

# GaussianEditor: Editing 3D Gaussians Delicately with Text Instructions

Junjie Wang\*, Jiemin Fang\*, Xiaopeng Zhang, Lingxi Xie, Qi Tian  
 Huawei Inc.

{is.wangjunjie, jaminfong, zxphistory, 198808xc}@gmail.com tian.qil@huawei.com

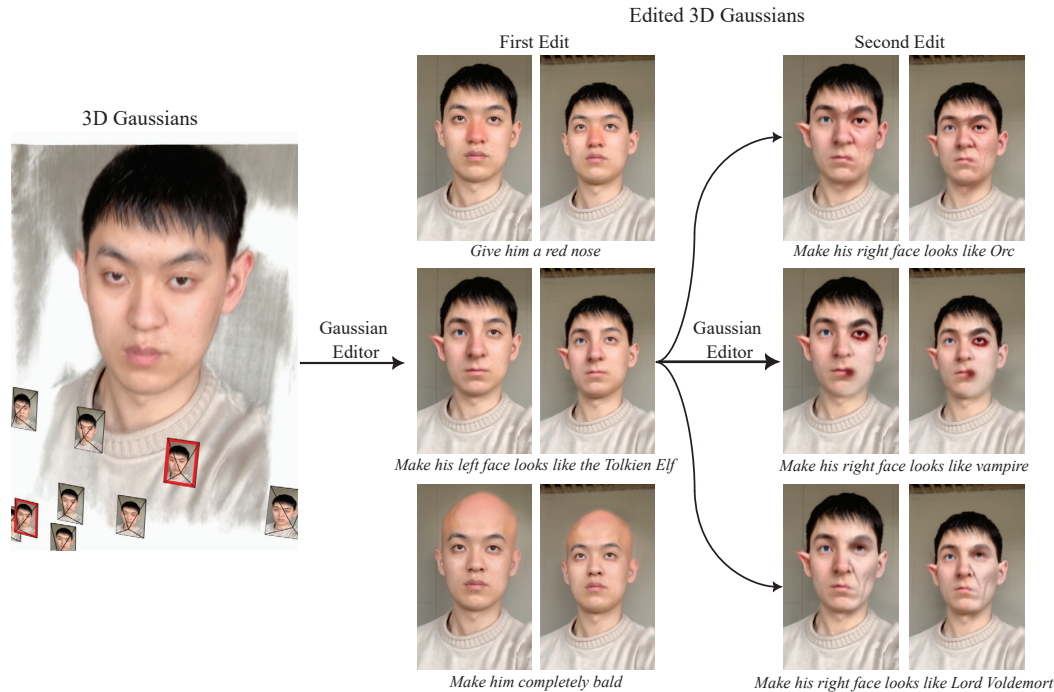


Figure 1. We propose GaussianEditor, an interactive framework to achieve delicate 3D scene editing following text instructions. As shown in this figure, our method can precisely control the editing region and achieve multi-round editing.

## Abstract

Recently, impressive results have been achieved in 3D scene editing with text instructions based on a 2D diffusion model. However, current diffusion models primarily generate images by predicting noise in the latent space, and the editing is usually applied to the whole image, which makes it challenging to perform delicate, especially localized, editing for 3D scenes. Inspired by recent 3D Gaussian splatting, we propose a systematic framework, named GaussianEditor, to edit 3D scenes delicately via 3D Gaussians with text instructions. Benefiting from the explicit property of 3D Gaussians, we design a series of techniques to achieve delicate editing. Specifically, we first extract the region of interest (RoI) corresponding to the text instruction, aligning it to 3D Gaussians. The Gaussian RoI is fur-

ther used to control the editing process. Our framework can achieve more delicate and precise editing of 3D scenes than previous methods while enjoying much faster training speed, i.e. within 20 minutes on a single V100 GPU, more than twice as fast as Instruct-NeRF2NeRF (45 minutes – 2 hours)<sup>1</sup>. The project page is at [GaussianEditor.github.io](https://github.com/GaussianEditor).

## 1. Introduction

Creating 3D assets has played a critical role in many applications and industries, e.g. movie/game production, artistic creation, AR, VR etc. However, this process is usually expensive and cumbersome, especially for traditional pipelines. Designers need to take a lot of labor and time to

\*Equal contributions.

†Corresponding author.

<sup>1</sup>The editing time varies in different scenes according to the scene structure complexity.

finish each step, *e.g.* sketching, building structures, creating textures *etc.* One cheap and effective way of creating high-quality 3D assets is to start from an existing scene, capturing, modeling, and editing the scene and obtaining the wanted one. This approach can be also used for user-interactive entertainment applications.

Neural radiance field methods [2, 3, 6, 29, 31, 42, 45] have shown great power in representing 3D scenes and synthesizing novel-view images. Past years have witnessed the rapid development of NeRF and its variants, from both quality and efficiency perspectives. Editing a pre-trained NeRF model becomes a promising way to edit 3D scenes. Represented by Instruct-NeRF2NeRF [11], researchers propose to use the image-conditioned 2D diffusion model, *e.g.* InstructPix2Pix [4], to edit 3D scenes simply with text instructions. Notable results have been achieved as real scenes can be changed following the text instruction. However, current 2D diffusion models face challenges in accurately localizing editing regions, which hinders the generation of finely edited scenes due to the change of unintended regions. Even though some works [30] propose to constrain the editing region on edited 2D images, the editing region is not accurately localized and hard to apply to the 3D representation. Besides, NeRF-based methods [9, 48] bear coupling effects between different spatial positions, *e.g.* different points are queried from the same MLP field (for implicit representations) or voxel vertices (for explicit representations).

Recent 3D Gaussian Splatting [18] (3D-GS) has been a groundbreaking work in the radiance field, which is the first to achieve a real sense of real-time rendering while enjoying high rendering quality and training speed. Besides its efficiency, we further notice its natural explicit property. 3D-GS has a great advantage for editing tasks as each 3D Gaussian exists individually. Editing 3D scenes by directly manipulating 3D Gaussians with desired constraints is easy.

