

Blended-NeRF: Zero-Shot Object Generation and Blending in Existing Neural Radiance Fields

Ori Gordon

The Hebrew University

ori.gordon@mail.huji.ac.il

Omri Avrahami

The Hebrew University

omri.avrahami@mail.huji.ac.il

Dani Lischinski

The Hebrew University

danix@mail.huji.ac.il




Figure 1: **Overview of the training and blending process.** (a) Given a NeRF scene F_{θ}^O , our pipeline trains a NeRF generator model F_{θ}^G , initialized with F_{θ}^O weights and guided by a similarity loss defined by a language-image model such as CLIP [48], to synthesize a new object inside a user-specified ROI. This is achieved by casting rays and sampling points for the rendering process [39] only inside the ROI box. Our method introduces augmentations and priors to get more natural results. (b) After training, we render the edited scene by blending the sample points generated by the two models along each view ray.

Abstract

Editing a local region or a specific object in a 3D scene represented by a NeRF or consistently blending a new realistic object into the scene is challenging, mainly due to the implicit nature of the scene representation. We present *Blended-NeRF*, a robust and flexible framework for editing a specific region of interest in an existing NeRF scene, based on text prompts, along with a 3D ROI box. Our method leverages a pretrained language-image model to steer the synthesis towards a user-provided text prompt, along with a 3D MLP model initialized on an existing NeRF scene to generate the object and blend it into a specified region in the original scene. We allow local editing by localizing a 3D ROI box in the input scene, and blend the content synthesized inside the ROI with the existing scene using a novel volumetric blending technique. To obtain natural looking and view-consistent results, we leverage existing and new geometric priors and 3D augmentations for improving the visual fidelity of the final result. We test our framework both qualitatively and quantitatively on a variety of real 3D scenes and text prompts, demonstrating real-

istic multi-view consistent results with much flexibility and diversity compared to the baselines. Finally, we show the applicability of our framework for several 3D editing applications, including adding new objects to a scene, removing/replacing/altering existing objects, and texture conversion.

1. Introduction

In the last few years we have witnessed exciting developments in neural implicit representations [56, 60, 15, 61, 35, 62]. In particular, implicit representations of 3D scenes [57, 37, 55, 27, 46, 39, 5, 4] have enabled unprecedented quality and reliability in 3D reconstruction and novel view synthesis. The pioneering work of Mildenhall *et al.* [39] introduced NeRFs, MLP-based neural models that implicitly represent a scene as a continuous volume and radiance fields from a limited number of observations, producing high-quality images from novel views via volume rendering.

However, editing a scene represented by a NeRF is non-trivial, mainly because the scene is encoded in an im-




Figure 2: **Large object replacement.** Here we preform object replacement to the blender ship scene by localizing the ROI box to include the sea and the bottom of the ship and training our model to steer the edit towards the given text prompts.

PLICIT manner by the model’s weights, in contrast to explicit representations, such as meshes, voxel grids, or point clouds. NeRFs offer no explicit separation between the various components that define the object, such as shape, color, or material. In contrast to local edits in images, *e.g.*, [3, 2, 6, 42, 51, 23, 8], where the edit is done in pixel space with all the required information appearing in a single view, editing a NeRF-represented scene is more challenging due to the requirement for consistency across multiple views between the new and the original NeRF scenes.

The first works attempting to edit NeRF scenes focused on the removal of local parts, changing color, or shape transfer on one class of synthetic data, guided by user scribbles or latent code of another object in the class [34]. In CLIP-NeRF [64], editing of the entire scene is preformed by text guidance and displacements to the latent representation of the input. They mainly focus on synthetic objects from one class, or global color changes for realistic scenes. Kobayashi *et al.* [28] perform semantic decomposition of the scene components by learning a feature field that maps each 3D coordinate to a descriptor representing a semantic feature, and allow zero-shot segmentation for local editing on a specific semantic class. Alternatively, Benaim *et al.* [7] suggest separating the volumetric representation of a foreground object from its background using a set of 2D masks per training view. These works have limited localization abilities and focus on the separation methods. They demonstrate manipulations such as object removal, color change, and transformations such as shift, rotation, and scale.

In this work, we present our approach for ROI-based editing of NeRF scenes guided by a text prompt or an image patch that: (1) can operate on any region of a real-world scene, (2) modifies only the region of interest, while preserving the rest of the scene without learning a new feature space or requiring a set of two-dimensional masks, (3) generates natural-looking and view-consistent results that blend with the existing scene, (4) is not restricted to a specific class or domain, and (5) enables complex text guided manipulations such as object insertion/replacement, objects blending and texture conversion.

