

NeRFiller: Completing Scenes via Generative 3D Inpainting

Ethan Weber^{1,2} Aleksander Holynski^{1,2} Varun Jampani¹ Saurabh Saxena¹
 Noah Snavely¹ Abhishek Kar¹ Angjoo Kanazawa²
¹Google Research ²UC Berkeley

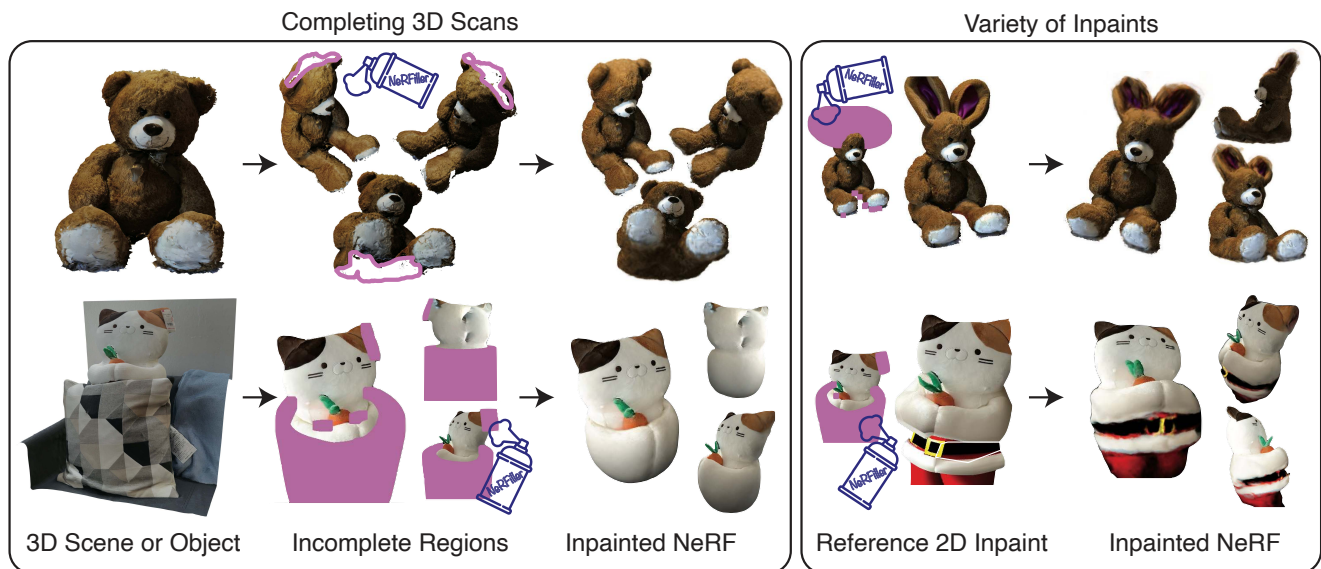


Figure 1. **NeRFiller**. We propose a generative 3D inpainting approach for scene or object completion. Given a 3D capture with incomplete regions (left), our approach completes scenes such as the incomplete teddy bear scan (top left) and deletes unwanted occluders such as the pillow and the price tag (bottom left). We can also control the completions using a reference inpainted exemplar (right) to guide the process.

Abstract

We propose *NeRFiller*, an approach that completes missing portions of a 3D capture via generative 3D inpainting using off-the-shelf 2D visual generative models. Often parts of a captured 3D scene or object are missing due to mesh reconstruction failures or a lack of observations (e.g., contact regions, such as the bottom of objects, or hard-to-reach areas). We approach this challenging 3D inpainting problem by leveraging a 2D inpainting diffusion model. We identify a surprising behavior of these models, where they generate more 3D consistent inpaints when images form a 2×2 grid, and show how to generalize this behavior to more than four images. We then present an iterative framework to distill these inpainted regions into a single consistent 3D scene. In contrast to related works, we focus on completing scenes rather than deleting foreground objects, and our approach does not require tight 2D object masks or text. We compare our approach to relevant baselines adapted to our setting on a variety of scenes, where *NeRFiller* creates the most 3D consistent and plausible scene completions. Our project page is at <https://ethanweber.me/nerfiller>.

1. Introduction

Consider the 3D scanned teddy bear and cat in Figure 1. In many 3D captures such as these, parts of the scene may not be as one desires: there may be unobserved regions such as the bottom of the bear and behind the cat, or there may be unwanted parts such as the price tag on the cat ear. Additionally, one may want to modify a feature, or generate a variety of alternative models, e.g., a bear with bunny ears or a santa cat. All of these tasks require the ability to edit and inpaint content in a 3D-aware and multi-view consistent manner. This is a challenge, since 2D generative inpainting models will not by default generate 3D consistent images. Our goal is to take a step in this direction and present a method that can create new content via scene completion conditioned on a set of multi-view images.

Specifically, we present a 3D scene completion framework called *NeRFiller*, which given a scene and specified parts of the scene to inpaint, returns a 3D scene that is completed in a multi-view consistent manner. Our approach not only completes missing regions (Figure 1, center), but can also generate multiple variations of the missing regions (Figure 1, right). Furthermore, our approach does not require




Figure 2. **Object removal vs. scene completion.** We focus on scene completion (right) as opposed to object removal (left). Prior work focuses on removing entire objects with tight masks, while we tackle the more general setting of completing scenes with arbitrary missing regions across wide baselines. More realistic scenarios include missing regions or parts of scenes to edit, as illustrated in Figure 1.

text prompting and can operate from the scene context alone.

We achieve this by proposing a novel approach to generate inpaints with an off-the-shelf 2D generative image model in a manner that encourages multi-view consistency. Specifically, we identify a useful phenomenon in text-to-image diffusion models that we refer to as a *Grid Prior*: denoising four images with missing observations that are tiled in a 2×2 grid results in more consistent multi-view inpaints than inpainting them independently, shown in Figure 3. We propose a method called *Joint Multi-View Inpainting* that generalizes this behavior to more than four images. While this technique results in more 3D consistent inpaints, it is still a 2D-based approach and 3D consistency is not guaranteed. Therefore, we propose a way to distill these inpaints in a global 3D scene representation in an iterative manner.