Aiming at editing 3D scenes delicately, we propose to represent the scene with 3D Gaussians, which can be edited with text instructions, and name our method as GaussianEditor. GaussianEditor is divided into three main parts to achieve precise control for editing regions. The first is the region of interest (RoI) extraction from the given text instruction. The instruction may be complex or indirect while this module helps extract the keywords matching the RoI for editing. The second part aligns the extracted text RoI to the 3D Gaussian space through the image space, where a grounding segmentation module is applied. The last part is to edit the original 3D Gaussians delicately with constraints in the obtained 3D Gaussian RoI. With the above processes, the region for editing can be precisely localized simply from text instructions, which constrains the 3D Gaussian updating to obtain a delicately edited new 3D scene. Besides, we enable interfaces for users to introduce more exact instructions for more delicate editing, *e.g.* Gaussian point selecting

and 3D boxes for modifying the editing regions<sup>2</sup>.

Our contributions can be summarized as follows.

- As far as we know, our GaussianEditor is one of the first systematic methods to achieve delicate 3D scene editing based on 3D Gaussian splatting.
- A series of techniques are designed and proposed to precisely localize the editing region of interest, which are aligned and applied to 3D Gaussians. Though some sub-modules are from existing works, we believe integrating these awesome techniques to work effectively is a valuable topic, which is what we focus on in this paper.
- Our method achieves a series of more delicate editing results compared with the previous representative work Instruct-NeRF2NeRF [11] with much shorter training time (within 20 minutes *v.s.* 45 minutes – 2 hours).

## 2. Related Work

**2D Image Editing with Diffusion Models.** By advancements in diffusion model technology [13, 43], numerous generative models [40] have attained remarkably impressive outcomes in image synthesis. Recent progressions in diffusion models have showcased their adeptness in fabricating lifelike images from arbitrary textual inputs [8, 14, 39, 41, 44]. Harnessing the robust semantic comprehension and image generation capabilities of foundational diffusion models, an escalating number of research explorations are currently employing diffusion models as a fundamental framework for implementing text-based image editing functionalities [33, 36, 37, 40]. Some of these methodologies necessitate the manual provision of captions for both the original and edited images [12], while others mandate specific scenario-based training for optimization [38]. These requisites have rendered it arduous for ordinary users to avail themselves of such techniques. Expanding upon this foundation, iP2P [4] introduces instruction-based capabilities to image editing, enabling users to simply input an image and apprise the model of the desired alterations. This user-friendly approach facilitates the democratization of image editing in a more accessible manner.

**3D Scene Editing of Radiance Fields.** 3D Scene Editing of Radiance Fields has gained significant popularity as a recent research direction [1, 10, 15, 20, 22–25, 28, 34, 46, 47, 51–53]. The main objective of such methods is to enable the manipulation of both the geometry and appearance of 3D scene representations. However, editing such scenes poses inherent challenges due to the implicit nature of traditional NeRF representations, which lack precise localization capabilities. As a result, previous works have primarily focused on achieving global style transformations of 3D scenes [7, 16, 17, 32, 48, 50, 54]. While some efforts

<sup>2</sup>These additional instructions are applied to generate the man with two different edited half faces in Fig. 1.