To this end, we utilize a pretrained language-image model, *e.g.*, CLIP [48], and a NeRF model [39] initialized on existing NeRF scene as our generator for synthesizing a new object and blend it into the scene in the region of interest (ROI). We use CLIP to steer the generation process towards the user-provided text prompt, enabling blended generation of diverse 3D objects.

To enable general local edits in any region, while preserving the rest of the scene, we localize a 3D box inside a given NeRF scene. To blend the synthesized content inside the ROI with the base scene, we propose a novel volumetric blending approach that merges the original and the synthesized radiance fields by blending the sampled 3D points along each camera ray.

We show that using this pipeline naively to perform the edit is insufficient, generating low quality incoherent and inconsistent results. Thus, we utilize the augmentations and priors suggested in [26] and introduce additional priors and augmentations, such as depth regularization, pose sampling, and directional dependent prompts to get more realistic, natural-looking and 3D consistent results. Finally, we conduct extensive experiments to evaluate our framework and the effect of our additional constraints and priors. We perform an in-depth comparison with the baseline and show the applicability of our approach on a series of 3D editing applications using a variety of real 3D scenes.

2. Related Work

Neural Implicit Representations have gained much popularity in the fields of computer vision and graphics in both 2D and 3D [56, 57, 55, 46, 37, 60, 15, 27]. Among their advantages is their ability to capture complex and diverse patterns and to provide a continuous representation of the underlying scene. They are resolution independent, yet compact, compared to explicit representations of high resolution 2D images, or meshes and point clouds in 3D. NeRFs [39, 4, 5] learn to represent a 3D scene as a continuous volume and radiance fields using the weights of a multilayer perceptron (MLP). Given a 3D position x and view direction (θ, ϕ) , NeRF outputs the density σ and color c at x .

Novel views of the scene can thus be rendered by accumulating the colors and densities along a view ray $\mathbf{r}(t)$ passing through each pixel, using an approximation to the classical volume rendering equation using the quadrature rule [36]:

$$C(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) c_i, T_i = \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j) \quad (1)$$

where $\delta_i = t_{i+1} - t_i$ is the distance between adjacent samples and T_i can be interpreted as the degree of transmittance at point x_i along the ray. The inputs are embedded into a high-dimensional space using a high frequency sinusoidal positional encoding $\gamma(x)$ to enable better fitting for high frequency variations in the data [49, 63]:

$$\gamma(x) = [\cos(2^l x), \sin(2^l x)]_{l=0}^{L-1} \quad (2)$$

NeRF 3D Generation. NeRFs inspired follow-up works to synthesize new NeRF objects from scratch. The first methods used NeRF combined with GANs [1, 19, 21] to design 3D-aware generators [20, 10, 14, 43, 44, 54, 72]. GRAF [54] adopts shape and appearance codes to conditionally synthesize NeRF and GIRAFF [44], StyleNeRF [20] utilizes NeRF to render features instead of RGB colors and adopt a two-stage strategy, where they render low-resolution feature maps first and then up-sample the feature maps using a CNN decoder. These models are category-specific and trained mostly on forward-facing scenes.

More recent works utilize the progress in contrastive representation learning [13, 48, 69, 31, 30], which enables easy and flexible control over the content of the generated objects using textual input. In Dream Fields [26], frozen image-text joint embedding models from CLIP [48] are used as a guidance to a NeRF model that generates 3D objects whose renderings have high semantic similarity with the input caption. To improve the visual quality, they introduce geometric priors and augmentations to enforce transmittance sparsity, object boundaries and multi-view consistency. In this paper, we utilize some of the priors from Dream Fields [26] and introduce improved augmentations and priors to edit existing NeRF scenes.

More recent works utilize the progress in diffusion models [24, 58, 59] and specifically in text-conditioned diffusion models [51, 52, 53]. DreamFusion [47] and its follow-ups [65, 38, 32, 50] optimize a NeRF model by replacing CLIP with score function losses using pretrained text-conditioned 2D diffusion-models applied on many different views of the generated scene to synthesize 3D objects aligned with the input text. These models synthesize new objects without considering how they can be inserted and blend into an existing scene.

