

Customize your NeRF: Adaptive Source Driven 3D Scene Editing via Local-Global Iterative Training

Runze He^{1,2*} Shaofei Huang^{1,2*} Xuecheng Nie³ Tianrui Hui⁴
Luoqi Liu³ Jiao Dai^{1,2} Jizhong Han^{1,2} Guanbin Li^{5†} Si Liu^{6,7†}

¹Institute of Information Engineering, Chinese Academy of Sciences

²School of Cyber Security, University of Chinese Academy of Sciences ³MT Lab, Meitu Inc.

⁴School of Computer Science and Information Engineering, Hefei University of Technology

⁵Sun Yat-sen University ⁶Institute of Artificial Intelligence, Beihang University

⁷Hangzhou Innovation Institute, Beihang University

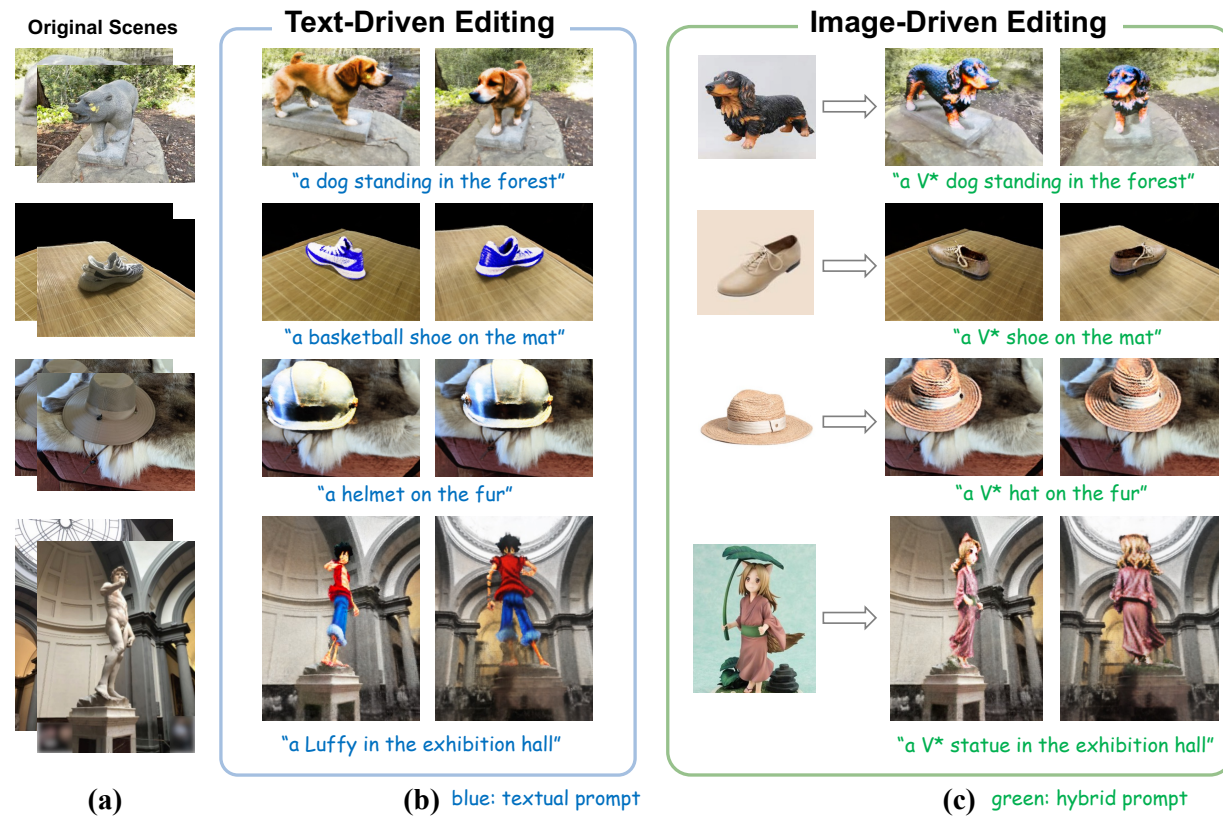


Figure 1. Adaptive source driven editing by CustomNeRF. Our method can accurately edit the original NeRF scene according to text descriptions or reference images, which are formulated as textual prompts or hybrid prompts. After editing, the foreground regions exhibit appropriate geometric and texture modifications consistent with the editing prompts. The layouts and background contents of the edited scenes remain almost the same as the original ones.

Abstract

In this paper, we target the adaptive source driven 3D scene editing task by proposing a CustomNeRF model that

unifies a text description or a reference image as the editing prompt. However, obtaining desired editing results conformed with the editing prompt is nontrivial since there exist two significant challenges, including accurate editing of only foreground regions and multi-view consistency given a single-view reference image. To tackle the first challenge,

*Equal contribution ({hrz010109, nowherespyfly}@gmail.com).

†Corresponding authors.

we propose a Local-Global Iterative Editing (LGIE) training scheme that alternates between foreground region editing and full-image editing, aimed at foreground-only manipulation while preserving the background. For the second challenge, we also design a class-guided regularization that exploits class priors within the generation model to alleviate the inconsistency problem among different views in image-driven editing. Extensive experiments show that our CustomNeRF produces precise editing results under various real scenes for both text- and image-driven settings. The code is available at: <https://github.com/hrz2000/CustomNeRF>.

1. Introduction

Characterized by the ease of optimization and continuous representation, neural fields (e.g., NeRF [23] and NeuS [43]) have garnered extensive attention, further inspiring numerous studies in the field of 3D scene editing, such as retexturing [34, 45], stylization [50], or deformation [7, 47] of 3D objects or scenes. In order to enhance the flexibility and accessibility of editing, recent works [8, 39, 51, 52] have proposed text-driven 3D scene editing, utilizing textual prompts to enable editing with more freedom. These methods take advantage of the powerful capabilities of pre-trained diffusion models [4, 35] to provide visual information conformed with the textual prompts, thereby achieving semantic-aware editing.