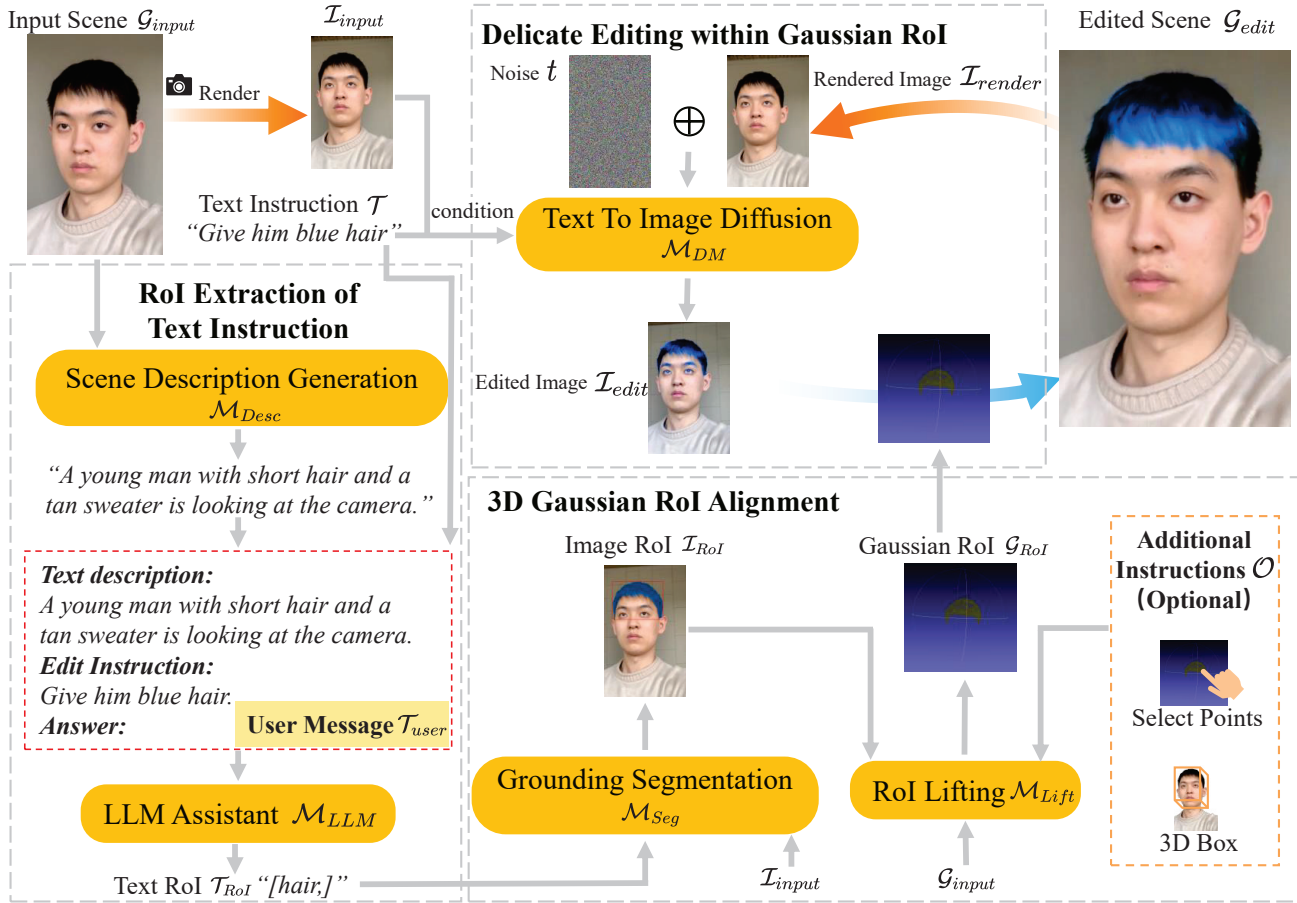


Figure 2. Our framework, named GaussianEditor, consists of three key steps. First, a module  $\mathcal{M}_{Desc}$  is used to get the description of the input scene, which is put to an LLM assistant  $\mathcal{M}_{LLM}$  with the text instruction  $\mathcal{T}$  provided by the user to obtain the text ROI  $\mathcal{T}_{RoI}$ . Second, a grounding segmentation module  $\mathcal{M}_{Seg}$  is used to convert  $\mathcal{T}_{RoI}$  to image ROI  $\mathcal{I}_{RoI}$ , which is then lifted to 3D Gaussians ROI  $\mathcal{G}_{RoI}$  by ROI lifting  $\mathcal{M}_{Lift}$ , where additional user instructions  $\mathcal{O}$  can be incorporated. Third, following the user instruction  $\mathcal{T}$ , rendered image  $\mathcal{I}_{render}$  from randomly chosen views is edited by a diffusion model  $\mathcal{M}_{DM}$ . The loss between  $\mathcal{I}_{render}$  and edited one  $\mathcal{I}_{edit}$  is calculated. Finally, gradient backpropagation and optimization are performed within the Gaussian ROI  $\mathcal{G}_{RoI}$  to get the edited scene  $\mathcal{G}_{edit}$ .

have been made towards object-centric scene editing[55], keeping the background unchanged has been a persistent challenge. For example, a recently proposed work, Instruct-NeRF2NeRF [11], implements text instruction-controlled 3D scene editing capabilities, achieving excellent editing effects while maintaining user-friendliness. However, it relies on the editing effect of 2D images, which may cause global changes to the 3D scene. A subsequent work [30] attempts to compute the relevance map between edited and unedited images to localize the editing area. The relevance map may be unreliable when the 2D IP2P [4] model fails. Other efforts [23] attempt to rely on the 3D coordinates entered by the user to decide the editing area. The introduction of 3D Gaussians [18] has provided an opportunity to address this limitation. Its explicit 3D representation methodology enables accurate selection and manipulation of editing areas. By incorporating LLMs, the whole process can be more automated.

### 3. Method

In this section, we first review 3D representation methods in Sec. 3.1. Subsequently, in Sec. 3.2, we overview our proposed approach, which mainly includes three modules. Sec. 3.3 delves into the precise Region of Interest (RoI) extraction of text instructions, using scene description generation module  $\mathcal{M}_{Desc}$  and LLM assistant  $\mathcal{M}_{LLM}$ . Sec. 3.4 introduces how to align the instruction RoI with 3D Gaussians, using grounding segmentation module  $\mathcal{M}_{Seg}$  and ROI lifting module  $\mathcal{M}_{Lift}$ . Finally, Sec. 3.5 describes the delicate editing process within the obtained Gaussian ROI, using text to image diffusion model  $\mathcal{M}_{DM}$ .

#### 3.1. Preliminaries

**3D Gaussian Splatting.** 3D Gaussian splatting [18] is a recent powerful 3D representation method. It represents the 3D scene with point-like 3D Gaussians  $\mathcal{G} = \{g_1, g_2 \dots g_N\}$ ,

where  $g_i = \{\mu, \Sigma, c, \alpha\}$  and  $i \in \{1, \dots, N\}$ . Among them,  $\mu \in \mathbb{R}^3$  is the position where the Gaussian centers,  $\Sigma \in \mathbb{R}^7$  denotes the 3D covariance matrix,  $c \in \mathbb{R}^3$  is the RGB color and  $\alpha \in \mathbb{R}^1$  is the opacity. Benefitting from the compact representation of Gaussians and efficient differentiable rendering approach, 3D Gaussian splatting achieves real-time rendering with high quality. The splatting rendering process can be formulated as

$$C = \sum_{i \in \mathcal{N}} c_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j), \quad (1)$$

where  $\sigma_i = \alpha_i e^{-\frac{1}{2}(x_i)^T \Sigma^{-1}(x_i)}$  represents the influence of the Gaussian to the image pixel and  $x_i$  is the distance between the 3D point and the center of the  $i$ -th Gaussian.

### 3.2. Overall Framework

Given a group of 3D Gaussians  $\mathcal{G}_{input}$  for an input scene and a text instruction  $\mathcal{T}$  for editing, our Gaussian editor  $\mathcal{E}$  can edit the 3D Gaussians delicately into a new one, denoted as  $\mathcal{G}_{edit}$ , with the guidance of the instruction. The whole process can be formulated as

$$\mathcal{G}_{edit} = \mathcal{E}(\mathcal{G}_{input}, \mathcal{T}). \quad (2)$$

Fig. 2 illustrates the overall framework of our approach, which consists of three main steps. First, the Region of Interest (RoI) is extracted from the text instruction. In this step, we employ a module named scene description generation  $\mathcal{M}_{Description}$  to get the description of the input scene. We then input the scene description  $\mathcal{T}_{scene}$  and text instruction  $\mathcal{T}$  into a large language model assistant  $\mathcal{M}_{LLM}$  to determine where we should make edits in the scene. The output of this step is referred to as the instruction RoI  $\mathcal{T}_{RoI}$ .

The next step is the 3D Gaussian RoI alignment. We use a grounding segmentation module  $\mathcal{M}_{Seg}$  to convert the RoI from text space, *i.e.*  $\mathcal{T}_{RoI}$ , to the image space, *i.e.*  $\mathcal{I}_{RoI}$ . Then the image RoI  $\mathcal{I}_{RoI}$  is lifted to the RoI of 3D Gaussians  $\mathcal{G}_{RoI}$  through RoI lifting module  $\mathcal{M}_{Lift}$ . The Gaussian RoI allows us to control the regions where edits will be applied precisely.