Editing NeRFs. The pioneering works [34, 64] were the first to tackle the challenge of editing NeRF scenes. They both define a conditional NeRF, where the NeRF model is

conditioned on latent shape and appearance codes, which enables separately editing the shape and the color of a 3D object. EditNeRF [34] only enables addition and removal of local parts or color changes guided by user scribbles and is limited to only one shape category. In ObjectNeRF [67] they enable editing tasks such as moving or adding new objects to the scene by introducing a neural scene rendering system with a scene branch which encodes the scene geometry and appearance and object branch which encodes each standalone object. CLIP-NeRF [64] leverage the joint language-image embedding space of CLIP [48] to perform text or image guided manipulation on the entire scene. During the optimization it uses two code mappers for the shape and appearance that receive the CLIP embedding and output shape and appearance codes which steer the input of the model and the model weights to apply the edit. The manipulation capabilities are demonstrated mainly on synthetic objects from one class and on global color changes for realistic scenes.

Later works focused on geometric edits [68], global style transfer [11, 12, 16, 25], recoloring [66, 18], and disentanglement of the scene to enable local edits [28, 7, 71]. Kobayashi [28] decomposes the scene to its semantic parts by training the NeRF model to learn a 3D feature field using supervision of pre-trained 2D image feature extractors [9, 29] in addition to learning of the volume density and the radiance field. After training, the model can perform zero-shot segmentation for local editing of a specific semantic class. Benaim *et al.* [7] disentangle the volumetric representation of a foreground object from its background using a set of 2D masks specifying the foreground object in each training view. They train two models for the full scene and the background scene, and subtract the background from the full scene in order to get the foreground. In both works the localization on the region of interest is incomplete and not flexible enough (does not enable editing parts of objects, empty regions or blending new densities into the area of existing object). They demonstrate manipulations such as object removal, transformations such as shift rotation and scale, and only basic optimization-based edits. Our work focuses on blending text generated objects with volume and color into any region of interest of an existing scene with more freedom and flexibility and without compromising on quality and visibility. For information regarding concurrent works, please refer to the supplement.

3. Method

Given an existing 3D scene x_o represented by a NeRF model F_θ^O , and a 3D region of interest (ROI), indicated by a box B localized inside the scene, our goal is to modify the scene inside the ROI, according to a user-provided text prompt. In other words, we aim to obtain a modified scene x_e , where $x_e \odot B$ is consistent with the user prompt from




Figure 3: **Distance Smoothing Operator.** We demonstrate our suggested smoothing operator in eq. (5) on a range of α values, When α is zero all the weight goes to the edited scene, and as we increase α , more attention is given to closer points from the original scene.

any point of view, while matching x_o outside the box ($x_e \odot (1 - B) = x_o \odot (1 - B)$).

To perform the edits inside the ROI we initialize a 3D MLP model F_θ^G with the weights of the original scene model F_θ^O and steer the weights towards the given prompt using a pretrained language-image model, such as CLIP [48]. We enable local edits in any region of the scene x_o using a simple GUI for localizing a 3D box inside the scene by rendering the original NeRF model F_θ^O from any view and using the output depth map of the model to obtain 3D understanding of the scene. Using the given ROI box we can disentangle the scene inside the box and outside it by decomposing the radiance fields accordingly. To obtain consistent results from any view direction, we perform volumetric blending of the original and the edited radiance fields by sampling 3D points along each camera ray \mathbf{r} in both F_θ^O and F_θ^G , and blending the samples while accounting for their densities, colors and distance from the center of the scene.

To get more realistic and natural-looking results we present existing [26] and novel augmentations and priors such as transmittance and depth regularization, background augmentations, pose sampling and directional dependent prompts. An overview of our approach is depicted in Figure 1.

In Section 3.1 we describe our 3D object generation and blending process, we continue and present the model objectives and proposed priors in Section 3.2.

3.1. Image-Text driven 3D synthesis and blending

Given a 3D scene represented by a NeRF model F_θ^O , an ROI box B , and a camera pose, we use a duplicate of F_θ^O , F_θ^G as our starting point for generating the content of B . The rest of the scene is preserved by rendering only the rays which have sample points inside B . The training of F_θ^G is guided by a language-image model, e.g., [48, 31, 30, 69] to align the content generated inside B with a user-provided text prompt.

To get a smoothly blended result, we query both models

F_θ^O, F_θ^G using the same set of rays. For sample points outside the ROI, we use the density and color inferred by F_θ^O , while for points inside the ROI, we blend the results of the two radiance fields using one of two modes, depending on the type of the edit: adding a new object in empty space, or completely replacing an existing one, vs. adding an object in a non-empty area.