However, these text-driven methods are only able to achieve **general editing** to the original 3D scene, leading to a wide range of possible editing outcomes. To meet more customized requirements of users, a reference image could be introduced to depict the desired outcome precisely, inspiring a more challenging but under-explored task of **specific editing**. In Figure 1, the textual prompt “a dog standing in the forest” is an example of general editing, which replaces the bear statue with an arbitrary dog. By further restricting the visual characteristics of the dog with a reference image such as a black dog, one can realize specific editing to replace the bear statue with the desired one.

In this paper, we present a unified framework named CustomNeRF to support adaptive source driven 3D scene editing, where a subject-aware Text-to-Image (T2I) model [16] is employed to embed the specific visual subject V^* contained in the reference image into a hybrid prompt, meeting both general and specific editing requirements as illustrated in Figure 1. The key to obtaining desired editing results lies in the precise identification of the foreground regions, thereby facilitating view-consistent foreground editing while maintaining the background. This poses two significant challenges: Firstly, since NeRF is an implicit representation of 3D scenes, editing NeRF directly is difficult in terms of accurately locating and manipulating only the fore-

ground regions, which necessitates a well-designed training scheme. Additionally, the lack of accurate inter-view calibration data during training may lead to inconsistency among views, i.e., the Janus problem. This problem worsens in the image-driven setting, as the diffusion model fine-tuned tends to overfit the input viewpoint of the reference images as encountered in [30].

To tackle the first challenge, we propose a Local-Global Iterative Editing (LGIE) training scheme that alternates between foreground region editing and full-image editing. Concretely, a foreground-aware NeRF is designed to estimate the editing probability in addition to color and density to distinguish between foreground and background regions, enabling the separate rendering of the foreground region. In the local stage, we feed the foreground-only rendered images and corresponding editing prompt into the T2I model, so that the parameters of NeRF are optimized to only edit the foreground region. In the global stage, the T2I model takes the fully rendered images and editing prompt as input to optimize NeRF, leveraging useful clues about sizes, positions, and postures contained in the background region to prevent layout drifting from the original scene. Through the local-global alternating, we can simultaneously constrain the foreground layout and maintain the background content. To tackle the second challenge, we design a class-guided regularization that only uses class words to denote the subject of the reference image to form a textual prompt during the local editing stage. In this way, the general class priors within the T2I model can be exploited to promote geometrically consistent editing of the novel subject. Together with the specific subject appearance learned by the hybrid prompt in the global editing stage, we can gradually transfer the subject’s visual characteristics to different views for producing view-consistent editing results.

Our contributions are summarized as follows: (1) We present a unified framework CustomNeRF to support adaptive source driven 3D scene editing, achieving accurate editing in both text-driven and image-driven settings. (2) A Local-Global Iterative Editing (LGIE) training scheme is proposed to accurately edit foreground regions while maintaining background content. (3) We design a class-guided regularization to exploit class priors within the T2I model for alleviating the Janus problem in image-driven editing.

2. Related Works

2.1. Text-to-Image Generation and Editing

Text-to-image generation [3, 5, 13, 31–33, 35, 37], which aims to generate high-quality images according to given text prompts, has gained considerable attention recently. Among these methods, diffusion model based methods, such as Stable Diffusion [35] and DALL-E 2 [32], have shown remarkable image generation ability as a result of

training on massive text-image datasets. Utilizing the rich cross-modal semantic knowledge inherent in pre-trained text-to-image generation models, several works [1, 4, 9, 20] have employed them in text-guided image editing, which aims to modify given images based on text descriptions. However, these methods are limited to editing styles or contents of images but fail to perform well when faced with more complex modifications, *e.g.*, layout adjustments, subject posture transformations, etc.

To preserve subject identity between the generated and original images while achieving more intricate control over the appearance and posture of these subjects, some works [6, 16, 36] have explored the task of subject-aware text-to-image generation. For example, Custom Diffusion [16] realizes faster tuning speed with efficient parameter optimization, and further explores multiple subject combinations through closed-form constrained optimization, thereby extending the capabilities of subject-aware text-to-image generation. In our work, we employ this method for reference image embedding, thus unifying text-driven and image-driven editing settings.

2.2. 3D Generation and Editing

With the significant progress made in text-to-image generation, there has been an increasing exploration of text-to-3D generation where pretrained visual-language models [38, 41] (*e.g.*, CLIP [29]) or text-conditioned diffusion models [12, 17, 28, 30, 42, 46] (*e.g.*, Stable Diffusion [35]) are leveraged to generate realistic 3D objects or scenes from textual descriptions. DreamFusion [28] proposes a Score Distillation Sampling (SDS) loss to use a pre-trained 2D diffusion model as a prior for optimization of neural fields, but cannot be used to edit existing real-world scenes due to lack of full alignment between images and text descriptions.

Given a pre-trained NeRF of a photo-realistic scene, neural field editing allows for manipulation at varying degrees, resulting in an edited NeRF that is visually coherent and cross-view continuous [2, 8, 18, 25, 34, 39, 41, 48, 52]. Leveraging the powerful text comprehension and generation capabilities of text-guided image generation and editing models, some works [8, 39, 51, 52] have enabled complex textual instructions to freely manipulate 3D scenes at local or global levels. In order to achieve accurate 3D editing, it is critical to determine precise editing regions while avoiding unnecessary modifications to non-editing regions in a 3D scene, which is relatively challenging for the implicit NeRF representation. To this end, Vox-E [39] and DreamEditor [52] adopt explicit 3D representations (*e.g.*, voxel and mesh) and utilize the cross-attention maps to selectively modify the target regions, RePaint-NeRF [51] constrains background contents with a semantic mask. Compared with these methods, our method edits directly on NeRF and conducts foreground-only edits with a well-designed Local

Global Iterative Editing training scheme. Furthermore, our method supports both text-driven and image-driven editing, which allows users to freely customize existing NeRF models according to their requirements.

