

# GeneAvatar: Generic Expression-Aware Volumetric Head Avatar Editing from a Single Image

Chong Bao<sup>1\*§</sup> Yinda Zhang<sup>2\*</sup> Yuan Li<sup>1\*</sup> Xiyu Zhang<sup>1</sup> Bangbang Yang<sup>4</sup>  
Hujun Bao<sup>1</sup> Marc Pollefeys<sup>3</sup> Guofeng Zhang<sup>1</sup> Zhaopeng Cui<sup>1†</sup>  
<sup>1</sup>State Key Lab of CAD&CG, Zhejiang University <sup>2</sup>Google <sup>3</sup>ETH Zürich <sup>4</sup>ByteDance

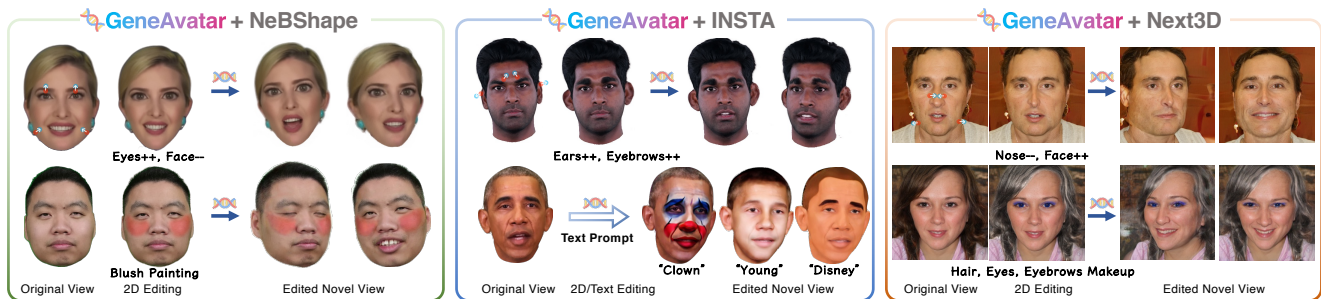


Figure 1. We propose a generic approach to edit 3D avatars in various volumetric representations (NeRFBlendShape [16], INSTA [76], Next3D [50]) from a single perspective using 2D editing methods with drag-style, text-prompt and pattern painting. Our editing results are consistent across multiple facial expression and camera viewpoints.

## Abstract

Recently, we have witnessed the explosive growth of various volumetric representations in modeling animatable head avatars. However, due to the diversity of frameworks, there is no practical method to support high-level applications like 3D head avatar editing across different representations. In this paper, we propose a generic avatar editing approach that can be universally applied to various 3DMM-driving volumetric head avatars. To achieve this goal, we design a novel expression-aware modification generative model, which enables lift 2D editing from a single image to a consistent 3D modification field. To ensure the effectiveness of the generative modification process, we develop several techniques, including an expression-dependent modification distillation scheme to draw knowledge from the large-scale head avatar model and 2D facial texture editing tools, implicit latent space guidance to enhance model convergence, and a segmentation-based loss reweight strategy for fine-grained texture inversion. Extensive experiments demonstrate that our method delivers high-quality and consistent results across multiple expression and viewpoints. Project page: <https://zju3dv.github.io/geneavatar/>.

## 1. Introduction

Recently various volumetric representations [3, 4, 15, 16, 57, 68, 69, 76] have achieved remarkable success in reconstructing personalized, animatable, and photorealistic head avatars using implicit [15, 16, 57, 68, 69] or explicit [3, 4, 76] conditioning of 3D Morphable Models

(3DMM) [6]. A popular demand, once with a created avatar model, is to edit the avatar, e.g., for face shape, facial makeup, or apply artistic effects, for the downstream applications, e.g., in virtual/augmented reality.

Ideally, the desired editing functionality on the animatable avatar should have the following properties. (1) **Adaptable**: The editing method should be applicable across various volumetric avatar representations. This is particularly valuable in light of the growing diversity of avatar frameworks [16, 50, 76]. (2) **User-friendly**: The editing should be user-friendly and intuitive. Preferably, the editing of geometry and texture of the 3D avatar could be accomplished on a single-perspective rendered image. (3) **Faithful**: The editing results should be consistent across various facial expression and camera viewpoints. (4) **Flexible**: Both intensive editing (e.g., global appearance transfer following style prompts) and delicate local editing (e.g., dragging to enlarge eyes or ears) should be supported as illustrated in Fig. 1.