F_θ^G is optimized using guidance from a language-image model, such as CLIP [48], by aiming to minimize the cosine similarity score between the user-provided text prompt y and rendered views of the generated content inside the ROI box, I_{ROI} :

$$L_{sim} = -E_{img}(I_{ROI})^T E_{txt}(y), \quad (3)$$

where E_{img}, E_{txt} are the image and text encoders of the image-language model. During optimization, we render I_{ROI} using only the 3D sample points contained inside B by sampling only along rays \mathbf{r} that pass through the box and setting the density to zero for all sample points outside B , according to eq. (1):

$$C(\mathbf{r}) = \begin{cases} \sum_{x_i \in B} T_i (1 - e^{-\sigma_i \delta_i}) c_i, & \exists x_i \in \mathbf{r} \text{ s.t. } x_i \in B \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

After training, we blend the scenes inside and outside the ROI with the same set of rays by querying both F_θ^O and F_θ^G where the points inside the box are rendered by F_θ^G and the points outside the box are rendered by F_θ^O . To get smooth blending between the two scenes we suggest distance smoothing operator per point inside the box considering its distance from the center of the ROI scene (center of mass, computed during training) and alpha compositing the density and color of the two scenes inside the ROI as follows:

$$f(\mathbf{x}) = 1 - \exp\left(\frac{-\alpha d(\mathbf{x})}{diag}\right) \quad (5)$$

$$\sigma_{blend}(\mathbf{x}) = f(\mathbf{x}) \cdot \sigma_O(\mathbf{x}) + (1 - f(\mathbf{x})) \cdot \sigma_G(\mathbf{x})$$

$$c_{blend}(\mathbf{x}) = f(\mathbf{x}) \cdot c_O(\mathbf{x}) + (1 - f(\mathbf{x})) \cdot c_G(\mathbf{x})$$

where σ_O and σ_G are the densities returned by each model, $d(\mathbf{x})$ is the Euclidean distance of a point \mathbf{x} inside the ROI from the center of the scene, $diag$ is the box diagonal and α is a hyperparameter which controls the strength of the blending, as can be seen intuitively in Figure 3. The resulted raw densities and RGB values inside and outside the ROI are then blended along each ray using eq. (1) to get the current rendered view of the edited scene x_e .

Object Insertion/Replacement. In this mode, a new synthetic object is added into an empty region of the scene, or entirely replaces another existing object inside the ROI. In this mode, we use the pipeline described above, when inside the ROI we consider only the radiance field of F_θ^G during training. After training, we blend the two scenes as described above.

Object Blending. In contrast to the above mode, here we aim to blend the new content with the existing scene inside the ROI. We query both the original F_θ^O and the edited F_θ^G fields inside the box and blend the resulting colors and densities at each ray sample. To blend the sample colors, we first compute the alpha values for each point x_i on the ray separately from each model:

$$\begin{aligned}\alpha_O(x_i) &= 1 - \exp(\phi(\sigma_O(x_i)) \cdot \delta_i) \\ \alpha_G(x_i) &= 1 - \exp(\phi(\sigma_G(x_i)) \cdot \delta_i)\end{aligned}\quad (6)$$

where ϕ is the activation function enforcing that these density values are non-negative. To blend the colors c_O and c_G obtained from the two models, we use the above alpha values, followed by a sigmoid function:

$$c(x_i) = S\left(\frac{c_O(x_i) \cdot \alpha_O(x_i) + c_G(x_i) \cdot \alpha_G(x_i)}{\epsilon + \alpha_O(x_i) + \alpha_G(x_i)}\right)\quad (7)$$

where ϵ is a small constant, for numerical stability and S is the sigmoid function.

For the density of the blended sample, we consider two options, which have different impact on the results of the blending:

$$\sigma(x_i) = \phi(\sigma_O(x_i) + \sigma_G(x_i))\quad (8)$$

$$\sigma(x_i) = \phi(\sigma_O(x_i)) + \phi(\sigma_G(x_i))\quad (9)$$

i.e., summing the densities inside or outside the activation function. When using eq. (8) we are summing inside the activation function thus allowing the generator F_θ^G to change the original scene density and even remove densities (if $\sigma_G(x_i) < 0$), while in eq. (9) we allow F_θ^G to only add new densities to the scene. We can choose either of these two options depending on the edit we wish to apply. We then compute the joint transmittance and alpha values according to eq. (1). The resulting blended image I_{ROI} is then used to guide F_θ^G during training by measuring its similarity to the input caption using eq. (3). The blending process after training is the same as in Object Insertion/Replacement mode. An illustration of our blending modes on the blender Lego scene is presented in Figure 4.