3. Preliminary

3.1. Latent Diffusion Model

Latent Diffusion Models (LDMs) [35] is a variation of Denoising Diffusion Probabilistic Models (DDPMs) [10] that operate in the latent space of images. It consists of an auto-encoder and a UNet-structured diffusion model. The auto-encoder is composed of an encoder $\mathcal{E}(\cdot)$ that maps an image \mathbf{x} into a spatial latent code $\mathbf{z} = \mathcal{E}(\mathbf{x})$ and a decoder $\mathcal{D}(\cdot)$ that maps the latent code back to image $\mathcal{D}(\mathcal{E}(\mathbf{x}))$. The diffusion model is trained to produce latent codes within the encoder’s latent space with:

$$\mathcal{L}_{LDM} = \mathbb{E}_{\mathbf{z} \sim \mathcal{E}(\mathbf{x}), y, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon_{\phi}(\mathbf{z}_t; t, y) - \epsilon\|_2^2], \quad (1)$$

where \mathbf{z}_t is the noisy latent code at timestep t , ϵ is the sampled noise, $\epsilon_{\phi}(\cdot)$ is the UNet that predicts the noise content, and y denotes the text description. During inference, a random noise \mathbf{z}_T is sampled and iteratively denoised to produce \mathbf{z}_0 , which is mapped back to the image space through image decoder $\mathbf{x}' = \mathcal{D}(\mathbf{z}_0)$.

3.2. Score Distillation Sampling Loss

Score Distillation Sampling (SDS) loss [28] optimizes a NeRF model under the guidance of a 2D text-to-image diffusion model mentioned in Section 3.1 to realize text-to-3D generation. Let θ denote the parameters of NeRF model, and $g(\cdot)$ denotes the process of image rendering. The NeRF model first renders a randomly viewed image $\mathbf{x} = g(\theta)$, which is then transformed to noisy latent code \mathbf{z}_t by feeding to the latent encoder $\mathcal{E}(\cdot)$ and adding noise sequentially. The SDS loss is calculated as follows:

$$\nabla_{\theta} \mathcal{L}_{SDS}(\phi, g(\theta)) = \mathbb{E}_{t, \epsilon} [\omega(t) (\epsilon_{\phi}(\mathbf{z}_t; t, y) - \epsilon) \frac{\partial \mathbf{x}}{\partial \theta}], \quad (2)$$

where $\omega(t)$ is a weighting function depending on timestep t (Equation 1). During training, the above gradients are back-propagated to the NeRF model only, enforcing it to render images that look like outputs of the frozen diffusion model.

4. Method

Given a set of well-calibrated images representing a 3D scene, our goal is to edit the reconstructed 3D scene under the guidance of a textual prompt or a reference image, enabling editing of the foreground regions while preserving the background regions simultaneously. The overall pipeline of our CustomNeRF is illustrated in Figure 2, which consists of three steps. First, to reconstruct the original 3D scene, we train a foreground-aware NeRF where