While there has been a surge of recent works that generate 3D scenes completely from scratch using text [14, 23] or image guidance [29, 30], our approach differs in that we focus on completing scenes given the context of an existing 3D scene. Our approach is related to recent methods that remove a specified object from a scene [37], but we can generate new content that goes beyond completing a textured background, as illustrated in Figure 2. We also do not assume a tight object mask, and can generate a diverse set of inpaints.

To demonstrate the efficacy of our approach, we experiment with a diverse set of scenes including 3D indoor photogrammetry captures lacking coverage in certain areas, 3D scenes with specified missing regions, and 3D objects. While our problem is challenging, we show that NeRFiller can recover more 3D consistent and plausible results compared to recent state-of-the-art methods adapted to our setting.

2. Related Work

Our goal is to complete missing parts of an existing 3D scene. There are several ways to approach this, via 2D inpainting or via distilling a 2D generative model for 3D generation.

2D inpainting. 2D inpainting methods take an image and mask and complete the missing content at the mask location. Early methods relied on inpainting by copying texture from known regions into the unknown regions [13]. A state-of-




Figure 3. **Grid Prior.** Here we inpaint the corner of the room (left illustrated in pink) with individual inpainting (top) and our Grid Prior method (bottom). Individual inpaints are diverse, while the Grid Prior encourages multi-view consistency.

the-art model is LaMa (Large Mask inpainting) [48], which is particularly good at infilling large missing areas. It uses fast Fourier convolutions, a large receptive field, and large training masks. This model is highly effective at completing plausible “background” textures within a specified mask (Fig. 2 left) but lacks diverse outputs as it is deterministic. Probabilistic diffusion models [22, 39] have recently produced remarkable results for image generation. They can also be used for inpainting and can generate diverse inpainted outputs. Pixel-based diffusion models do not have to be trained explicitly for inpainting, but can be modified at test-time by setting known regions before each denoising iteration [32]. Latent diffusion models (LDMs) [44] are also effective at inpainting and are efficient because they operate in latent space. However, they require fine-tuning for inpainting with image and mask conditioning. 2D inpainting models can be prompted [1] and/or fine-tuned [24, 50] enabling additional flexibility for downstream applications.

3D generation. 3D generation takes as input text or images and outputs 3D content. The Infinite Nature line of work [27, 29, 30, 42, 60], takes as input a single image and generates immersive fly-through content using a 2D inpainting model queried in an autoregressive manner [29, 30]. Cai et. al. [6] follow this path with a diffusion model, however, none of these approaches can recover a global 3D scene representation. Persistent Nature [7] and related work [9, 12] maintain a latent scene but are completely generative and not conditioned on input image sets. SceneScape [14] and Text2Room [23] use text prompts and 2D inpainters to create a 3D mesh by using an inpainter and depth predictor to successively stitch a mesh. These approaches cannot fix a mistake in the scene if a bad inpaint is made during the successive stitching because no global optimization is performed.

Other methods create 3D content via a global optimization strategy. DreamFusion [40] and related works [56, 63] use the NeRF framework to optimize a 3D volume given a text prompt. Others train models to have 3D consistent

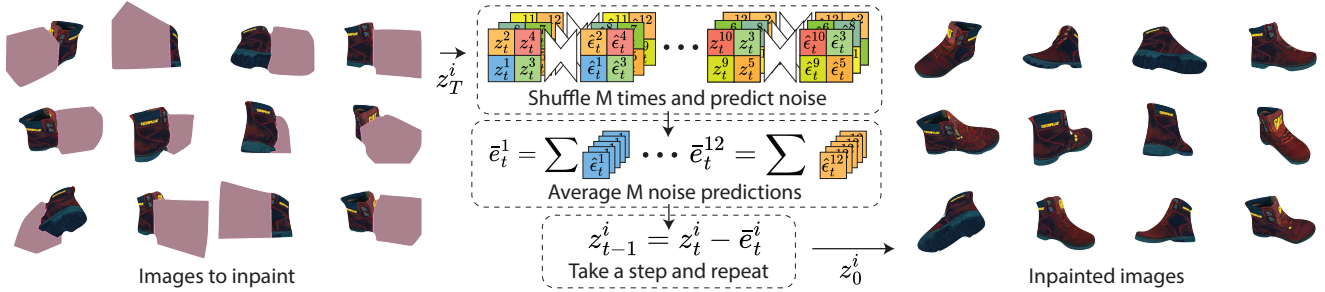


Figure 4. **Joint Multi-View Inpainting.** We use the Grid Prior with more than four images by averaging diffusion model predictions. We take N images (left), create $N/4$ grids, and obtain a noise prediction from SD [44]. We do this M times and average the noise predictions before taking a denoising step. At z_0 , the images (right) are fairly consistent and can be used to train a NeRF with our Inpaint DU method.

properties [28, 31, 51, 58, 63]. Follow-up works leverage 2D diffusion models techniques [15, 45] to create 3D content conditioned on real images [8, 33, 41]. These approaches are not designed for the inpainting task.

3D inpainting. Unlike most 3D generation methods, we ground our inpaints with an actual 3D scene or object that has missing regions (Figure 1). Casual capture [5, 19, 20, 34, 52] or NeRFs [26, 35] is a use-case as they often contain artifacts when rendered from novel views [17, 57]. Most relevant to scene the completion setting is the object removal setting (see Figure 2). These works remove foreground objects from NeRF captures [37, 53, 59]. They do this by inpainting each image in a NeRF dataset once and training with various losses including patch-based perceptual losses and depth regularization. [36] enables inpainting from a reference image. A variety of these methods are evaluated on the SPIN-NeRF dataset, which is in the forward-facing LLFF [34] format and has small parallax. Our work uses datasets with a significantly larger baseline.