However, 3D-aware avatar editing is still underexplored in both geometry and texture. One plausible way is to perform 3D editing via animatable 3D GAN [50, 52, 54], but the editing results may not be consistently reflected when expression and camera viewpoint change. Alternatively, the editing can be done on the generated 2D video using 2D personalized StyleGAN [28]; however, the identity shift is often observed. Some face-swapping methods [11, 12, 42] are capable of substituting the face in a video with another face derived from a reference image or video; however, they do not support texture editing and local geometry editing.

To this end, we propose GeneAvatar – a generic approach to support fine-grained 3D editing in various volumetric avatar representations from a single perspective by

\* Authors contributed equally.

† Corresponding authors.

§ The work was partially done when visiting ETHZ.

leveraging 2D editing methods, such as drag-based methods [30, 34, 39, 47], text-driven methods [7, 17, 20, 40, 41], or image editing tools like Photoshop (see Fig. 1). We adopt a novel editing framework that formulates the editing as predicting expression-aware 3D modification fields applied in the geometry and texture space of the volumetric avatars, which makes editing independent with the original representation as long as they are in parametric-driven radiance field, e.g., 3DMM-based neural avatar. Second, to ensure that the 2D image editing can be faithfully transferred into the 3D space, we propose to learn a generative model for modification fields, which produces 3DMM conditional modification fields from a compact latent space. Given the rendered avatar image and its edited counterpart, we conduct auto-decoding optimization on this generative model to search for the latent code that best explains the editing, obtaining consistent 3DMM conditional modification fields across various viewpoints and expression. Third, inspired by the spirit of learning from the pre-trained large-scale generative model [7, 18, 32, 74], we design a novel distillation scheme to learn the expression-dependent modification from a 3DMM-based GAN [50] and 2D face editing tools [22, 27, 36]. The scheme addresses the issue of insufficient real training data (i.e., avatars with a wide range of geometry and texture changes). Besides, we develop several techniques to enhance the editing effects, including the implicit latent space guidance to stabilize the initialization and convergence of learning, and a segmentation-based loss reweight strategy for fine-grained texture inversion.

The contributions of our paper are summarized as follows. **1)** We propose a generic avatar editing approach that can be applied to various 3DMM driving head avatars in the neural radiance field. To achieve this, we design a novel expression-aware modification generative model, which lifts the geometry and texture editing from a single image to a consistent 3D modification field. **2)** To bootstrap the training of the modification generator with limited real paired training data, we design a distillation scheme to learn the expression-dependent geometry and texture modification from the large-scale head avatar generative model [50] and 2D face texture editing tools [22, 27, 36], and develop several techniques, including implicit guidance in latent space to improve training convergence, and a loss reweight strategy based on segmentation for fine-grained texture inversion. **3)** Extensive experiments on various head avatar representations demonstrate that our method delivers high-quality editing results and the editing effects are consistent under different viewpoints and expression.

## 2. Related Work

**2D Head Avatar Editing.** The manipulation of 2D head avatars has made significant strides in recent years. Various GAN-based methods [23–25] can result in precise and high-resolution human face editing by leveraging image space se-

mantic information [29, 72] or controlling latent space explorations [10, 19, 46, 70]. Some approaches [22, 27, 36, 61] focus on the task of makeup transferring by exploiting the GAN to learn the transferring ability from a large unaligned makeup and non-makeup face datasets. The drag-based GAN editing approach [39, 71] gained vast popularity due to providing a user-friendly editing way. PVP [28] uses a monocular video to fine-tune StyleGAN [23, 24] to obtain the personalized image generator and provide various editing functions. The diffusion models [7, 21] also show the capability of achieving fine-grained face editing with text prompt and other conditional input. For editing an avatar in a video, lots of face-swapping methods [11, 12, 42] have emerged to provide high-quality and properly aligned face-swapping results. However, these approaches typically suffer from multi-view consistency and identity preservation.

**3D Head Avatar Editing.** Neural Radiance Field [33] has exhibited great reconstruction and rendering qualities in SLAM [62, 73], scene editing [5, 58–60, 64] and re-lighting [63, 66, 67], especially promoting the emergence of many 3D avatar reconstruction [4, 16, 53, 68, 69, 76] and generation [50, 52, 54]. Some methods [2, 49, 56, 65] exploit the powerful editing ability of GAN to edit a 3D static head portrait. However, they cannot be trivially extended to the dynamic avatars. The methods [37, 38, 45] focus on style transfer of the avatar using text prompt or style image but reach a poor identity-preserving. We propose a novel 3D avatar editing approach with an expression-aware modification generative model, which can be applied to various 3DMM-based volumetric avatars and render consistent novel views with fine-grained editing across multiple viewpoints and expression while preserving identity of person.