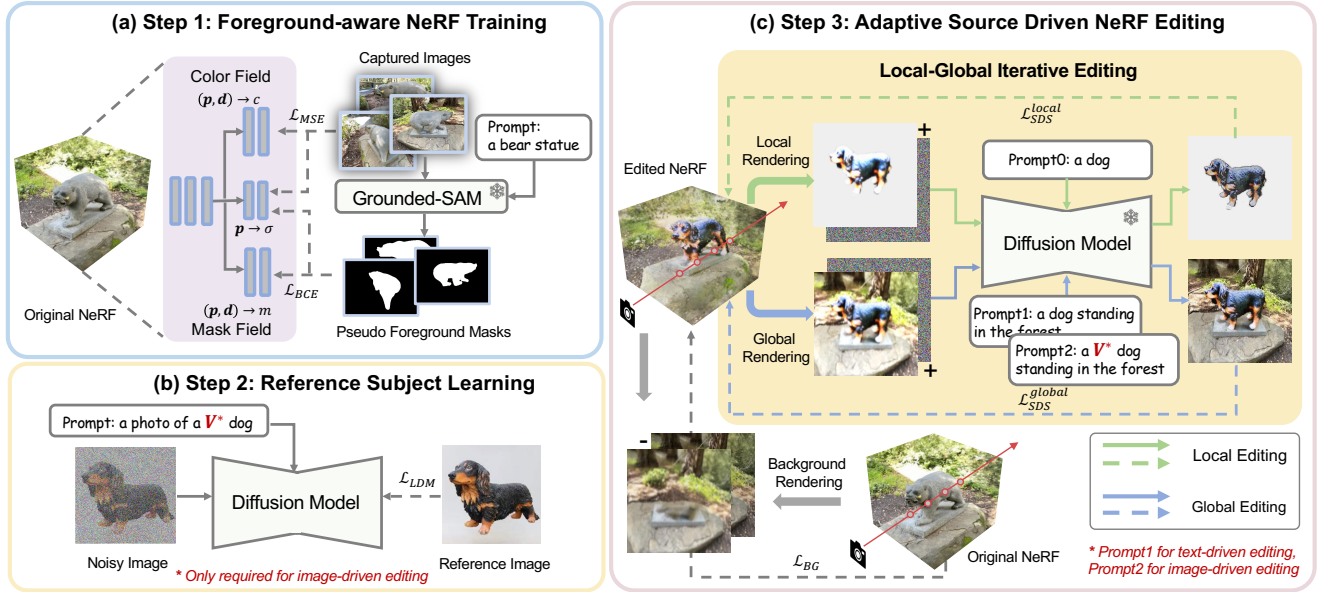


Figure 2. The overall pipeline of our CustomNeRF. Given an original 3D scene, CustomNeRF aims to edit it according to a textual description or a reference image. Our pipeline consists of three steps in total. (a) Reconstruct the original 3D scene with a foreground-aware NeRF, which enables the identification of foreground regions. (b) Finetune the pre-trained text-to-image diffusion model to embed the visual subject in the reference image into a special token V^* . This step is only required in image-driven editing. (c) Optimize the given NeRF with our proposed Local-Global Iterative Editing (LGIE) training scheme to align with the textual description or the reference image. The local and global data flows occur in an alternating manner. $Prompt2$ is employed in image-driven editing whereas $Prompt1$ is utilized in text-driven editing. $Prompt0$ is shared in both editing settings.

an extra mask field is introduced to estimate the editing probability besides regular color and density prediction as shown in Figure 2(a). The editing probability is utilized in the last step to discriminate regions to be edited (*i.e.*, foreground regions) from irrelevant regions (*i.e.*, background regions). Second, as illustrated in Figure 2(b), we employ a subject-aware Text-to-Image (T2I) diffusion model to encode the subject in the reference image as a special token V^* , thus forming a hybrid prompt for image-driven editing (not required for text-driven editing). Finally, we train the edited NeRF conditioned on the textual prompt or the hybrid prompt using the Local-Global Iterative Editing (LGIE) scheme, allowing it to render images that conform with the editing prompt while preserving background contents, which is presented in Figure 2(c). The class-guided regularization is further adopted in image-driven editing to alleviate the Janus problem.

4.1. Foreground-aware NeRF Training

Given a 3D point \mathbf{p} and a view direction \mathbf{d} , a typical NeRF [22] defines MLPs to output view-dependent color $c(\mathbf{p}, \mathbf{d})$ and the point density $\sigma(\mathbf{p})$. Considering a pixel's camera ray $\mathbf{r}(k) = \mathbf{o} + k\mathbf{d}$ starting from origin \mathbf{o} in direction \mathbf{d} , with k representing the distance from \mathbf{o} . The color

of this pixel can be obtained by volume rendering:

$$C(\mathbf{r}) = \int_{k_n}^{k_f} T(k)\sigma(\mathbf{r}(k))c(\mathbf{r}(k), \mathbf{d})dk, \quad (3)$$

where $T(k) = \exp\left(-\int_{k_n}^k \sigma(\mathbf{r}(k))dk\right)$,

where k_n and k_f are the near plane and the far plane, $T(k)$ can be regarded as transparency. MSE loss between the rendered color $\hat{C}(\mathbf{r})$ and the captured color $C(\mathbf{r})$ is calculated for supervision.

In our method, to identify foreground regions, we extend the typical NeRF with an extra mask field estimating the editing probability $m(\mathbf{p}, \mathbf{d}) \in (0, 1)$ of point \mathbf{p} from the view \mathbf{d} , which is shown in Figure 2(a). Pixel editing probability of ray \mathbf{r} , denoted by $\hat{M}(\mathbf{r})$ can be rendered in the same way as Equation 3. To acquire the ground truth of $\hat{M}(\mathbf{r})$, we incorporate Grounded SAM [11, 15, 19] to obtain the binary segmentation mask $M(\mathbf{r})$ of each captured image according to a textual prompt describing the edited region, *e.g.*, “a bear statue” in Figure 2(a). BCE loss is calculated between $\hat{M}(\mathbf{r})$ and the pseudo label $M(\mathbf{r})$.

4.2. Reference Subject Learning

In order to unify the formulation of image-driven and test-driven editing tasks, we adopt an efficient subject-aware

text-to-image generation model Custom Diffusion [16] for reference subject encoding. Custom Diffusion is built upon a pre-trained text-conditioned diffusion model and fine-tuned on several reference images to generate images of the specific subject while preserving its key identifying characteristics. As illustrated in Figure 2(b), a special word V^* is defined to uniquely identify the reference subject in the reference image, which is optimized by the LDM loss (Equation 1) to map it to a proper word embedding. After training, V^* can be treated as a regular word token to form a hybrid prompt, e.g., “a photo of a V^* dog”, for customized generation of the specific black dog. In this way, our CustomNeRF is able to conduct consistent and effective editing with adaptive types of source data including images or texts.

4.3. Adaptive Source Driven NeRF Editing

In this step, we aim to obtain an edited NeRF that can re-create the 3D scene aligning with the text description or the reference image. An intuitive way to obtain this edited NeRF is to first initialize it as a copy of the original NeRF obtained in Section 4.1 and then optimize it with the SDS loss calculated using the pre-trained T2I model to align with the textual or hybrid prompt. However, due to the implicit NeRF representation, naively applying SDS loss to the whole 3D scene results in changes to the background regions, which should be kept the same as the original scene. To this end, we propose a Local-Global Iterative Editing (LGIE) scheme that alternates between foreground region editing and full-image editing. Furthermore, to mitigate the Janus problem in image-driven editing, we design a class-guided regularization to exploit general class priors contained in the T2I model to inject novel concepts into scenes in a more geometrically consistent way. Details of the above training scheme can be found in Figure 2(c).