3.2. Objectives and Priors

Previous works [26, 7, 64] and our experiments indicate that a scene representation depending on similarity loss alone (eq. (3)) is too unconstrained, resulting in a scene that is not visually compatible to a human, but still satisfies the loss. Thus, we utilize the priors and augmentations mentioned in DreamFields [26] and suggest additional priors to get more realistic results.

Pose Sampling. CLIP-NeRF [64] shows the multi-view consistency evaluation of CLIP [48]. When using different camera poses and rendering different views of the same object, they still have high similarity, in contrast to different objects which have low similarity even in identical view. DreamFields [26] shows that sampling different camera poses is a good regularizer and improves the realism of the object geometry. Thus, each iteration we sample a random camera pose around the scene depending on the scene type (360° and forward-facing scenes) including its azimuth and elevation angles (θ, ϕ). We found it beneficial to be relatively close to the object during training to get a bigger object in the rendered view, which in turn yields larger gradients from eq. (3). We set the initial distance d from the ROI according to the camera $AFOV = 2\gamma$ and the maximum dimension of the box e_{max} and we randomly sample the radius r around this value:

$$d = \frac{e_{max}}{2 \tan(\gamma/2)}\quad (10)$$

Background Augmentation. DreamFields [26] note that when using white or black background during opti-




Figure 4: **Blending Modes.** Guided by “plant with green leaves and white and blue flowers”. When using eq. (8) (second column), we allow F_θ^G to change the density of the original scene, in this case removing parts of the wheel. When utilizing eq. (9) (third column), we can only add additionally density to the scene, so the plant warps around the wheel without changing it.

mization, the scene populates the background, and eventually we get a diffused scene. Thus, we use the same random backgrounds as in DreamFields: Gaussian noise, checkerboard patterns and random Fourier textures from [41] to get more sharp and coherent objects.

Directional Dependent Prompts. Due to the fact that there’s no constraint on F_{θ}^G to describe the object differently in different views, we concatenate to the original caption a text prompt depending on the current view. For more details, please refer to the supplementary materials.

Transmittance loss. Same as in DreamFields [26], in order to get more sparse and coherent results we encourage the generator to increase the average transmittance of the scene inside the box by adding a transmittance loss to the generator objective:

$$L_T = -\min(\tau, \text{mean}(T(\mathbf{P}))) \quad (11)$$

Where $\text{mean}(T(\mathbf{P}))$ is the average transmittance of a rendered view from pose \mathbf{P} and τ is the max transmittance.

Depth loss. When blending in forward-facing scenes (such as LLFF dataset [40]) and due to the limited viewing intervals, for some captions we get a flat billboard geometry effect and the resulting edit does not seem to have a volume. We encourage the generator to synthesize volumetric 3D shapes by adding a depth loss to the generator objective:

$$L_D = -\min(\rho, \sigma^2(D(\mathbf{P}))) \quad (12)$$

Where $\sigma^2(D(\mathbf{P}))$ is the variance of the disparity map of a rendered view from pose \mathbf{P} and ρ is the max variance we allow during training. We gradually introduce L_T and L_D during training using annealing strategy to prevent completely transparent or amorphous scenes. In summary, the final objective for the generator F_{θ}^G is:

$$L_{total} = L_{sim} + \lambda_T L_T + \lambda_D L_D \quad (13)$$

Where λ_T, λ_D are the weights for L_T, L_D accordingly. For more information on implementation details and hyperparameters, please refer to the supplement.

4. Experiments

In Section 4.1 we begin by comparing our method both qualitatively and quantitatively to the baseline Volumetric Disentanglement for 3D Scene Manipulation [7]. Next, in Section 4.2 we demonstrate the effect of our suggested priors and augmentations on improving fidelity and visual quality. Finally, in Section 4.3 we demonstrate several applications enabled by our framework.

4.1. Comparisons

Our qualitative comparisons to Volumetric Disentanglement [7] are shown in Figure 5. Since the implementation




Figure 5: **Comparison to [7] for object replacement.** We compare our editing capabilities to [7] in the fern scene from llff dataset [40]. The left and right images in each row are [7] and ours, accordingly. Our proposed method exhibits more realistic results that agrees better with the text. For example the edit for the text “aspen tree” indeed looks like a trunk of an aspen tree in our edit.