Our focus is on the more general scene completion setting, which is related to editing. IN2N [18] edits a scene using InstructPix2Pix [4], but it cannot hallucinate new geometry. The video editing literature is also relevant, with techniques such as extended attention from Tune-A-Video [16, 61] to encourage consistency in edited video frames. However, editing 2D images does not guarantee consistency when lifted to 3D. In our method, we encourage 2D inpaints to converge via iterative NeRF optimization and dataset updates. Moreover, we achieve this with *off-the-shelf* 2D generative models without the need for expensive purpose-trained diffusion models or model fine-tuning.

3. Preliminaries

3.1. Neural Radiance Fields (NeRFs)

Neural radiance fields (NeRFs) [35] represent the 3D geometry and radiance of a scene with neural networks. NeRFs take as input a 3D position (x, y, z) and a viewing direction (θ, ϕ) , and output a color and density (c, σ) . To train a NeRF f_Θ , a set of calibrated, posed images are used to construct

a set of 3D rays $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ for each pixel with known color $C(\mathbf{r})$. During training, these rays are sampled and rendered via volumetric rendering to obtain a color estimate $\hat{C}(\mathbf{r})$. Rays are sampled from training images and the field is optimized with photometric losses $\mathcal{L}_{nerf}(C(\mathbf{r}), \hat{C}(\mathbf{r}))$, e.g., MSE or LPIPS [62]. During inference, a full image is rendered with all rays of the desired camera.

3.2. 2D Diffusion Models

Diffusion models consist of two processes: a forward process q that gradually adds noise to a data sample $z_0 \sim p_{data}(z)$, and a learned reverse process to iteratively denoise a pure Gaussian noise sample $z_T \sim \mathcal{N}(0, 1)$ into a clean image z_0 . An intermediate noisy z_t can be obtained from the clean image by adding noise ϵ with scaling $\bar{\alpha}_t$, where $z_t = \sqrt{\bar{\alpha}_t}z_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon$. The diffusion model ϵ_θ predicts noise $\hat{\epsilon}$ present in the image z_t as $\hat{\epsilon} = \epsilon_\phi(z_t, t, c)$. t is a time indicating how much noise is in the sample, and c is a general form of conditioning (e.g., images, masks, or text). During training, random noise ϵ and t are sampled and the objective $\mathcal{L}_{diff} = \|\hat{\epsilon} - \epsilon\|^2$ is minimized. With the prediction $\hat{\epsilon}$ at time t , a reduced noise z_{t-1} can be obtained by $z_{t-1} = z_t - \hat{\epsilon}$ (where we omit the scaling of $\hat{\epsilon}$ for simplicity). Repeating this until z_0 yields a fully denoised sample. Stochastically training with c (conditionally) and without c (unconditionally) enables classifier-free guidance (CFG) [21] during inference time. In practice, z is latents since we are using SD (Stable Diffusion) [44], but in general we can map from z to higher resolution pixels x with an encoder $\mathcal{E}(x)$ and decoder $\mathcal{D}(z)$.

Using a diffusion model as a prior. Diffusion models have advantages over other models (e.g., GANs and deterministic inpainters [48]) because they can be used a a prior to optimize underlying variables such as the parameters of a 3D NeRF f_Θ with methods like score distillation sampling (SDS) [40, 55]. When used as a prior for NeRFs, the objective is to find the best Θ such that a rendered image x has high likelihood under the diffusion model prediction ϵ_ϕ . SDS involves rendering an image, adding partial noise, and updating the NeRF such that the diffusion model can pre-



Figure 5. **Joint Multi-View Inpainting Examples.** The top images are inpainted with SD [44] without any text conditioning (middle) and with our Joint Multi-View Inpainting method (bottom). Our joint inpaints are more multi-view consistent.

dict the added noise. IN2N [18] introduced a variant of this method coined Dataset Update (DU). Instead of backproping based on the diffusion model prediction, DU renders an image, adds partial noise, and takes multiple steps to recover an estimated clean image x_0 . The clean image is added to the dataset and used to supervise the NeRF. Every S iteration, another image is replaced. The DU supervision signal will be slightly delayed since images are cached for several iterations, while SDS provides immediate gradients corresponding to the current render. However, we use the DU method in our work because it has a few advantages over SDS in terms of implementation: We can obtain higher-resolution supervision (albeit slightly delayed) with less GPU memory and we can update a large batch of images simultaneously (e.g., 40 images). We find that large batch updates are important for our inpainting task since we are changing the NeRF geometry, unlike prior works that focus purely on modifying appearance [18, 38].

4. Method

Our method, NeRFiller, aims to complete a missing region within a 3D scene by using an inpainting prior from a generative 2D diffusion model. This problem statement poses a number of challenges. First, the inpainted estimates from a 2D diffusion model are diverse, and may vary from sample to sample. This requires a consolidation mechanism to ensure that the completed 3D scene contains one salient inpainted result, as opposed to the average of all possibilities. Second, 2D inpainting models are not trained for 3D consistency, and will therefore provide estimates that cannot be explained by a single 3D scene, even if they correspond to the same approximate style or content. In the following, we describe our approaches to tackle these problems. In Section 4.1, we describe how we encourage the inpainted outputs from a diffusion sampling process to be 3D consistent. In Section 4.2, we describe an iterative 3D scene optimization method that uses these inpainted images to optimize for a globally consistent inpainted 3D scene. NeRFiller builds on an observation that inpainting a grid of images encourages the outputs to have similar appearance, and we extend this idea to an arbitrarily large collection of images through a joint sampling approach.