Local-Global Iterative Editing. Our LGIE training scheme alternates between the local editing stage which separately renders and edits the foreground regions only, and the global editing stage which edits foreground regions contextualized on background clues. We first take text-driven editing as an example for illustration. As illustrated in Figure 2(c), given a 3D scene of the bear statue, our goal is to edit it with the textual prompt “a dog standing in the forest” to replace the bear with a dog. In the local editing stage, we modify Equation 3 to render foreground regions using the editing probability $m(\mathbf{r}(k), \mathbf{d})$:

$$C_f(\mathbf{r}) = \int_{k_n}^{k_f} T_f(k) \sigma_f(\mathbf{r}(k), \mathbf{d}) c(\mathbf{r}(k), \mathbf{d}) dk, \quad (4)$$

where $T_f(k) = \exp\left(-\int_{k_n}^k \sigma_f(\mathbf{r}(k), \mathbf{d}) dk\right)$,

where point density $\sigma(\mathbf{r}_t)$ is adjusted by $m(\mathbf{r}_t, \mathbf{d})$ to eliminate the background contents in the rendered image as fol-

lows:

$$\sigma_f(\mathbf{r}(k), \mathbf{d}) = m(\mathbf{r}(k), \mathbf{d}) \cdot \sigma(\mathbf{r}(k)). \quad (5)$$

The rendered foreground image \mathbf{x}_f is then fed into the frozen T2I model together with the foreground part of the above editing prompt (e.g., “a dog”). Following [28], view-dependent words, e.g., “front view”, “side view”, and “back view”, etc, are appended to the prompt to alleviate Janus problem in text-driven editing. Due to the absence of a standardized definition of direction in the captured dataset, instead of manually rotating the scene to the standard front-view orientation, we conduct matching between the CLIP features of images and view-dependent descriptions to assign an appropriate view-dependent word to the rendered image. Afterward, the local SDS loss $\mathcal{L}_{SDS}^{local}$ formulated in Equation 2 is calculated to optimize the NeRF parameters.

In the global editing stage, we render the full image and feed it into the T2I model together with the complete editing prompt “a dog standing in the forest” to calculate a global SDS loss $\mathcal{L}_{SDS}^{global}$. In this stage, the context information in the background regions serves as a useful clue for editing, realizing foreground-background harmony. To further avoid changes to the context of the background region, we detach the gradients of the background region in fully rendered images to prevent them from affecting the NeRF parameters.

Class-guided regularization For image-driven editing where hybrid prompt (e.g., “a V^* dog standing in the forest”) is provided to replace the bear statue with a specific black dog, we apply a class-guided regularization method to mitigate the Janus problem which becomes severe due to the single-view reference image. Concretely, in the local editing stage, we remove the special token V^* from the foreground part of the editing prompt, so that the class word “dog” plays a major role in the local optimization process. In this manner, the general class priors embedded within the T2I model are effectively leveraged to guide geometrically reasonable optimization. The global stage runs in the same way as the text-driven setting, except that the hybrid prompt is adopted to inform the NeRF model with the appearance information of the reference subject. By alternating between local and global stages, the appearance of the reference subject can be transferred to various views gradually, producing view-consistent editing results.

Loss functions. Apart from the above two SDS losses, we further explicitly constrain the rendered pixel color of the background region to be the same as the original one for the goal of background preservation. We render the background regions before and after editing, which are denoted as \mathbf{x}_b^o and \mathbf{x}_b^e . Concretely, we first calculate the background point density $\sigma_b(\mathbf{r}(k), \mathbf{d})$ for background rendering:

$$\sigma_b(\mathbf{r}(k), \mathbf{d}) = (1 - (m(\mathbf{r}(k), \mathbf{d}))) \cdot \sigma(\mathbf{r}(k)), \quad (6)$$

where subscript o and e are omitted for simplicity. The background image is rendered using Equation 4 except that

$\sigma_f(\mathbf{r}(k), \mathbf{d})$ is replaced with $\sigma_b(\mathbf{r}(k), \mathbf{d})$. An MSE loss \mathcal{L}_{bg} is calculated between \mathbf{x}_f^o and \mathbf{x}_f^e to enforce the edited NeRF rendering similar background contents with the original NeRF.

The overall loss function for the third step is summarized as follows:

$$\mathcal{L} = \lambda_{SDS} \mathcal{L}_{SDS} + \lambda_{bg} \mathcal{L}_{bg}, \quad (7)$$

where λ_{SDS} and λ_{bg} are hyperparameters to balance different losses, which are empirically set as 0.01 and 1000 respectively. We use \mathcal{L}_{SDS} to represent both $\mathcal{L}_{SDS}^{local}$ and $\mathcal{L}_{SDS}^{global}$ which are alternately calculated during training.

5. Experiments

5.1. Experimental Setup

Datasets. To validate the effectiveness of our CustomNeRF, we test on 8 real scenes from BlendedMVS [49], LLFF [21], IBRNet [44], and Bear Statue [8]. These scenes feature complex backgrounds, including outdoor statues and objects frequently encountered in everyday life, *etc.* For text-driven editing, 4-6 textual prompts are adopted for each scene. For image-driven editing, we use images downloaded from the Internet as reference images, with 3-5 reference images available for each scene.

Baselines. For text-driven editing, we compare our method with two recent editing methods based on NeRF representation, including: (1) Instruct-NeRF2NeRF [8], which edits NeRF by iteratively updating dataset with images generated by a 2D editing model [4]. (2) RePaint-NeRF [51], which uses SDS loss for NeRF optimization and keeps non-editing regions unchanged by 2D semantic masks.