Method	CLIP	CLIP	LPIPS↓
	Direction Similarity↑	Direction Consistency↑	
[Benaim 2022]	0.128	0.736	0.3
Ours	0.143	0.787	0.024

Table 1: **Quantitative Evaluation.** Quantitative comparison to [7] using the metrics described in Section 4.1. Our method demonstrates edits that are better align to the input captions and consistent between views, while preserving the background of the scene.

of [7] is not currently available, we preform the comparisons using the examples included on their project page ¹. As can be seen from the results in Figure 5, our results exhibit richer and more natural colors and are aligned better with the text. To test these observations quantitatively, in Table 1 we compare our proposed method to [7] using three metrics:

(1) **CLIP Direction Similarity**, a metric originally introduced in StyleGAN-NADA [17], measures how well the change between the original and edited views is aligned with the change in the texts describing them (in the CLIP

¹<https://sagiebenaim.github.io/volumetric-disentanglement/>



Figure 6: **Depth Loss Impact.** Comparison of synthesizing a “donut covered with glaze and sprinkles” from COCO dataset [33] on a limited view scene with and without our suggested depth prior. The first column displays a view of the edited scenes and the second column displays the disparity map of the synthesized objects. In (a) the results are more flat, which can be clearly seen in the disparity map.

embedding space).

(2) **CLIP Direction Consistency**, introduced by Haque [22], measures the cosine similarity of the CLIP embeddings of a pair of adjacent frames. For each edit, we take 6 consecutive frames, compute the metric for each consecutive pair, and average the results among all pairs.

Finally, we use (3) **LPIS** [70] to measure the difference between the original and edited scenes, with the ROI masked, for comparing the background preservation. As can be seen from Table 1, our model outperforms the baseline in all metrics, which implies that our generated objects match better to the input text captions, they are more consistent from any view and, on the other hand, our method manages to keep the rest of the scene untouched.

4.2. Ablation Study

To show the importance of our proposed augmentations and priors, we use the R-Precision score [45] using both CLIP and BLIP [48, 31, 30] as the metric language-image model to measure how well the generated images align with the true caption. Similar to DreamFields [26], we use a randomly selected subset of 20 samples (due to time and resources limitations) from the object-centric dataset which contains 153 images and captions from COCO dataset [33] as our ground truth. The objects are synthesized using the given captions and blended into an empty region in the lfff

Method	CLIP	BLIP
	R-Precision \uparrow	R-Precision \uparrow
COCO GT	0.933	0.98
Ours(full pipeline)	0.86	0.8
Ours(no dir prompts)	0.85	0.8
Ours(no depth prior)	0.81	0.78

Table 2: **Ablation study.** We test our proposed priors and augmentations on a subset of captions and images from COCO dataset [33]. The CLIP and BLIP R-Precision scores utilize CLIP B-32 and BLIP2 architecture accordingly. The first row shows the scores of the GT COCO image, the second row shows our method scores using all the priors and augmentations as described in Section 3 and the last two rows present the scores when taking out the directional dependent prompts and the depth loss.

fern scene. Due to the fact we are training on the same CLIP model, we test our results with a different language-image model, BLIP2 [30]. The results of both metrics are presented in Table 2. The directional dependent prompts seem to only slightly improve the results, probably due to the forward-facing nature of the scene. When rendering from limited camera positions and viewing angles and without our proposed depth priors, the results deteriorate. To test this conclusion visually, in Figure 6 we compare the task of inserting a new object into an empty region of the fern lfff scene [40] with and without the depth loss. As can be seen from the figure, when using our proposed depth prior, the generated object has more volume and looks more natural and consistent. For additional details, please refer to the supplement.

4.3. Applications

In this section, we demonstrate the applicability of our framework for several 3D editing scenarios.

New Object Insertion. Using the method described in Section 3, and by placing the ROI box in an empty space of the scene, we can synthesize a new object given a text prompt and blend it into the original scene. Visual example of this application can be seen in Figure 6 and in the supplement.

Object Replacement. To replace an existing object in the scene with new synthesized content, we place the ROI 3D box in the required area (enclosing the object to be replaced), and perform the training process described in Section 3. In Figure 2 we demonstrate the replacement of the sea in the blender ship scene, while in Figure 5 we replace the fern’s trunk.