The last step is delicate editing within the Gaussian RoI. In this step, we randomly sample the view to obtain the rendered image  $\mathcal{I}_{render}$ . A 2D diffusion model  $\mathcal{M}_{DM}$  is used to perform the editing process on the rendered image  $\mathcal{I}_{render}$ , with the user instruction  $\mathcal{T}$  and the image  $\mathcal{I}_{input}$  of input scene as conditions. The resulting edited image is denoted as  $\mathcal{I}_{edit}$ . Subsequently, we calculate the loss between  $\mathcal{I}_{edit}$  and  $\mathcal{I}_{render}$  and make gradient back-propagation within  $\mathcal{G}_{RoI}$ . This implies that only the regions specified by the RoI can receive corresponding gradients during the back-propagation process. Finally, optimization is executed based on these gradients. The final optimized scene representation  $\mathcal{G}_{edit}$  is obtained through several rounds of iterative optimization.

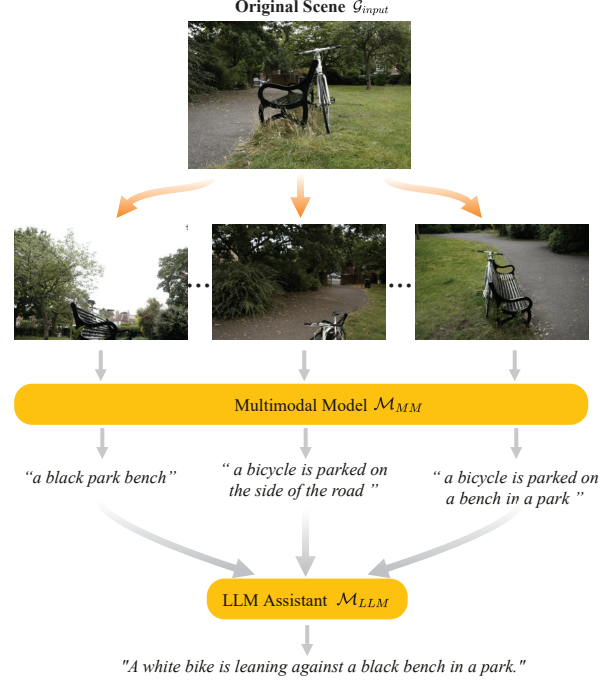


Figure 3. The process of obtaining scene description.

### 3.3. RoI Extraction of Text Instruction

The instruction RoI is extracted for the editing regions from both the input 3D scene  $\mathcal{G}_{input}$  and the text instruction  $\mathcal{T}$  provided by the user. To achieve this, we employ a multimodal model  $\mathcal{M}_{MM}$  in conjunction with the large language model assistant  $\mathcal{M}_{LLM}$ . The first step is the scene description generation  $\mathcal{M}_{Desc}$ , which aims to get the scene description  $\mathcal{T}_{scene}$  from 3D Gaussians  $\mathcal{G}_{input}$ :

$$\mathcal{T}_{scene} = \mathcal{M}_{Desc}(\mathcal{G}_{input}). \quad (3)$$

The process of the scene description generation  $\mathcal{M}_{Desc}$  is shown in Fig. 3. By leveraging the technique of differentiable splatting as shown in Eq. 1, a set of 2D image samples  $\{\mathcal{I}_{sample}\}$  are generated and then inputted into a multimodal model  $\mathcal{M}_{MM}$  to generate corresponding text descriptions  $\{\mathcal{T}_{sample}\}$ :

$$\mathcal{T}_{sample} = \mathcal{M}_{MM}(\mathcal{P}_{MM}, \mathcal{I}_{sample}), \quad (4)$$

where  $\mathcal{P}_{MM}$  is a prompt, such as “What is the content of the image”, for multimodal model  $\mathcal{M}_{MM}$  to get precise description. Subsequently, these descriptions  $\{\mathcal{T}_{sample}\}$  are fed into a large language model  $\mathcal{M}_{LLM}$ , which is specifically instructed by a prompt  $\mathcal{P}_{merge}$  to merge descriptions of diverse views into one detailed scene description  $\mathcal{T}_{scene}$ :

$$\mathcal{T}_{scene} = \mathcal{M}_{LLM}(\mathcal{P}_{merge}, \{\mathcal{T}_{sample}\}). \quad (5)$$

After that, the scene description  $\mathcal{T}_{scene}$  and the user instruction  $\mathcal{T}$  are combined with a predefined tem-

plate  $\mathcal{T}_{template}$ : “Text description:  $\mathcal{T}_{scene}$  Edit Instruction:  $\mathcal{T}_{Answer}$ .” to form the user message  $\mathcal{T}_{user} = \mathcal{T}_{template}(\mathcal{T}_{scene}, \mathcal{T})$ . The LLM model  $\mathcal{M}_{LLM}$  is used to extract the instruction RoI  $\mathcal{T}_{RoI}$  from user message  $\mathcal{T}_{user}$  with a new prompt  $\mathcal{P}_{extract}$ :

$$\mathcal{T}_{RoI} = \mathcal{M}_{LLM}(\mathcal{P}_{extract}, \mathcal{T}_{user}). \quad (6)$$

### 3.4. 3D Gaussian RoI Alignment

To confine the 3D editing region within the instruction RoI, 3D Gaussian RoI  $\mathcal{G}_{RoI}$  is aligned with the text RoI  $\mathcal{T}_{RoI}$ . First, The RoI in the text space is transformed into the image space via a grounding segmentation module  $\mathcal{M}_{Seg}$ :

$$\mathcal{I}_{RoI} = \mathcal{M}_{Seg}(\mathcal{I}_{input}, \mathcal{T}_{RoI}), \quad (7)$$

where  $\mathcal{I}_{input}$  is rendered image of the input scene  $\mathcal{G}_{input}$ .