For image-driven editing, due to the absence of existing methods for this setting, we modified several existing works to serve as baselines for comparison with our method, including: (1) Ours+Splice Loss: It is hard to directly compare with SINE due to different task settings and its non-public training code. As an alternative, we combine Splice Loss [40], which plays an important role in transferring reference image information to the edited NeRF in SINE, with our method to build a feasible baseline. We apply this loss to our foreground-aware NeRF during the editing stage to transfer appearance information from the reference image while maintaining the structure layout of the original scene. (2) RePaint-NeRF*: Since the original framework of RePaint-NeRF only supports text-driven editing, we replace its employed T2I model with the Custom Diffusion finetuned on reference images. In this way, the reference image can be embedded as a special token in the editing prompt, thereby equipping RePaint-NeRF with the capability for image-driven editing.

Evaluation metrics. Following [8, 52], we use the CLIP Text-Image Directional Similarity to evaluate the alignment

between the images and text prompts before and after editing, reflecting whether the image editing direction is consistent with the text changing direction. The DINO [26] Similarity used in subject-driven image generation is also adopted as a metric. By calculating the average of the DINO similarity between the edited images from different views and the reference image, we can measure whether the edited NeRF is consistent with the reference image. Furthermore, to compensate for the lack of objective metrics in subjective editing tasks, we also conduct user studies in text-driven and image-driven editing.

Implementation details. In our experiments, we employ Instant-NGP [24] to learn the original neural field and make necessary extensions to enable the prediction of editing probabilities. The training of original NeRF takes 3K iterations with the supervision of both rendered image and editing probability. We use Stable Diffusion v1.5 and follow the Custom Diffusion training process to conduct reference subject learning with 250 steps. In the editing step, we use the Adam optimizer [14] with a learning rate of $5e-4$ to optimize the NeRF parameters. The duration of the editing process varies from 10 minutes to 1 hour depending on the complexity of the scene and the editing type on a single NVIDIA RTX 3090 GPU. We implement our model using PyTorch [27].

5.2. Qualitative Results

Image-driven editing. We provide qualitative results of image-driven editing in Figure 3 to compare our CustomNeRF with other baseline methods. It is shown in Figure 3(b) that Splice Loss can only transfer superficial attributes such as color and texture from the reference image to the edited NeRF, while it fails to modify the geometry of the edited object. For instance, in the 3rd row, where the statue of David is replaced with a comic character, it only turns the statue into red. Although RePaint-NeRF* (Figure 3(c)) learns the texture and geometry information of the reference image to some extent, it is still limited in the richness of details and similarity to the reference image, *e.g.*, the hairstyle and clothes texture of the comic character are not presented in the edited NeRF. Moreover, background contents are not well-preserved in this method. In contrast, our method is able to generate edited results that not only maintain the background effectively but also bear a closer resemblance to the reference image compared to the two baselines. For example, in the 1st row, our method successfully maintains the ring of yellow fur on the leg of the black dog. In the more challenging scenario depicted in the 3rd row, our approach preserves intricate details such as the hairstyle of the comic character and the folds in the clothes, which further demonstrates the superiority of our CustomNeRF. More visualization results are included in supplementary materials.

Text-driven editing. We also provide qualitative results

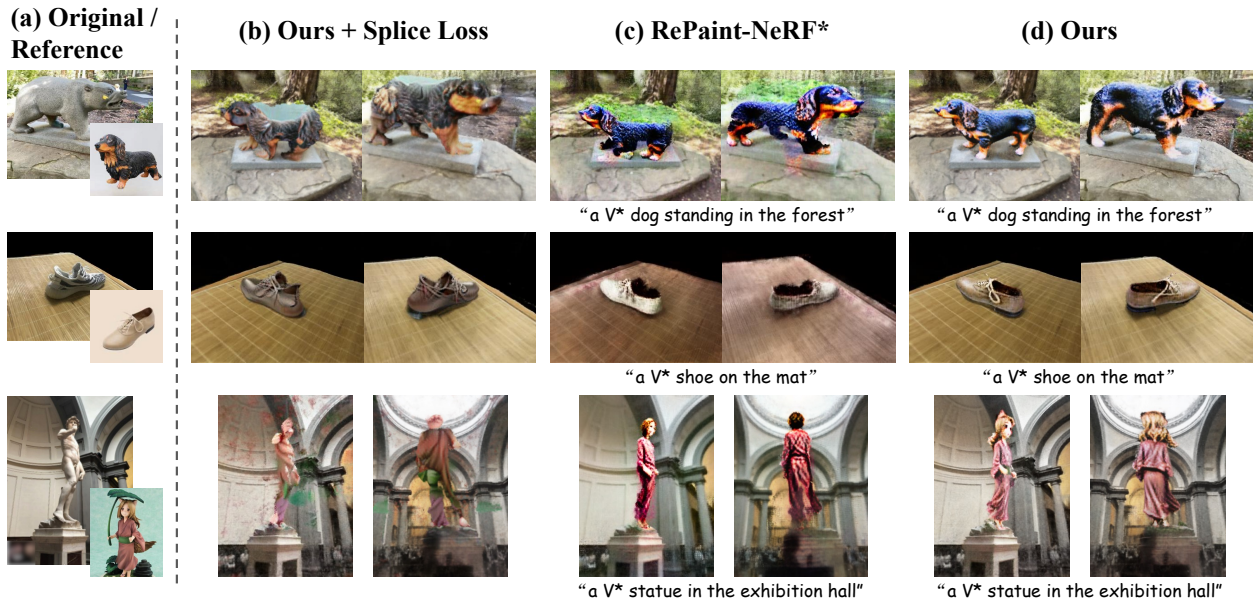


Figure 3. Qualitative comparison on the image-driven editing setting. * denotes replacing the employed T2I model with Custom Diffusion.