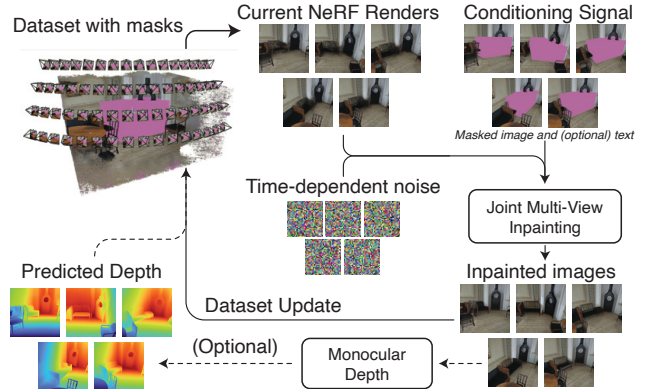


Figure 6. **Inpaint Dataset Update.** Every S iterations, we update the unknown pixels of the NeRF training images. We render N images, add partial noise, and jointly inpaint with a conditioning signal. We (optionally) predict the depth and update the dataset.

4.1. Multi-view consistent inpainting

A core challenge in 3D inpainting with a 2D generative model is getting the outputs of the 2D model to be consistent across views. This challenge stems from the multimodality of the output distribution: in most cases, there are many plausible inpaintings, and sampling multiple consistent images remains an open research problem.

Grid Prior. While inpainting multiple viewpoints independently may produce inconsistent results, one interesting discovery is that consistency can be achieved by tiling the input images into a grid and treating the *grid* of images (and their corresponding masks) as a single inpainting target. This grid-based prior can produce more 3D consistent views, both in coarse appearance and approximate scene structure (illustrated in Figure 3). We hypothesize that this phenomenon results from similarly structured examples in Stable Diffusion’s training dataset: sets of observations depicting the same scene or object organized as a grid (e.g., screenshots of online product photos). Similar properties were also explored in visual-prompting [1].

More specifically, in order to inpaint four images consistently, one can downsample them and their corresponding inpainting masks to quarter-resolution and tile them as a 2×2 grid. This grid is fed through the 2D inpainting model (as a single image would) to get as output four inpainted




Figure 7. **Inpainting methods.** Inpainting methods produce inconsistent inpaints. We show various inpainting methods (boxed) and use a collection of them to train a NeRF. A resulting render is shown on the right of each method. Using our Grid Prior and Joint Multi-View Inpainting creates reasonably consistent inpaints and a plausible NeRF. \mathcal{X} means the NeRF failed and resulted in white everywhere.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Masked NeRF	7.76	0.71	0.37
LaMask	19.58	0.89	0.20
SD Text Cond	12.56	0.73	0.32
SD Image Cond	14.15	0.76	0.28
Extended Attention	14.57	0.77	0.27
Grid Prior	14.43	0.80	0.25
Joint Multi-View Inpainting	15.89	0.82	0.23

Table 1. **Multi-view consistent inpainting.** We inpaint images and train a NeRF for the 8 scenes of the NeRF synthetic dataset. Better metrics indicate more consistency of the NeRF 3D reconstruction with the 2D inpaints. Note that LaMask achieves the best results as it often copies the white background into the hole (see Figure 7) and results in a failed NeRF that is totally white.

images with consistent content. More formally, let \mathcal{G} be the downsampling and grid operation for four images and let \mathcal{G}^{-1} undo this. We can grid four images latents $z_t^1, z_t^2, z_t^3, z_t^4$ as follows:

$$\{\hat{\epsilon}_t^1, \hat{\epsilon}_t^2, \hat{\epsilon}_t^3, \hat{\epsilon}_t^4\} = \mathcal{G}^{-1}(\epsilon_\phi(\mathcal{G}(\{z_t^1, z_t^2, z_t^3, z_t^4\}))) \quad (1)$$

and take a denoising step with $z_{t-1}^i = z_t^i - \hat{\epsilon}_t^i$. In principle, this approach is similar to recent methods that use *extended attention*, i.e., shared keys and values in the attention operations across a set of parallel sampling processes [61]. Our approach does not share attention features but instead shares context with other images via the diffusion U-Net receptive field that sees 4 tiled images at a time. In our experiments, we compare to *extended attention* and demonstrate that our grid prior more effectively inpaints 3D consistent content when used with our Joint Multi-View Inpainting method.

Joint Multi-View Inpainting. While effective at inpainting a set of images consistently, applying the grid prior to a larger set of images poses additional challenges. Increasing the number of images in the grid proportionally decreases the output resolution of each image (e.g., arranging a set of 2×2 images in a grid reduces each image’s resolution by 4, a 3×3 grid by 9, and so on). For the purpose of producing high-quality inpainting results, we would like to minimize any loss in image detail. Therefore, we propose a method that uses the above grid prior in a joint sampling process, inspired by MultiDiffusion [2]. In each sampling step, we shuffle all the input images into a set of 2×2 grids. We repeat this M times and before taking a sampling step, the

score estimate for each image is combined across all the grid combinations in which it was seen. This causes the inpainting estimates to be gradually shared across the entire dataset, effectively increasing the grid size without further reducing effective resolution.

Figure 4 describes this procedure. More formally, for a batch of N images, we randomly permute their order, construct $N/4$ grids, and predict the noise for M iterations ($j \in [1, M]$), as follows:

$$\{\hat{\epsilon}_t^{1j} \dots \hat{\epsilon}_t^{Nj}\} = \mathcal{G}^{-1}(\epsilon_\phi(\mathcal{G}(\{z_t^{1j} \dots z_t^{Nj}\}))) \quad (2)$$

and step with $z_{t-1}^i = z_t^i - \sum_{j \in M} \hat{\epsilon}_t^{ij}$ from z_T to z_0 . Qualitative results are shown in Figure 5.