Blending of Objects. To preform blending between the original and the generated object inside the ROI, we utilize the object blending process described in Section 3. In Figure 4 and Figure 8 we demonstrate this blending on blender



Figure 7: **Texture Editing.** We can change only the texture of an object by freezing the layers responsible for the density and training only the layers that impact the color of the scene. To get a smooth blending, we utilize eq. (5) to blend the scene inside and outside the ROI.

lego and lfff fern scenes.

Texture Editing. We enable texture editing by training only the color-related layers of F_{θ}^G and freezing all the other layers in a similar way as in [64]. For seamless blending results, we utilize eq. (5). In Figure 7 we demonstrate this edit method on 360 scenes. For additional results and videos please refer to supplement.



Figure 8: **Blending Densities Inside Activation.** We demonstrate our suggested blending procedure for blending the original and synthesized objects inside the ROI in lfff fern scene [40] using eq. (8) for summing the densities.

5. Limitations and Conclusions

We introduced a novel solution to blend new objects into an existing NeRF scene with natural looking and consistent results by utilizing a language-image model to steer the generation process towards the edit and by introducing novel priors, augmentations and volumetric blending techniques for improving the final edited scene. We tested our method on a variety of scenes and text prompts and showed the applicability of our framework on several editing applications. We believe that our framework can be utilized in a variety of applications due to the ease and intuitive interaction enabled by our interface.

One of the limitations of our framework is that currently it can’t edit multiple objects in a given scene, such as changing two wheels of a 3D car without impacting the rest of the scene. Additionally, the use of a box as our ROI scene shape can be sometimes limiting; for example, when trying to edit a circular scene like the blender ship scene in Figure 2, a cylinder could be preferable. Due to the fact we are rendering one view in each training step, we may get artifacts like multiple heads on the generated object. The quality of our generated objects can be improved by utilizing the recent progress in diffusion models, we leave it as a future work to combine our suggested blending framework with these models.

Acknowledgements: This work was supported in part by the Israel Science Foundation (grants No. 2492/20, and 3611/21).

References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, 2017. 3
- [2] Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *ArXiv*, abs/2206.02779, 2022. 2
- [3] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218, 2022. 2
- [4] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 1, 2
- [5] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 1, 2
- [6] David Bau, Alex Andonian, Audrey Cui, YeonHwan Park, Ali Jahani, Aude Oliva, and Antonio Torralba. Paint by word. *arXiv preprint arXiv:2103.10951*, 2021. 2
- [7] Sagie Benaim, Frederik Warburg, Peter Ebert Christensen, and Serge Belongie. Volumetric disentanglement for 3d scene manipulation. *ArXiv*, abs/2206.02776, 2022. 2, 3, 5, 6
- [8] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. *arXiv preprint arXiv:2211.09800*, 2022. 2
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 3
- [10] Eric R. Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. Pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5799–5809, 2021. 3
- [11] Yaosen Chen, Qi Yuan, Zhiqiang Li, Yuegen Liu, Wei Wang, Chaoping Xie, Xuming Wen, and Qien Yu. Upst-nerf: Universal photorealistic style transfer of neural radiance fields for 3d scene. *arXiv preprint arXiv:2208.07059*, 2022. 3
- [12] Pei-Ze Chiang, Meng-Shiun Tsai, Hung-Yu Tseng, Wei-Sheng Lai, and Wei-Chen Chiu. Stylizing 3d scene via implicit representation and hypernetwork. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1475–1484, 2022. 3
- [13] Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11162–11173, 2021. 3
- [14] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W. Taylor, and Joshua M. Susskind. Unconstrained scene generation with locally conditioned radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14304–14313, 2021. 3
- [15] Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021. 1, 2
- [16] Zhiwen Fan, Yifan Jiang, Peihao Wang, Xinyu Gong, Dejia Xu, and Zhangyang Wang. Unified implicit neural stylization. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XV*, pages 636–654, 2022. 3
- [17] Rinon Gal, Or Patashnik, Haggai Maron, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *arXiv preprint arXiv:2108.00946*, 2021. 6
- [18] Bingchen Gong, Yuehao Wang, Xiaoguang Han, and Qi Dou. Recolornerf: Layer decomposed radiance fields for efficient color editing of 3d scenes. *arXiv preprint arXiv:2301.07958*, 2023. 3
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2014. 3
- [20] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleneRF: A style-based 3d aware generator for high-resolution image synthesis. In *Advances in Neural Information Processing Systems*, 2022. 3
- [21] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, volume 30, 2017. 3
- [22] Ayaan Haque, Matthew Tancik, Alexei A Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. *arXiv preprint arXiv:2303.12789*, 2023. 7
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851, 2020. 3
- [25] Yi-Hua Huang, Yue He, Yu-Jie Yuan, Yu-Kun Lai, and Lin Gao. Stylizednerf: consistent 3d scene stylization as stylized nerf via 2d-3d mutual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18342–18352, 2022. 3
- [26] Ajay Jain, Ben Mildenhall, Jonathan T. Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3, 4, 5, 6, 7
- [27] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3d scenes. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020. 1, 2
- [28] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. In *NeurIPS*, 2022. 2, 3
- [29] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022. 3
- [30] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, 2023. 3, 4, 7
- [31] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 3, 4, 7
- [32] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *Proceedings ECCV 2014, Part V 13*, pages 740–755. Springer, 2014. 7
- [34] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 2, 3
- [35] Xian Liu, Yinghao Xu, Qianyi Wu, Hang Zhou, Wayne Wu, and Bolei Zhou. Semantic-aware implicit neural audio-driven video portrait generation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVII*, pages 106–125. Springer, 2022. 1
- [36] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. 3
- [37] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019. 1, 2
- [38] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. *arXiv preprint arXiv:2211.07600*, 2022. 3
- [39] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2
- [40] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019. 6, 7, 8
- [41] Alexander Mordvintsev, Nicola Pezzotti, Ludwig Schubert, and Chris Olah. Differentiable image parameterizations. *Distill*, 3(7):e12, 2018. 6
- [42] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning*, 2021. 2
- [43] Michael Niemeyer and Andreas Geiger. Campari: Camera-aware decomposed generative neural radiance fields. In *Proc. of the International Conf. on 3D Vision (3DV)*, 2021. 3
- [44] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11453–11464, 2021. 3
- [45] Dong Huk Park, Samaneh Azadi, Xihui Liu, Trevor Darrell, and Anna Rohrbach. Benchmark for compositional text-to-image synthesis. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. 7
- [46] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. 1, 2
- [47] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*, 2023. 3
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, 2021. 1, 2, 3, 4, 5, 7
- [49] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred A. Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, 2019. 3
- [50] 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 *arXiv preprint 2303.13508*, 2023. 3
- [51] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv*, abs/2204.06125, 2022. 2, 3
- [52] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 3
- [53] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed

- Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, 2022. 3
- [54] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In *Advances in Neural Information Processing Systems*, volume 33, pages 20154–20166, 2020. 3
- [55] Vincent Sitzmann, Eric R. Chan, Richard Tucker, Noah Snavely, and Gordon Wetzstein. Metasdf: Meta-learning signed distance functions. In *arXiv*, 2020. 1, 2
- [56] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020. 1, 2
- [57] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, 2019. 1, 2
- [58] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 3
- [59] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *Advances in Neural Information Processing Systems*, volume 33, pages 12438–12448, 2020. 3
- [60] Yannick Strümpfer, Janis Postels, Ren Yang, Luc Van Gool, and Federico Tombari. Implicit neural representations for image compression. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*, pages 74–91. Springer, 2022. 1, 2
- [61] Kun Su, Mingfei Chen, and Eli Shlizerman. Inras: Implicit neural representation for audio scenes. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 8144–8158. Curran Associates, Inc., 2022. 1
- [62] Filip Szatkowski, Karol J Piczak, Przemysław Spurek, Jacek Tabor, and Tomasz Trzciniński. Hypersound: Generating implicit neural representations of audio signals with hypernetworks. *arXiv preprint arXiv:2211.01839*, 2022. 1
- [63] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, volume 33, 2020. 3
- [64] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18–24, 2022*, 2022. 2, 3, 5, 8
- [65] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. *CoRR*, abs/2212.00774, 2022. 3
- [66] Qiling Wu, Jianchao Tan, and Kun Xu. Palettenerf: Palette-based color editing for nerfs. *arXiv preprint arXiv:2212.12871*, 2022. 3
- [67] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *International Conference on Computer Vision (ICCV)*, October 2021. 3
- [68] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18353–18364, 2022. 3
- [69] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18123–18133, 2022. 3, 4
- [70] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 7
- [71] Xiaoshuai Zhang, Abhijit Kundu, Thomas Funkhouser, Leonidas Guibas, Hao Su, and Kyle Genova. Nerflets: Local radiance fields for efficient structure-aware 3d scene representation from 2d supervision. *CVPR*, 2023. 3
- [72] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. Cips-3d: A 3d-aware generator of gans based on conditionally-independent pixel synthesis, 2021. 3