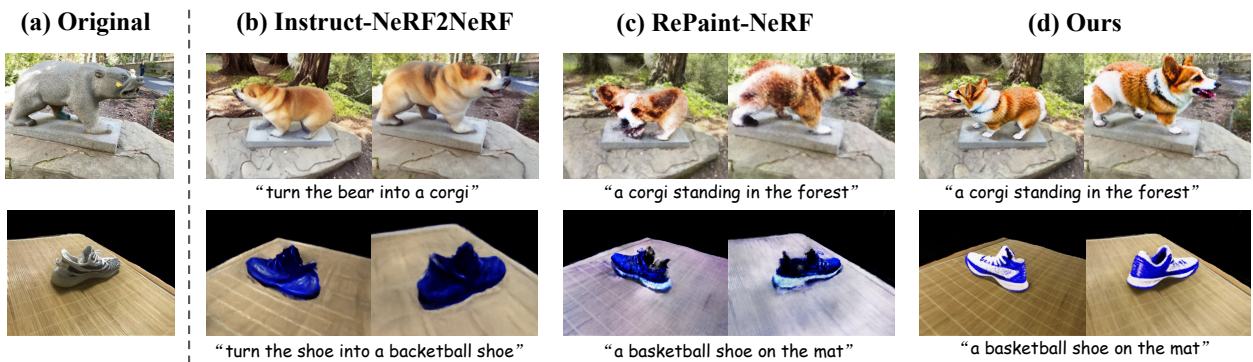


Figure 4. Qualitative comparison on the text-driven editing setting.

of text-driven editing in Figure 4. Compared to baseline methods, our approach is capable of realizing editing with more substantial shape and semantic changes and maintaining consistency between different views in complex scenes, such as transforming a bear statue into a corgi dog. Instruct-NeRF2NeRF still retains the bear’s body shape and RePaint-NeRF tends to produce abnormal renderings. In terms of background preservation, our method also significantly outperforms the baseline approaches, particularly RePaint-NeRF, which tends to have its background styles easily influenced by the foreground content.

5.3. Quantitative results

We present the quantitative comparison with other methods on the image-driven editing and text-driven editing tasks in Table 1 and Table 2 respectively. For user study, we distribute 50 questionnaires for both image-driven editing and text-driven editing, including 10 image-driven edits across

Method	CLIP _{dir} ↑	DINO _{sim} ↑	Vote Percentage ↑
Ours+Splicing Loss	15.47	46.51	10.6%
RePaint-NeRF*	16.46	38.74	13.8%
CustomNeRF (Ours)	20.07	47.44	75.6%

Table 1. Quantitative comparison on the image-driven editing setting. CLIP_{dir} is short for CLIP Text-Image Directional Similarity and DINO_{sim} is short for DINO Similarity.

Method	CLIP _{dir} ↑	Vote Percentage ↑
Instruct-NeRF2NeRF	11.48	10.6%
RePaint-NeRF	16.34	4.2%
CustomNeRF (Ours)	22.66	85.2%

Table 2. Quantitative comparison on the text-driven editing setting. CLIP_{dir} is short for CLIP Text-Image Directional Similarity.

5 scenes and 10 text-driven edits across 2 scenes, asking

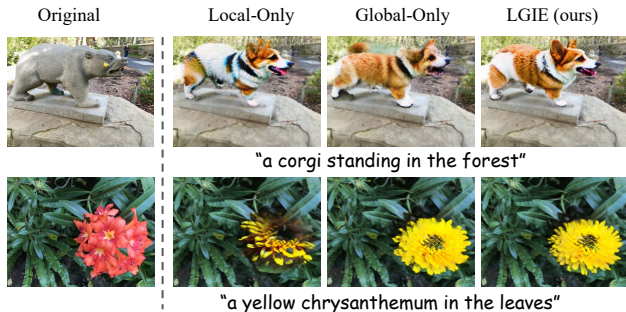


Figure 5. Ablation study of Local-Global Iterative Editing. Local-only and Global-only denote editing with only local SDS loss and global SDS loss respectively.

users to choose the best editing results. Our CustomNeRF significantly outperforms baseline methods and achieves the highest editing scores in both tasks, especially in user evaluation performances. Although Ours+Splicing Loss in Table 1 also achieves a relatively high DINO Similarity, this is largely due to the fact that it is optimized based on DINO-ViT features, giving it a significant advantage when computing DINO feature similarity. However, as indicated by the other metrics and qualitative results in Figure 3, it is substantially inferior in performance compared to our method.

5.4. Ablation study

Effectiveness of Local-Global Iterative Editing. We conduct ablation studies on our LGIE training scheme in the text-driven editing task in Figure 5. For the local-only setting, conducting local editing without considering the global context can lead to a discordance between the generated foreground content and the background region. For instance, in the second row, the right part of the sunflower is obscured by leaves. Conversely, performing global-only editing may change the background region, especially in complex scenes, *e.g.*, the blurred background area above the dog in the first row. By iteratively optimizing between these two stages, we can achieve better background preservation and foreground-background harmony.

Class-guided regularization. We conduct ablation studies on class-guided regularization and present the results in Figure 6. Removing the class prior regularization from both local and global stages results in edited outputs that more closely resemble the reference image, but also lead to abnormal views and greater influence from the background in the reference image. Applying class prior to the global stage leads to an obvious Janus problem, indicating that class priors tend to work in scenarios with simpler backgrounds. We include more ablation results in supplementary materials.



Figure 6. Ablation study of class-guided regularization. In the first column, large images depict the original scene, and small images with black borders are reference images. w/o Class Reg. denotes apply V^* to both global and local stages. Global Class Reg. denotes removing V^* from the global stage, and Local Class Reg. applies similar removals in the local stage.

