [29] Zuoyue Li, Tianxing Fang, Zhenqiang Li, Zhaopeng Cui, Yoichi Sato, Marc Pollefeys, and Martin R. Oswald. CompNVS: Novel view synthesis with scene completion. In *ECCV*, 2022. (pages 3, 6, and 7)
- [30] Zhen Li, Cheng-Ze Lu, Jianhua Qin, Chun-Le Guo, and Ming-Ming Cheng. Towards an end-to-end framework for flow-guided video inpainting. In *CVPR*, 2022. (pages 5, 6, and 7)
- [31] Miao Liao, Feixiang Lu, Dingfu Zhou, Sibozhang, Wei Li, and Ruigang Yang. DVI: Depth guided video inpainting for autonomous driving. In *ECCV*, 2020. (page 2)

- [32] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *ECCV*, 2018. (page 2)
- [33] Shichen Liu, Shunsuke Saito, Weikai Chen, and Hao Li. Learning to infer implicit surfaces without 3D supervision. *arXiv:1911.00767*, 2019. (page 2)
- [34] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *ICCV*, 2021. (page 1)
- [35] Wen Liu, Zhixin Piao, Jie Min, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping GAN: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *ICCV*, 2019. (page 2)
- [36] Yu-Lun Liu, Wei-Sheng Lai, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Learning to see through obstructions. In *CVPR*, 2020. (page 2)
- [37] Yu-Lun Liu, Wei-Sheng Lai, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Learning to see through obstructions with layered decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11), 2021. (page 2)
- [38] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. SparseNeuS: Fast generalizable neural surface reconstruction from sparse views. In *ECCV*, 2022. (pages 1 and 2)
- [39] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, 2022. (page 2)
- [40] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *CVPR*, 2019. (page 2)
- [41] 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. (pages 1, 2, 3, and 6)
- [42] Arthur Moreau, Nathan Piasco, Dzmitry Tsishkou, Bogdan Stanciulescu, and Arnaud de La Fortelle. LENS: Localization enhanced by NeRF synthesis. In *Conference on Robot Learning*, 2022. (page 1)
- [43] Shohei Mori, Okan Erat, Wolfgang Broll, Hideo Saito, Dieter Schmalstieg, and Denis Kalkofen. InpaintFusion: incremental RGB-D inpainting for 3D scenes. *IEEE TVCG*, 2020. (page 2)
- [44] Shohei Mori, Dieter Schmalstieg, and Denis Kalkofen. Good keyframes to inpaint. *IEEE Transactions on Visualization and Computer Graphics*, 2022. (page 2)
- [45] Thomas Müller, Alex Evans, Christoph Schied, Marco Foco, András Bódis-Szomorú, Isaac Deutsch, Michael Shelley, and Alexander Keller. Instant neural radiance fields. In *ACM SIGGRAPH 2022 Real-Time Live! 2022*. (page 1)
- [46] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3D representations from natural images. In *CVPR*, 2019. (page 2)
- [47] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. RegNeRF: Regularizing neural radiance fields for view synthesis from sparse inputs. In *CVPR*, 2022. (pages 2 and 5)
- [48] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. In *CVPR*, 2021. (page 2)
- [49] Selcen Ozturkcan. Service innovation: Using augmented reality in the IKEA Place app. *Journal of Information Technology Teaching Cases*, 11(1), 2021. (page 1)
- [50] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. (page 2)
- [51] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. (page 2)
- [52] Julien Philip and George Drettakis. Plane-based multi-view inpainting for image-based rendering in large scenes. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pages 1–11, 2018. (page 2)
- [53] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. In *CVPR*, pages 12932–12942, 2022. (page 2)
- [54] Chris Rockwell, David F Fouhey, and Justin Johnson. Pixel-Synth: Generating a 3D-consistent experience from a single image. In *ICCV*, 2021. (pages 2, 3, 6, 7, and 8)
- [55] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *CVPR*, 2022. (pages 2 and 3)
- [56] Martin Runz, Kejie Li, Meng Tang, Lingni Ma, Chen Kong, Tanner Schmidt, Ian Reid, Lourdes Agapito, Julian Straub, Steven Lovegrove, et al. FroDO: From detections to 3D objects. In *CVPR*, 2020. (page 2)
- [57] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. (page 2)
- [58] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *CVPR*, 2020. (page 2)
- [59] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2021. (page 2)
- [60] Mohamed Sayed, John Gibson, Jamie Watson, Victor Prisacariu, Michael Firman, and Clément Godard. SimpleRecon: 3D reconstruction without 3D convolutions. In *ECCV*, 2022. (pages 6 and 7)
- [61] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. (pages 2 and 3)
- [62] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, 2016. (pages 2 and 3)

- [63] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3D-aware image synthesis. *NeurIPS*, 2020. (page 2)
- [64] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3D photography using context-aware layered depth inpainting. In *CVPR*, 2020. (page 2)
- [65] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew Davison. iMAP: Implicit mapping and positioning in real-time. In *ICCV*, 2021. (page 2)
- [66] Mohamad Zaidi Sulaiman, Mohd Nasiruddin Abdul Aziz, Mohd Haidar Abu Bakar, Nur Akma Halili, and Muhammad Asri Azuddin. Matterport: virtual tour as a new marketing approach in real estate business during pandemic COVID-19. In *International Conference of Innovation in Media and Visual Design (IMDES)*, 2020. (page 1)
- [67] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *WACV*, 2022. (pages 2, 4, 5, 6, and 7)
- [68] Theo Thonat, Eli Shechtman, Sylvain Paris, and George Drettakis. Multi-view inpainting for image-based scene editing and rendering. In *3DV*, 2016. (page 2)
- [69] Jonathan Tremblay, Moustafa Meshry, Alex Evans, Jan Kautz, Alexander Keller, Sameh Khamis, Charles Loop, Nathan Morrical, Koki Nagano, Towaki Takikawa, and Stan Birchfield. RTMV: A ray-traced multi-view synthetic dataset for novel view synthesis. *ECCV Workshop (Learn3DG ECCV)*, 2022. (page 2)
- [70] Qianqian Wang, Zhengqi Li, David Salesin, Noah Snavely, Brian Curless, and Janne Kontkanen. 3D moments from near-duplicate photos. In *CVPR*, 2022. (page 2)
- [71] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4), 2004. (page 6)
- [72] Qianyi Wu, Xian Liu, Yuedong Chen, Kejie Li, Chuanxia Zheng, Jianfei Cai, and Jianmin Zheng. Object-compositional neural implicit surfaces. In *ECCV*, 2022. (pages 1 and 2)
- [73] Rui Xu, Xiaoxiao Li, Bolei Zhou, and Chen Change Loy. Deep flow-guided video inpainting. In *CVPR*, 2019. (pages 2 and 5)
- [74] Xiaogang Xu, Ying-Cong Chen, and Jiaya Jia. View independent generative adversarial network for novel view synthesis. In *ICCV*, 2019. (page 2)
- [75] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *ICCV*, October 2021. (pages 1, 2, 6, 7, and 8)
- [76] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. (page 2)
- [77] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *CVPR*, 2018. (page 2)
- [78] 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 *CVPR*, 2022. (page 1)
- [79] Yanhong Zeng, Jianlong Fu, and Hongyang Chao. Learning joint spatial-temporal transformations for video inpainting. In *ECCV*, 2020. (page 2)
- [80] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Baining Guo. Learning pyramid-context encoder network for high-quality image inpainting. In *CVPR*, 2019. (page 2)
- [81] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. (page 6)
- [82] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew Davison. In-place scene labelling and understanding with implicit scene representation. In *ICCV*, 2021. (page 2)