Then we lift the the RoI  $\mathcal{I}_{RoI}$  in the image space to 3D Gaussian  $\mathcal{G}_{RoI}$  through training. To achieve this, an additional RoI attribute  $r \in \mathbb{R}^1$  was added to 3D Gaussian  $g_i = \{\mu_i, \Sigma_i, c_i, \alpha_i, r_i\}$ .  $r$  is initialized to 0, which means it is not in the Gaussians RoI, and 1 means it is inside the RoI. The set of  $r$  is denoted as  $\mathcal{R} \in \mathbb{R}^{\mathcal{N}, 1}$ , where the  $\mathcal{N}$  is the number of 3D Gaussians  $\mathcal{G}_{input}$ .

Then the color  $c_i$  in Eq. 1 was rewritten with  $r_i$  to get the rendered RoI  $\mathcal{I}_{RoI}^{render}$ :

$$\mathcal{I}_{RoI}^{render} = \sum_{i \in \mathcal{N}} r_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j). \quad (8)$$

Taking inspiration from SA3D [5], to get the trained Gaussians RoI  $\mathcal{G}_{RoI}^{train}$ , we adopt a similar loss function to supervise the training process:

$$\mathcal{L}_{proj} = \lambda_1 \sum (\mathcal{I}_{RoI}^{render} \cdot \mathcal{I}_{RoI}) + \lambda_2 \sum ((1 - \mathcal{I}_{RoI}^{render}) \cdot \mathcal{I}_{RoI}), \quad (9)$$

where  $\lambda_1$  and  $\lambda_2$  are hyperparameters. The  $r$  in Eq. 8 is updated via  $r \leftarrow r - \eta \frac{\partial \mathcal{L}_{proj}}{\partial r}$  with gradient descent, where  $\eta$  denotes the learning rate. Eq. 9 encourages rendered RoI to cover the Image RoI and not exceed it. Additionally, the user can modify the trained Gaussian RoI  $\mathcal{G}_{RoI}^{train}$  by giving added Gaussian RoI  $\mathcal{G}_{RoI}^{add}$ , deleted Gaussian RoI  $\mathcal{G}_{RoI}^{del}$  and 3D box  $\mathcal{B}_{3D}$ :

$$\mathcal{G}_{RoI} = (\mathcal{G}_{RoI}^{train} \cup \mathcal{G}_{RoI}^{add} - \mathcal{G}_{RoI}^{del}) \cap \mathcal{B}_{3D}, \quad (10)$$

$\mathcal{G}_{RoI}^{add}$  represents the 3D Gaussians user wants to edit,  $\mathcal{G}_{RoI}^{del}$  means 3D Gaussians user wants to keep from editing,  $\mathcal{B}_{3D}$  is the coordinates of 3D cuboid it limits RoI to inside the box.  $\mathcal{G}_{RoI}$  is the aligned RoI with the text RoI. For example, when editing the left face of the man in Fig. 1, grounding segmentation failed to ground “left face”, instead, it grounded the whole face. In this scenario, the user can use the interactive interface to set the right face as  $\mathcal{G}_{RoI}^{del}$  or enter

the rectangular box where the left face is located as  $\mathcal{B}_{3D}$ . The lifting process  $\mathcal{M}_{Lift}$  can be represented as:

$$\mathcal{G}_{RoI} = \mathcal{M}_{Lift}(\mathcal{I}_{RoI}, \mathcal{O}), \quad (11)$$

where  $\mathcal{O} = \{\mathcal{G}_{RoI}^{add}, \mathcal{G}_{RoI}^{del}, \mathcal{B}_{3D}\}$  is optional instructions.

### 3.5. Delicate Editing within Gaussian RoI

To achieve delicate editing in 3D scenes, we use the Gaussian RoI to constrain the editing area. In particular, we randomly sample viewpoints from the 3D scene and render 2D image  $\mathcal{I}_{render}$ . After that,  $\mathcal{I}_{render}$  and noise level  $t$  are put into 2D diffusion model  $\mathcal{M}_{DM}$ , with the user instruction  $\mathcal{T}$  and image  $\mathcal{I}_{input}$  of input scene as conditions, to get edited image  $\mathcal{I}_{edit}$ :

$$\mathcal{I}_{edit} = \mathcal{D}(\mathcal{I}_{render}, t; \mathcal{T}, \mathcal{I}_{input}), \quad (12)$$

where  $t$  is a randomly chosen noise level from  $[t_{min}, t_{max}]$ .

Similar to 3D-GS [18], we apply the  $\mathcal{L}_1$  and D-SSIM loss functions during editing.

$$\mathcal{L} = (1 - \beta)\mathcal{L}_1 + \beta\mathcal{L}_{D-SSIM}. \quad (13)$$

the two losses are calculated between the 2D edited image  $\mathcal{I}_{ed}$  and the rendered image  $\mathcal{I}_{rd}$ . Then, gradient backpropagation is performed within Gaussian RoI  $\mathcal{G}_{RoI}$ :

$$\nabla \mathcal{G} = \frac{\partial \mathcal{L}}{\partial \mathcal{G}} \cdot \mathcal{R}, \quad (14)$$

where  $\mathcal{R}$  is the set of RoI attributes. That means only Gaussians in RoI can receive gradients. Finally, we utilize the Adam algorithm to optimize the 3D Gaussians. After many rounds of training, the edited 3D scene  $\mathcal{G}_{edit}$  is obtained.

## 4. Experiments

### 4.1. Implementation Details

Our method is implemented in PyTorch [35] and CUDA, based on 3D Gaussian splatting. The multimodal model we used in our method is BLIP2 [21], and we use GPT-3.5 Turbo to ground the text RoI. For grounding segmentation, We use the cascade strategy, *i.e.* first using Grounding DINO [26] to get the box on the image corresponding to the text, and then using SAM [19] to get the corresponding image RoI. The 2D diffusion model used in our method is Instruct Pix2Pix [4]. We leave more details in the Appendix.