4.2. Completing 3D Scenes

The proposed Joint Multi-View Inpainting enables inpainting images in a more 3D consistent manner than other inpainting methods. Next, we describe how to distill these 2D inpainting results into a single 3D reconstruction. We refer to our method as Inpaint Iterative Dataset Update (Fig. 6), or **Inpaint DU**, as it derives from IN2N’s [18] Iterative DU method (see Sec. 3.2). In contrast to IN2N, which begins with a complete NeRF reconstruction and uses a model conditioned on complete 2D observations, our task requires us to train a complete NeRF from images with masked unknown regions. As in IN2N, we begin training with a dataset of original (known) pixels, and update the dataset over training by adding or replacing the set of initially unknown pixels with the inpainted estimates. Fig. 6 illustrates this procedure. Specifically, every S steps, we render the set of N training views, encode them into latents z_0^i and then partially noise them before feeding them to SD. We sample from these partially noised inputs using the proposed Joint Multi-View Inpainting strategy, then use the resulting images to replace the corresponding images in the dataset. This process is both prefixed and suffixed by an encode and decode operation, since the base model is a latent diffusion model. We repeat this process many times while linearly annealing t from full noise $t = 1$ to $t = t_{\min}$. In practice, we set $t_{\min} = 0.4$, since we find that low noise values result in quality degradation. We observe that over the course of optimization, our inpainted images become gradually more

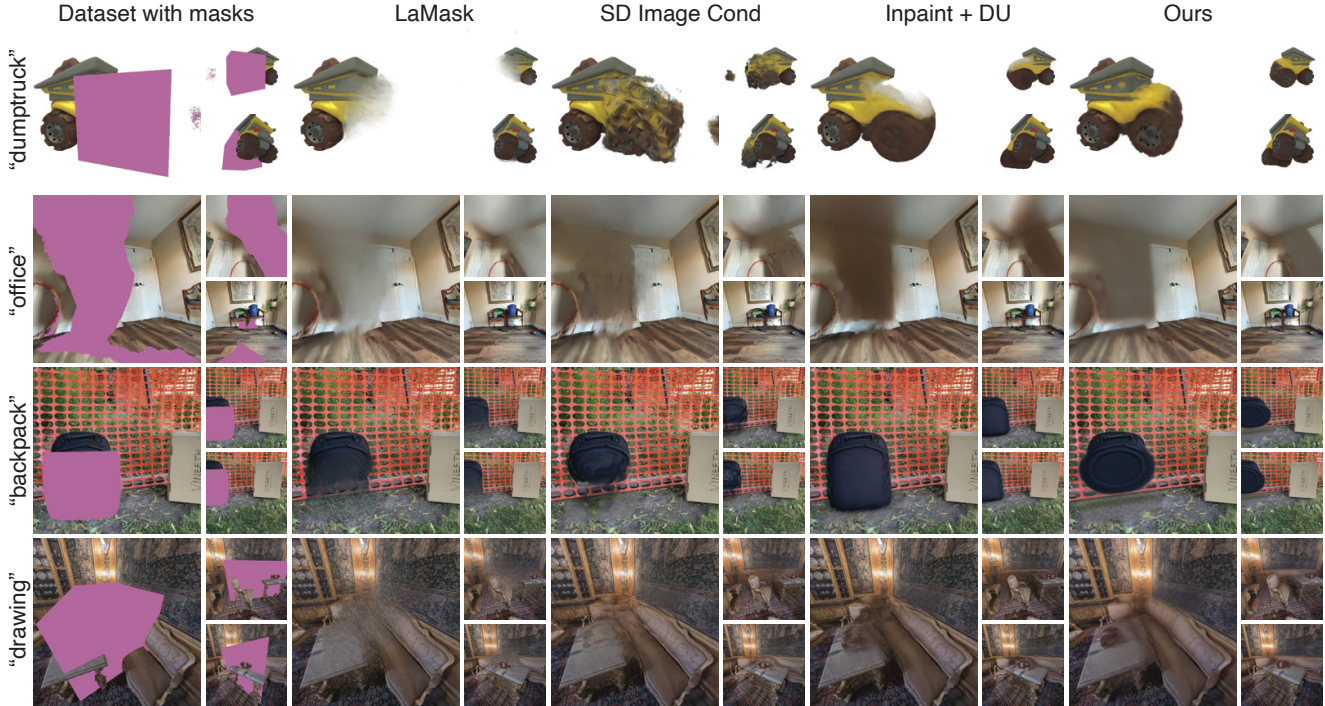


Figure 8. **Qualitative NeRF results.** On the left, we show various scenes with pink regions to be completed. We compare NeRFiller (far right) against baselines adapted to our scene completion setting. The “office” scene is missing parts of the wall, floor, and under the chairs.

consistent (Fig. 9) as the added geometry and texture begin to take form. Annealing t helps encourage the inpainting to converge to a single result rather than making large changes late in training.

Depth regularization. Inpaint DU optionally incorporates depth supervision to improve inpainted scene geometry. After each dataset update, we predict the depth for all images with ZoeDepth [3]. We use a relative depth ranking loss [54] in the inpainted regions (but not on the known pixels). We use a ranking loss because it’s a softer constraint than metric depth supervision, where errors in scale-and-shift alignment could more easily harm the 3D scene geometry. *We only apply depth supervision for our main method to indoor scenes and not objects, since we empirically noticed that [3] performs less consistently when the background is a solid color (e.g., white or black).*

5. Experiments

We compare NeRFiller for 3D scene completion to various inpainting baselines. We first investigate various inpainting strategies on multi-view synthetic scenes to gauge how effective a deterministic inpainter [48] is compared to SD [44], sampled in various ways. After establishing that our *Joint Multi-View Inpainting* demonstrates multi-view inpainting properties, we evaluate our full method on 10 scans with missing regions. We compare NeRFiller to various object-removal baselines, *adapted to our setting*, to complete

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MUSIQ \uparrow	Corrs \uparrow
Masked NeRF	14.71	0.78	0.26	3.71	675
LaMask	27.39	0.90	0.05	3.76	643
SD Image Cond	22.03	0.86	0.11	3.68	665
Inpaint + DU	26.60	0.89	0.08	3.76	660
Ours w/o depth	28.41	0.92	0.06	3.72	682
Ours	28.28	0.91	0.06	3.73	696

Table 2. **Quantitative NeRF results.** We report various metrics averaged over our 10 scenes to quantify consistency of the 3D NeRF with the dataset inpaints (left), as well as novel-view metrics (right), such as “Corrs” (number of high-quality correspondences between random pairs of frames) for geometry.

missing regions. Finally, we analyze the parameters of our method and show an application to reference-guided scene completion. We conduct experiments with Nerfstudio [49] and provide implementation specifics in the appendix.