6. Conclusion and Limitation

In this study, we present the CustomNeRF model, designed for adaptive source driven 3D scene editing with unified editing prompts from either texts or reference images. We tackle two key challenges: ensuring precise foreground-only editing, and maintaining consistency across multiple views when using single-view reference images. Our approach includes a Local-Global Iterative Editing (LGIE) training scheme for editing that focuses on the foreground while keeping the background unchanged, and a class-guided regularization to mitigate view inconsistency in image-driven editing. Extensive experiments validate CustomNeRF’s capability to accurately edit in diverse real-world scenarios for both text and image prompts.

Our method transfers subject appearance from reference images, yet it faces limitations with Custom Diffusion’s inability to perfectly replicate references, resulting in some inconsistencies in details in our image-driven editing results. In addition, we are also facing issues inherited from DreamFusion, such as oversaturation of generated results. Currently, our approach is also limited to text and image prompts, but future work could expand to include other editing sources like audio, sketches, etc.

7. Acknowledgments

This research is supported in part by National Key R&D Program of China (2022ZD0115502), National Natural Science Foundation of China (Grant No. 62122010, U23B2010, 62322608), Zhejiang Provincial Natural Science Foundation of China under Grant No. LDT23F02022F02, Key Research and Development Program of Zhejiang Province under Grant 2022C01082, CAAI-MindSpore Open Fund, developed on OpenI Community, and the Open Project Program of State Key Laboratory of Virtual Reality Technology and Systems, Beihang University (No.VRLAB2023A01).

References

- [1] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *CVPR*, 2022. 3
- [2] Chong Bao, Yinda Zhang, Bangbang Yang, Tianxing Fan, Zesong Yang, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Sine: Semantic-driven image-based nerf editing with prior-guided editing field. In *CVPR*, 2023. 3
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 2
- [4] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023. 2, 3, 6
- [5] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023. 2
- [6] 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. *arXiv preprint arXiv:2208.01618*, 2022. 3
- [7] Stephan J Garbin, Marek Kowalski, Virginia Estellers, Stanislaw Szymanowicz, Shideh Rezaeifar, Jingjing Shen, Matthew Johnson, and Julien Valentin. Voltmorph: Real-time, controllable and generalisable animation of volumetric representations. *arXiv preprint arXiv:2208.00949*, 2022. 2
- [8] 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. 2, 3, 6
- [9] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 3
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 3
- [11] IDEA-Research. Grounded-sam. <https://github.com/IDEA-Research/Grounded-Segment-Anything>. 4
- [12] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *CVPR*, 2022. 3
- [13] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [15] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023. 4
- [16] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *CVPR*, 2023. 2, 3, 5
- [17] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, 2023. 3
- [18] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *ICCV*, 2021. 3
- [19] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 4
- [20] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 3
- [21] 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. *TOG*, 2019. 6
- [22] 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. *Communications of the ACM*, 2021. 4
- [23] 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. *Communications of the ACM*, 2021. 2
- [24] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *TOG*, 2022. 6
- [25] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *ICCV*, 2021. 3
- [26] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023. 6
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 2019. 6
- [28] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 3, 5
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learn-

- ing transferable visual models from natural language supervision. In *ICML*, 2021. 3
- [30] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Nataniel Ruiz, Ben Mildenhall, Shiran Zada, Kfir Aberman, Michael Rubinstein, Jonathan Barron, et al. Dreambooth3d: Subject-driven text-to-3d generation. *arXiv preprint arXiv:2303.13508*, 2023. 2, 3
- [31] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. 2
- [32] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [33] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *NeurIPS*, 2019. 2
- [34] Elad Richardson, Gal Metzer, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. Texture: Text-guided texturing of 3d shapes. *arXiv preprint arXiv:2302.01721*, 2023. 2, 3
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3
- [36] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 3
- [37] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 2022. 2
- [38] Aditya Sanghi, Hang Chu, Joseph G Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshah. Clip-forge: Towards zero-shot text-to-shape generation. In *CVPR*, 2022. 3
- [39] Etai Sella, Gal Fiebelman, Peter Hedman, and Hadar Averbuch-Elor. Vox-e: Text-guided voxel editing of 3d objects. In *CVPR*, 2023. 2, 3
- [40] Narek Tumanyan, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Splicing vit features for semantic appearance transfer. In *CVPR*, 2022. 6
- [41] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *CVPR*, 2022. 3
- [42] 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
- [43] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 2
- [44] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. 6
- [45] Fanbo Xiang, Zexiang Xu, Milos Hasan, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Hao Su. Neutex: Neural texture mapping for volumetric neural rendering. In *CVPR*, 2021. 2
- [46] Jiale Xu, Xintao Wang, Weihao Cheng, Yan-Pei Cao, Ying Shan, Xiaohu Qie, and Shenghua Gao. Dream3d: Zero-shot text-to-3d synthesis using 3d shape prior and text-to-image diffusion models. In *CVPR*, 2023. 3
- [47] Tianhan Xu and Tatsuya Harada. Deforming radiance fields with cages. In *ECCV*, 2022. 2
- [48] Bangbang Yang, Chong Bao, Junyi Zeng, Hujun Bao, Yinda Zhang, Zhaopeng Cui, and Guofeng Zhang. Neumesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In *ECCV*, 2022. 3
- [49] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *CVPR*, 2020. 6
- [50] Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields. In *ECCV*, 2022. 2
- [51] Xingchen Zhou, Ying He, F Richard Yu, Jianqiang Li, and You Li. Repaint-nerf: Nerf editing via semantic masks and diffusion models. *arXiv preprint arXiv:2306.05668*, 2023. 2, 3, 6
- [52] Jingyu Zhuang, Chen Wang, Lingjie Liu, Liang Lin, and Guanbin Li. Dreameditor: Text-driven 3d scene editing with neural fields. *arXiv preprint arXiv:2306.13455*, 2023. 2, 3, 6