### 4.2. Qualitative Evaluation

**Visualization Results.** In Fig. 1 and Fig. 4, we present the visual results of GaussianEditor, demonstrating the precise editing effects while ensuring 3D consistency. Fig. 1 shows the editing capabilities for characters. The first column displays the original scenes. In the second column, the first row “Give him a red nose” illustrates

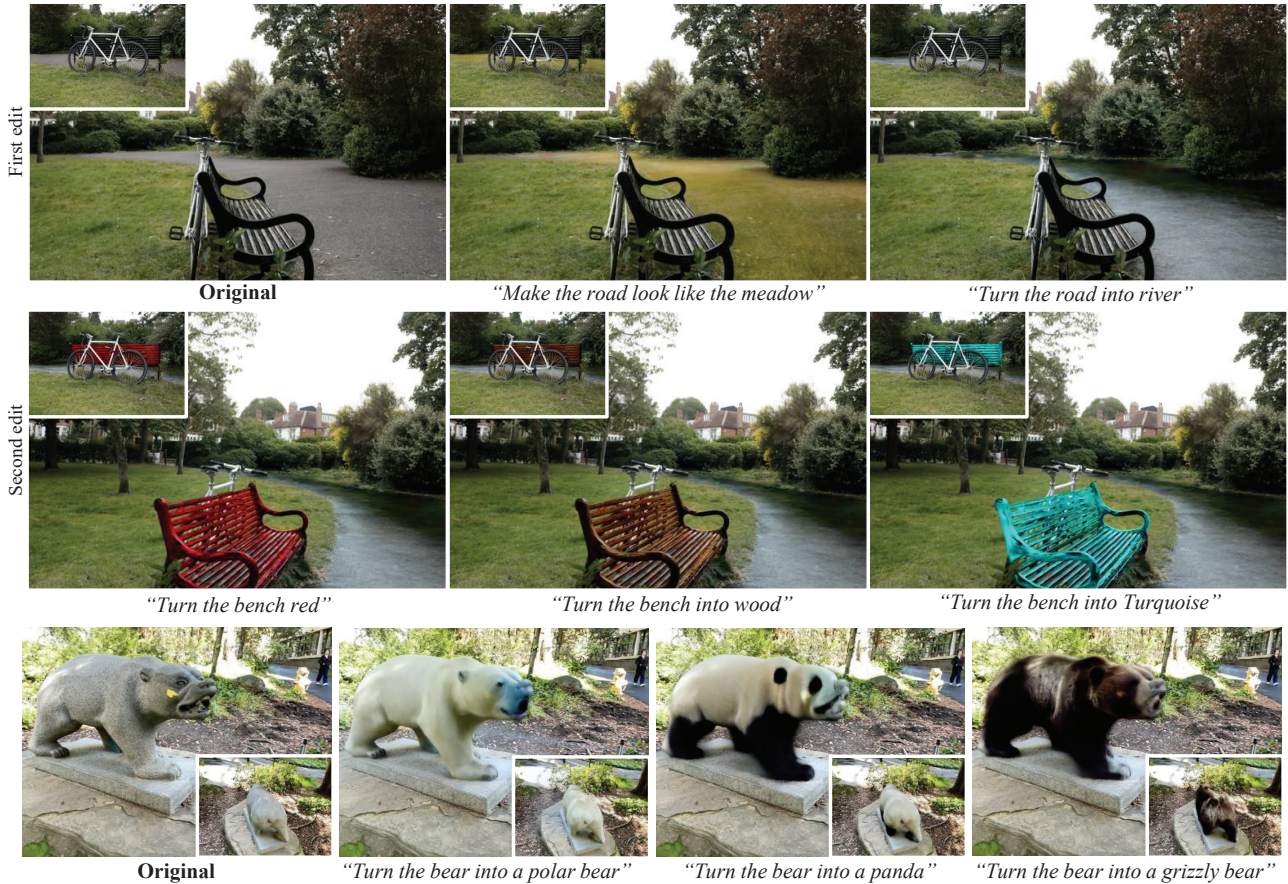


Figure 4. Qualitative results on outdoor scenes. Our method supports separate foreground and background editing in real-world scenes.

color-changing ability, while the third row, “Make him completely bald”, showcases capabilities of retexturing and slight geometry editing. The second row in the second column demonstrates precise editing ability by exclusively editing the left side of the face. Based on that, we achieve editing in the third column, focusing on the right side of the face, showcasing the ability of multi-round edits, and accurately fulfilling user instructions. Fig. 4 showcases the precise editing capabilities in open 3D scenes. In the upper portion, the bicycle scene allows us to accurately locate the position of the road and edit its texture, transforming it into the grass, a river. In the experiment where we change the texture to a river, our method accurately constructs the reflection, making it appear realistic. Based on editing the road into a river, we further edited the bench, proving that our method can achieve multiple rounds of editing. The lower portion demonstrates the results of editing the bear, which fully preserves the original appearance of the background area and focuses the edits on the bear.

**Comparisons with Instruct-NeRF2NeRF.** Fig. 5 compares the results of our method with those of IN2N [11], on the scenes presented in IN2N. From the figure, it is evident that our method changes the texture of the pants without affecting the clothes, and vice versa, demonstrating

the effectiveness of our method in distinguishing different objects within the foreground. Additionally, when editing the clothes and pants, the background remains unaffected, indicating our method’s effective separation of foreground and background. Furthermore, the last column reveals that IN2N, limited by 2D diffusion, distorts the face, while our method maintains a superior rendering quality of faces.

**Complex Multi-Object Scenes.** Furthermore, we present the results of our editing in a complex scene featuring multiple objects, as depicted in Fig 6. We specifically selected three distinct object types for editing purposes. The first object type is the background. In a particular sub-scene, the desktop shares similarities with the background. Employing a caption-based approach, we successfully transformed the desktop into a wooden material. As evident in the image, the edited result exhibits a distinct and pronounced wood texture. The second object type is a foreground object, the flowerpot. We opted to change the color of the flowerpot to red, and the outcome was highly successful. Lastly, the most intricate editing task involved the rolling pin, which was occluded by multiple objects from various perspectives. As illustrated in the lower right corner of the picture, we managed to edit it into a cucumber without any adverse impact on the other objects.



Figure 5. Comparisons with Instruct-NeRF2NeRF (IN2N) [11] on the scene presented in their paper.



Figure 6. Qualitative results on complex multi-object scenes. The background “desk”, the foreground “flower pot”, and the multi-view blocked foreground “rolling pin” are edited separately.

Table 1. Quantitative evaluation on the bicycle scene of the Mip-NeRF360 dataset [3].