5.1. 3D consistent image inpainting

Our goal is to evaluate various 2D inpainting models and strategies to quantify their 3D consistency.

Setting and evaluation. For these experiments, we take the testing split of the NeRF synthetic dataset [35] (8 scenes of 200 images each). We resize each image to 512×512 resolution and mask out the center of each image with a 256×256 region to inpaint (see Figure 7 top left). We inpaint all images with various methods and train a Nerfacto NeRF model [49] on 180 equally spaced images and evaluate metrics on the remaining 20 images. We use the standard NeRF metrics




Figure 9. **Noise schedule.** We anneal the amount of noise we add to the NeRF renders when making a dataset update, while I-N2N chooses random noise each time. Annealing the noise produces sharper results (top) while I-N2N’s noise schedule introduces significant blur (bottom).

because they capture how similar the 3D reconstruction is to the 20 hold-out evaluation images. When rendering for evaluation (right of inpainted images in Figure 7), we push the near plane slightly forward to avoid including any floaters hiding in front of the cameras.

Baselines. Our baselines are the following:

- *Masked NeRF* - No inpainting, only train on known pixels.
- *LaMask* - LaMa [48] inpainting model.
- *SD Text Cond* - SD using a text CFG with prompt “a photo of {description}”. See the appendix for text prompts.
- *SD Image Cond* - SD with only image CFG.
- *Extended Attention* - SD with only image CFG and extended attention [61].

For *Extended Attention* and *Grid Prior*, we inpaint in batches of 5 and 4, respectively, and for *Joint Multi-View Inpainting*, we inpaint 40 images simultaneously with $M = 8$ diffusion averaging steps. To inpaint additional batches of 40 images, we set 20 in the batch as known. This enables fitting within the memory constraints of a 16 GB GPU.

Results. Our results are shown in Tab. 1, where we see that *Joint Multi-View Inpainting* achieves the most favorable metrics in multi-view consistent inpainting. We note that *LaMask* has good metrics too, but this is due to failure cases described in the caption. Our results show that we have achieved some level of multi-view consistency. Our images look the most consistent in Figure 7 (bottom right), and the trained NeRF looks plausible. The other methods yield significant blur in the NeRF reconstruction. We provide videos on the project page showing the NeRF results for each method.

5.2. Completing large unknown 3D regions

In this setting, our goal is to complete missing regions in 3D content. We construct a set of 10 datasets consisting of various 3D content. For some scenes e.g., the backpack from [37], we modify their provided mask to include part of the object (Fig. 2) to convert it to the scene completion setting. For other scenes, we simply want to fill in any miss-



Figure 10. **Reference-based completion.** Given a reference inpaint (top row), we propagate it into a 3D NeRF (bottom row).

ing details (e.g., parts of walls). Some of the meshes are missing vertices after multi-view stereo reconstruction, e.g., “bear” and “office”. For others, we place a large 3D occluder in the scene to simulate the scene completion setting. The pink regions in Fig. 8 (left) shows the areas to complete. We create the datasets by rendering ~ 60 novel views looking at the occluded region. Importantly, the rendered images have enough known pixels to provide context to the inpainting model that the images observe the same scene from different camera viewpoints. Our datasets have much more parallax than the forward-facing scenes of [37]. Our appendix provides details on our data, including where we obtained our 3D content, mostly from Objaverse [10, 11] and Sketchfab.

Evaluation. The evaluation is similar to Sec. 5.1 for the dataset images. However, in this case our task is to construct a scene for good novel-view synthesis, so we use all images for both training and evaluation. We compute NeRF metrics on the entire images, where we compare the final rendered images with the latest version of the inpainted region. For methods that inpaint once without DU (e.g., *LaMask*), we compare against the first and only inpaints. For DU methods, we compare against the latest round of inpaints. *Our metrics are against inpainted images which serves to evaluate the consistency of the scene because there is no ground-truth solution.* We also report novel-view metrics, computed on a custom 10 second 30 FPS camera path novel views that moves around the scene. We also report an image quality metric MUSIQ [25] and a geometry metric. For geometry, we report the number of high-quality LoFTR [46] correspondences between 100 randomly sampled pairs of frames. More high-quality matches should correlate with better multi-view consistency and fewer extreme view-dependent effects that destroy realism. Please see the Appendix for more details.

Baselines. We implemented the following baselines:

- *Masked NeRF* - no inpainting, where we train a Nerfacto model [49] only in the known pixel locations,
- *LaMask* - Inpaint once with LaMa [48] and train with patch-based perceptual losses. This is our adaptation of SPIn-NeRF [37].
- *SD Image Cond* - Inpaint once with SD. This is similar to InpaintNeRF360 [53] but without text since we find in

Sec. 5.1 that text CFG produces very inconsistent inpaints.

- *Inpaint + DU* - An adaption of IN2N [18] for our setting, which inpaints one image at a time with the SD inpainting model and our annealed noise schedule.

Results. Some qualitative results are shown in Fig. 8 and full videos for all 10 scenes are provided in the appendix. *LaMask* and *SD Image Cond* are both inpaint-once methods and therefore create large blurry regions in the NeRF, but between the two, *LaMask* is smoother since its deterministic inpainter [48] is less creative than SD in its outputs. *LaMa* [48] tends to copy background textures into the mask region. From certain views, *Inpaint + DU* looks sharp due to inpainting individually at full resolution; however, it has geometric inconsistencies and view-dependent effects which are crisp from some angles and blurry in others. Our method looks the most consistent with plausible outputs, although its ability to create consistent high-frequency texture details may be improved. We provide quantitative results showing that our final renders are most similar with the latest round of inpaints (Tab. 2 left). For novel-view metrics, we obtain the most correspondences (Tab. 2 right).