Methods	CTIDS $\uparrow$	IIS $\uparrow$	FID $\downarrow$	Time $\downarrow$
IN2N [11]	0.22	0.85	103	51 min
Ours-DVGO [45]	0.11	0.82	148	40 min
Ours-3DGS	<b>0.28</b>	<b>0.95</b>	<b>51</b>	<b>20 min</b>

### 4.3. Quantitative Evaluation

**Metric Comparisons.** Table 1 shows quantitative results on the bicycle scene of the Mip-NeRF360 dataset [3], comparing with IN2N [11] and Direct Voxel Grid Optimization (DVGO) [45] as the representation. The metrics include CLIP Text-Image Direction Similarity (CTIDS), Image-Image Similarity (IIS), FID, and training time. GaussianEditor achieves the best results in all metrics. The test data is shown in the supplementary material.

**User Study.** We perform a user study comparing with IN2N on the bear scene in Fig. 4 and the human scene in Fig. 5, involving 21 participants. GaussianEditor gets a

87.07% voting percentage, while IN2N gets 12.93%.

### 4.4. Ablation Study and Analysis

**Ablation of Gaussian RoI, Text RoI, RoI Lifting.** To validate the effectiveness of each module in our framework, we design three variant approaches: (1) w/o Gaussian RoI: We discontinued the use of Gaussian RoI  $\mathcal{G}_{RoI}$  to control the gradients of Gaussian points, as mentioned in Eq. 14. (2) w/o Text ROI: In this scenario, we ceased the selection of text ROI  $\mathcal{T}_{RoI}$  using LLM assistant  $\mathcal{M}_{LLM}$ . Instead, all the words in the user’s instruction are put to  $\mathcal{M}_{Seg}$  to get  $\mathcal{I}_{RoI}$ . (3) w/o RoI lifting: Instead of lifting the image RoI  $\mathcal{I}_{RoI}$  to 3D Gaussians, the image RoI  $\mathcal{I}_{RoI}$  is used to govern the calculation of the loss. That is, only the pixels within the image RoI  $\mathcal{I}_{RoI}$  are taken into account for the loss computation. Fig. 7 showcases the outcomes of our ablation experiment, which aimed to edit the doll based on the instruction “Turn its mouth into red.” The results reveal the following findings. (1) When the Gaussian RoI is not used, the 3D scene is all turned red because the 2D diffusion fails to control the editing area. (2) In cases where

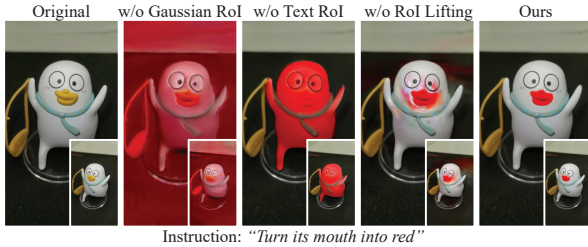


Figure 7. Ablation experiment of RoI.

text ROI  $\mathcal{T}_{RoI}$  is not utilized, the grounding segmentation model tends to segment the entire foreground object, leading to the doll being entirely edited to red. (3) When ROI lifting  $\mathcal{M}_{Lift}$  is not employed, the doll’s mouth is successfully turned red, but other facial areas are also affected. Because the grounding segmentation model may fail to parse specific views, noise exists on the image ROI  $\mathcal{I}_{RoI}$ . Consequently, leakage occurs during the editing process. Our proposed ROI lifting module effectively addresses this issue during training. In conclusion, our ablation experiment demonstrates the effectiveness of several ROI-related modules in our method.

**Ablation of Scene Description Generation.** We further conduct experiments to evaluate the role of scene description generation, employing three distinct experimental setups. The first one composes the user message  $\mathcal{T}_{user}$ , without employing scene description. The second method randomly samples a view and extracts the corresponding image’s text description as the scene description. The third one represents the complete version of our approach. The test scene involves a park where a bike and a bench are positioned closely together. The editing instruction is “Turn the thing next to the bike orange”. The obtained results are presented in Fig. 8. As shown in the image, when scene description is not employed, the LLM fails to acquire the Text ROI according to the user instructions, resulting in editing failure. The second one randomly samples images to obtain scene descriptions, resulting in incomplete descriptions and leading to an incorrect text ROI prediction by the LLM. Consequently, the final editing result turns the road into an orange color. In contrast, our method flawlessly executes the editing task. This success can be attributed to scene description generation, which obtains an accurate text description encompassing the relative positional relationship between the bicycle and the bench. This enables the LLM to analyze and determine the user’s intention to edit the bench. Consequently, the desired color change of the bench is successfully implemented.

#### 4.5. Limitations

Although our framework has solved some problems inherited from the integrated sub-modules, *e.g.* noise in the results of grounding segmentation, there are still some prob-

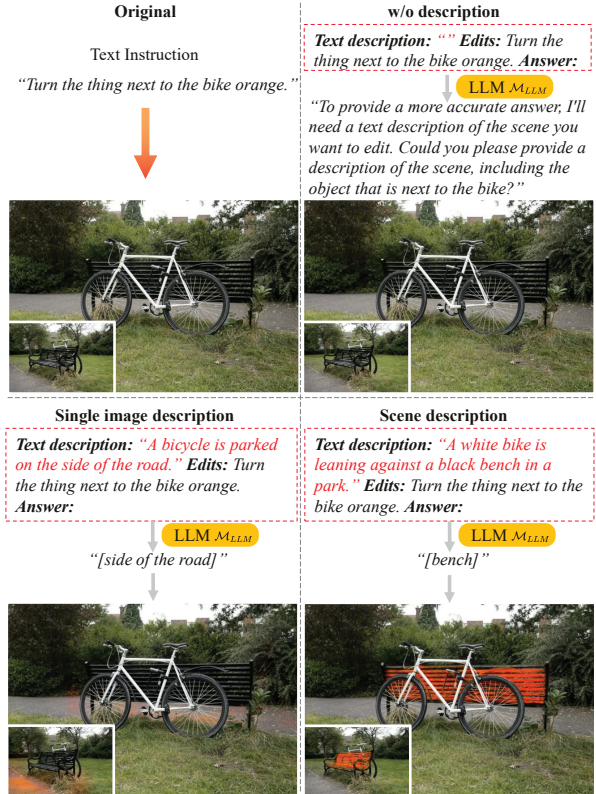


Figure 8. Ablation results about the scene description generation.

lems that the current system cannot completely avoid. In scene description generation, the descriptions from different views of the same object may differ from each other. When the differences are large enough, the LLM may misunderstand these descriptions as those from multiple objects. This issue does not affect the results in the current experiment, but we would like to optimize this in the future. In addition, our system cannot achieve good editing results in scenes where the grounding segmentation or diffusion model completely fails, such as drastic geometric editing.

## 5. Conclusion

This paper proposes a systematic framework, named GaussianEditor, for text-guided delicate 3D scene editing. As we know, GaussianEditor is one of the first works to edit 3D Gaussians, taking advantage of the explicit property of 3D Gaussians and making it easy to control the editing area precisely. Several techniques are proposed to achieve delicate editing, including extracting instruction ROI from texts, aligning the ROI to 3D Gaussians, and editing the scene with the Gaussian ROI. GaussianEditor achieves notably more delicate editing results than IN2N [11] with much shorter training time (within 20 minutes *v.s.* 45 minutes – 2 hours). Noticing recent works [27, 49] have extended Gaussian splatting to dynamic scenes, we leave the delicate editing in dynamic scenes as future work.



## References

- [1] Chong Bao, Yinda Zhang, Bangbang Yang, Tianxing Fan, Zesong Yang, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Sine: Semantic-driven nerf editing with prior-guided editing field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 2
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, 2021. 2
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, 2022. 2, 7
- [4] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023. 2, 3, 5
- [5] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Chen Yang, Wei Shen, Lingxi Xie, Dongsheng Jiang, Xiaopeng Zhang, and Qi Tian. Segment anything in 3d with nerfs. In *NeurIPS*, 2023. 5
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *ECCV*, 2022. 2
- [7] 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*, 2022. 2
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 2021. 2
- [9] Shuangkang Fang, Yufeng Wang, Yi Yang, Yi Yang, Yi-Hsuan Tsai, Wenrui Ding, Ming-Hsuan Yang, and Shuchang Zhou. Text-driven editing of 3d scenes without retraining. *Arxiv preprint arXiv:2309.04917*, 2023. 2
- [10] William Gao, Noam Aigerman, Thibault Groueix, Vova Kim, and Rana Hanocka. Textdeformer: Geometry manipulation using text guidance. In *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. 2
- [11] Ayaan Haque, Matthew Tancik, Alexei Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *CVPR*, 2023. 2, 3, 6, 7, 8
- [12] 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. 2
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 2
- [14] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *The Journal of Machine Learning Research*, 2022. 2
- [15] Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. Avatarclip: Zero-shot text-driven generation and animation of 3d avatars. *arXiv preprint arXiv:2205.08535*, 2022. 2
- [16] Hsin-Ping Huang, Hung-Yu Tseng, Saurabh Saini, Maneesh Singh, and Ming-Hsuan Yang. Learning to stylize novel views. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [17] 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 *CVPR*, 2022. 2
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)*, 2023. 2, 3, 5
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 5
- [20] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. *Advances in Neural Information Processing Systems*, 2022. 2
- [21] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 5
- [22] Yuan Li, Zhi-Hao Lin, David Forsyth, Jia-Bin Huang, and Shenlong Wang. Climatenerf: Physically-based neural rendering for extreme climate synthesis. *arXiv e-prints*, 2022. 2
- [23] Yuhan Li, Yishun Dou, Yue Shi, Yu Lei, Xuanhong Chen, Yi Zhang, Peng Zhou, and Bingbing Ni. Focaldreamer: Text-driven 3d editing via focal-fusion assembly. *arXiv preprint arXiv:2308.10608*, 2023. 3
- [24] Hao-Kang Liu, I Shen, Bing-Yu Chen, et al. Nerf-in: Free-form nerf inpainting with rgb-d priors. *arXiv preprint arXiv:2206.04901*, 2022.
- [25] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2021. 2
- [26] 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. 5
- [27] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023. 8
- [28] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2mesh: Text-driven neural stylization for meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [29] 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

- [30] Ashkan Mirzaei, Tristan Aumentado-Armstrong, Marcus A. Brubaker, Jonathan Kelly, Alex Levinshtein, Konstantinos G. Derpanis, and Igor Gilitschenski. Watch your steps: Local image and scene editing by text instructions. In *arXiv preprint arXiv:2308.08947*, 2023. 2, 3
- [31] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics (ToG)*, 2022. 2
- [32] Thu Nguyen-Phuoc, Feng Liu, and Lei Xiao. Snerf: stylized neural implicit representations for 3d scenes. *arXiv preprint arXiv:2207.02363*, 2022. 2
- [33] 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. *arXiv preprint arXiv:2112.10741*, 2021. 2
- [34] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [35] 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. 5
- [36] 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
- [37] 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*, 2022. 2
- [38] 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. 2
- [39] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, 2022. 2
- [40] 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. *Advances in Neural Information Processing Systems*, 2022. 2
- [41] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2
- [42] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 2
- [43] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2015. 2
- [44] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 2019. 2
- [45] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022. 2, 7
- [46] Vadim Tschernezki, Iro Laina, Diane Larlus, and Andrea Vedaldi. Neural feature fusion fields: 3d distillation of self-supervised 2d image representations. In *2022 International Conference on 3D Vision (3DV)*, 2022. 2
- [47] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [48] Can Wang, Ruixiang Jiang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Nerf-art: Text-driven neural radiance fields stylization. *TVCG*, 2023. 2
- [49] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Wang Xinggang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. 8
- [50] Qiling Wu, Jianchao Tan, and Kun Xu. Palettenerf: Palette-based color editing for nerfs. *arXiv preprint arXiv:2212.12871*, 2022. 2
- [51] Jiale Xu, Xintao Wang, Yan-Pei Cao, Weihao Cheng, Ying Shan, and Shenghua Gao. Instructp2p: Learning to edit 3d point clouds with text instructions. *arXiv preprint arXiv:2306.07154*, 2023. 2
- [52] Tianhan Xu and Tatsuya Harada. Deforming radiance fields with cages. In *European Conference on Computer Vision*, 2022.
- [53] 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 *European Conference on Computer Vision*, 2022. 2
- [54] Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields. In *European Conference on Computer Vision*, 2022. 2
- [55] 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. 3