5.3. Reference-based inpainting

In some situations it is desirable to have control over the content used to complete the scene. NeRFiller can be easily adapted to 3D inpaint with respect to a user-provided reference inpaint. To do this, we first inpaint from one view and use it to prompt our Grid Prior update method. We ensure that each grid has the one reference inpaint when passed through SD. This ensures that all U-Net predictions are influenced by the reference inpaint so that new inpaints are more likely to be consistent with the reference. Figure 1 and Figure 10 have examples.

5.4. Parameter choices

Noise schedule. It is important to anneal the amount of noise added the rendered images. We start by adding full noise (1.0) and decrease it to 0.4 over the 30K iterations of training. IN2N [18], in contrast, uses a random schedule of choosing between 0.98 and 0.02 each update. This likely works because the InstructPix2Pix model is image conditioned and geometry of the NeRF does not change much. In our case, we are changing the geometry and using their schedule leads to blurry results, shown in Figure 9.

Depth regularization. We find that adding depth ranking supervision [54] in *Ours* improves geometry but it hardly changes the quantitative or visual results. Our method without depth supervision is still favorable compared to the baselines. In Figure 11, we see that the geometry is significantly improved but the RGB NeRF renderings are nearly indistinguishable. Consequently, we use depth supervision on indoor scenes since having better geometry is favorable for downstream applications such as mesh export.




Figure 11. **Relative depth supervision.** Relative depth supervision cleans up geometry (top) without affecting visual quality (bottom).

6. Limitations

Low resolution and blur. Our method recovers coarse geometry quite well but struggles to recover high-resolution detail in regions far way from the training cameras. We suspect this is because our *Joint Multi-View Inpainting* method downsamples images to construct the 2×2 grids. Perhaps a post-processing method to fine-tune SD [47] or a GANeRF [43] like optimization could improve the fidelity, but recovering high-frequency details remains a fundamental issue for 3D generative methods, e.g., DreamFusion [40].

Inpainting NeRF casual captures. An application of our approach would be to inpaint deleted content from Nerfbusters [57] or Bayes’ Rays [17]. However, these masked-out regions are large (limiting scene context to an inpainter) and furthermore, their mask patterns cause SD [44] to fail without multiple iterations of dilation, as pointed out in Text2Room [23] and in our appendix. One could retrain SD with these mask distributions, but this is out of scope of our method which uses an *off-the-shelf* model.

7. Conclusion

In this paper, we propose a generative 3D inpainting method called NeRFiller, which leverages an *off-the-shelf* 2D inpainting model [44] to complete missing parts of 3D scenes and objects. We discover a unique property of these models where tiling four images into a 2×2 grid produces more consistent inpaints than inpainting them independently. We exploit this property and propose *Joint Multi-View Inpainting*, which enables inpainting many images simultaneously with more consistency by averaging noise predictions. We show how to use it in the NeRF setting by performing iterative dataset updates. We evaluate against relevant state-of-the-art baselines adapted to our problem setting on a variety of 3D captures. Our approach also enables users to specify how to fill in the missing regions. Many 3D captures are incomplete with holes, and our work presents a framework for completing these missing regions.

References

- [1] Amir Bar, Yossi Gandelsman, Trevor Darrell, Amir Globerson, and Alexei Efros. Visual prompting via image inpainting. In *ANeurIPS*, 2022. [2](#), [4](#)
- [2] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation. In *ICML*, 2023. [5](#)
- [3] Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. In *arXiv preprint arXiv:2302.12288*, 2023. [6](#)
- [4] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023. [3](#)
- [5] Michael Broxton, John Flynn, Ryan Overbeck, Daniel Erickson, Peter Hedman, Matthew DuVall, Jason Dourgarian, Jay Busch, Matt Whalen, and Paul Debevec. Immersive light field video with a layered mesh representation. In *SIGGRAPH*, 2020. [3](#)
- [6] Shengqu Cai, Eric Ryan Chan, Songyou Peng, Mohamad Shahbazi, Anton Obukhov, Luc Van Gool, and Gordon Wetzstein. Diffdreamer: Consistent single-view perpetual view generation with conditional diffusion models. In *ICCV*, 2023. [2](#)
- [7] Lucy Chai, Richard Tucker, Zhengqi Li, Phillip Isola, and Noah Snavely. Persistent nature: A generative model of unbounded 3d worlds. In *CVPR*, 2023. [2](#)
- [8] Eric R Chan, Koki Nagano, Matthew A Chan, Alexander W Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. Generative novel view synthesis with 3d-aware diffusion models. In *arXiv*, 2023. [3](#)
- [9] Zhaoxi Chen, Guangcong Wang, and Ziwei Liu. Scenedreamer: Unbounded 3d scene generation from 2d image collections. 2023. [2](#)
- [10] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. In *arXiv*, 2023. [7](#)
- [11] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *CVPR*, 2023. [7](#)
- [12] 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. [2](#)
- [13] Alexei A Efros and Thomas K Leung. Texture synthesis by non-parametric sampling. In *ICCV*, 1999. [2](#)
- [14] Rafail Fridman, Amit Abecasis, Yoni Kasten, and Tali Dekel. Scenescape: Text-driven consistent scene generation. In *arXiv*, 2023. [2](#)
- [15] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *arXiv*, 2022. [3](#)
- [16] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Tokenflow: Consistent diffusion features for consistent video editing. In *arXiv*, 2023. [3](#)
- [17] Lily Goli, Cody Reading, Silvia Sellán, Alec Jacobson, and Andrea Tagliasacchi. Bayes’ Rays: Uncertainty quantification in neural radiance fields. In *arXiv*, 2023. [3](#), [8](#)
- [18] Ayaan Haque, Matthew Tancik, Alexei Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *ICCV*, 2023. [3](#), [4](#), [5](#), [8](#)
- [19] Peter Hedman and Johannes Kopf. Instant 3D Photography. In *SIGGRAPH*, 2018. [3](#)
- [20] Peter Hedman, Suhil Alsian, Richard Szeliski, and Johannes Kopf. Casual 3D Photography. In *SIGGRAPH Asia*, 2017. [3](#)
- [21] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS Workshop on Deep Generative Models and Downstream Applications*, 2022. [3](#)
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. [2](#)
- [23] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *ICCV*, 2023. [2](#), [8](#)
- [24] Yash Kant, Aliaksandr Siarohin, Michael Vasilkovsky, Riza Alp Guler, Jian Ren, Sergey Tulyakov, and Igor Gilitschenski. invs: Repurposing diffusion inpainters for novel view synthesis. In *SIGGRAPH Asia Conference Proceedings*, 2023. [2](#)
- [25] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *ICCV*, 2021. [7](#)
- [26] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. In *SIGGRAPH*, 2023. [3](#)
- [27] Jing Yu Koh, Honglak Lee, Yinfei Yang, Jason Baldridge, and Peter Anderson. Pathdreamer: A world model for indoor navigation. In *ICCV*, 2021. [2](#)
- [28] Jialu Li and Mohit Bansal. Panogen: Text-conditioned panoramic environment generation for vision-and-language navigation. In *NeurIPS*, 2023. [3](#)
- [29] Zhengqi Li, Qianqian Wang, Noah Snavely, and Angjoo Kanazawa. Infineture-zero: Learning perpetual view generation of natural scenes from single images. In *ECCV*, 2022. [2](#)
- [30] Andrew Liu, Richard Tucker, Varun Jampani, Ameet Makadia, Noah Snavely, and Angjoo Kanazawa. Infinite nature: Perpetual view generation of natural scenes from a single image. In *ICCV*, 2021. [2](#)
- [31] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Learning to generate multiview-consistent images from a single-view image. In *arXiv*, 2023. [3](#)
- [32] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, 2022. [2](#)
- [33] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *CVPR*, 2023. [3](#)

- [34] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. In *SIGGRAPH*, 2019. 3
- [35] 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. 3, 6
- [36] Ashkan Mirzaei, Tristan Aumentado-Armstrong, Marcus A. Brubaker, Jonathan Kelly, Alex Levinshtein, Konstantinos G. Derpanis, and Igor Gilitschenski. Reference-guided controllable inpainting of neural radiance fields. In *ICCV*, 2023. 3
- [37] Ashkan Mirzaei, Tristan Aumentado-Armstrong, Konstantinos G. Derpanis, Jonathan Kelly, Marcus A. Brubaker, Igor Gilitschenski, and Alex Levinshtein. SPIn-NeRF: Multiview segmentation and perceptual inpainting with neural radiance fields. In *CVPR*, 2023. 2, 3, 7
- [38] Thu Nguyen-Phuoc, Feng Liu, and Lei Xiao. Snerf: stylized neural implicit representations for 3d scenes. In *SIGGRAPH*, 2022. 4
- [39] Ryan Po, Wang Yifan, Vladislav Golyanik, Kfir Aberman, Jonathan T Barron, Amit H Bermano, Eric Ryan Chan, Tali Dekel, Aleksander Holynski, Angjoo Kanazawa, et al. State of the art on diffusion models for visual computing. *arXiv preprint arXiv:2310.07204*, 2023. 2
- [40] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICML*, 2023. 2, 3, 8
- [41] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Ben Mildenhall, Nataniel Ruiz, Shiran Zada, Kfir Aberman, Michael Rubenstein, Jonathan Barron, Yuanzhen Li, and Varun Jampani. Dreambooth3d: Subject-driven text-to-3d generation. In *ICCV*, 2023. 3
- [42] Chris Rockwell, David F Fouhey, and Justin Johnson. Pixel-synth: Generating a 3d-consistent experience from a single image. In *ICCV*, 2021. 2
- [43] Barbara Roessle, Norman Müller, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, and Matthias Nießner. Ganerf: Leveraging discriminators to optimize neural radiance fields. In *SIGGRAPH Asia*, 2023. 8
- [44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3, 4, 6, 8
- [45] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubenstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 3
- [46] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *CVPR*, 2021. 7
- [47] Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. In *arXiv*, 2023. 8
- [48] 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. 2, 3, 6, 7, 8
- [49] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, et al. Nerfstudio: A modular framework for neural radiance field development. In *SIGGRAPH Conference Proceedings*, 2023. 6, 7
- [50] Luming Tang, Nataniel Ruiz, Qinghao Chu, Yuanzhen Li, Aleksander Holynski, David E Jacobs, Bharath Hariharan, Yael Pritch, Neal Wadhwa, Kfir Aberman, et al. Realfill: Reference-driven generation for authentic image completion. In *arXiv*, 2023. 2
- [51] Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. In *arXiv*, 2023. 3
- [52] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *CVPR*, 2020. 3
- [53] Dongqing Wang, Tong Zhang, Alaa Abboud, and Sabine Süsstrunk. Inpaintnerf360: Text-guided 3d inpainting on unbounded neural radiance fields. In *arXiv*, 2023. 3, 7
- [54] Guangcong Wang, Zhaoxi Chen, Chen Change Loy, and Ziwei Liu. Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. In *ICCV*, 2023. 6, 8
- [55] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *CVPR*, 2023. 3
- [56] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In *NeurIPS*, 2023. 2
- [57] Frederik Warburg*, Ethan Weber*, Matthew Tancik, Aleksander Holynski, and Angjoo Kanazawa. Nerfbusters: Removing ghostly artifacts from casually captured nerfs. In *ICCV*, 2023. 3, 8
- [58] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. In *arXiv*, 2022. 3
- [59] Silvan Weder, Guillermo Garcia-Hernando, Áron Monszpart, Marc Pollefeys, Gabriel Brostow, Michael Firman, and Sara Vicente. Removing objects from neural radiance fields. In *CVPR*, 2023. 3
- [60] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *CVPR*, 2020. 2
- [61] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, 2023. 3, 5, 7
- [62] 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. 3

- [63] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In *CVPR*, 2023. [2](#), [3](#)