## 3. Method

As shown in Fig. 2, given a volumetric head avatar, we edit the avatar using a single-view image and synthesize consistent novel views across multiple expression and viewpoints. To achieve this goal, we propose a novel expression-aware modification generator to generate 3D modification fields, which can be seamlessly integrated into various representations and animated with facial expression (see Sec. 3.2). Furthermore, to bootstrap the training with limited pairwise data, we propose a novel expression-aware distillation scheme to learn the expression-dependent modifications from large-scale generative models [7, 50] (see Sec. 3.3). During the editing process, given a single edited image of a 3D avatar, we perform an auto-decoding optimization to lift 2D editing effect to the 3D space (see Sec. 3.4).

### 3.1. Preliminaries

The current implicit volumetric representations of head avatar [16, 50, 76] are mostly built upon the NeRF [33] or its variants [8, 9, 31, 35, 44, 48]. In general, the neural architecture can be simplified as an implicit field  $\mathbf{F}$  that

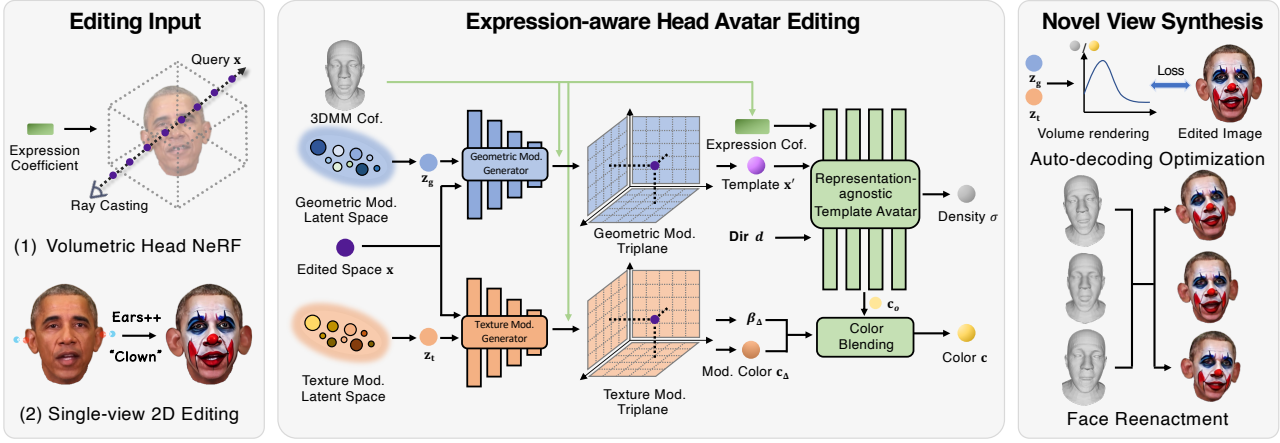


Figure 2. We use an expression-aware generative model that accepts a modification latent code  $\mathbf{z}_{g/t}$  and 3DMM coefficients and outputs a modification field of a tri-plane structure. The modification field modifies the geometry and texture of the template avatar by deforming the sample points  $\mathbf{x}$  and blending the color  $\mathbf{c}_o$  with the modification color  $\mathbf{c}_\Delta$  respectively. We lift the 2D editing effect to 3D using an auto-decoding optimization and synthesize novel views across different expression.

takes position  $\mathbf{x}$  and view direction  $\mathbf{d}$  as inputs and predicts the geometry  $\sigma$  and the texture  $\mathbf{c}$  of the avatar, i.e.,  $(\sigma, \mathbf{c}) = \mathbf{F}(\mathbf{x}, \mathbf{d})$ . Then, the volume rendering technique is used to render images as follows:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i \mathbf{c}_i, \quad T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma'_j \delta_j\right), \quad (1)$$

where  $\alpha_i = 1 - \exp(-\sigma'_i \delta_i)$ , and  $\delta_i$  is the distance between adjacent samples along the ray. In order to animate the head avatar, 3DMM [6] is incorporated to describe the deformation implicitly [15, 16, 57, 68, 69] or explicitly [4, 76].

### 3.2. Expression-aware Modification Generator

To enable the modification animated with the facial expressions, we follow the architecture of 3DMM-based 3D GAN [50] to build our expression-aware modification generator. As shown in Fig. 2, our generator consists of a geometry generator  $\mathbf{G}_{\Delta g}$  and a texture generator  $\mathbf{G}_{\Delta t}$ .  $\mathbf{G}_{\Delta g}$  encodes the expression-dependent geometry modification by deforming the query points in the edited space to the original template space under each expression.  $\mathbf{G}_{\Delta t}$  encodes the expression-dependent modification color of query points under each expression:

$$\mathbf{x}' = \mathbf{G}_{\Delta g}(\mathbf{x}, \mathbf{z}_g, \mathbf{v}), \quad (\mathbf{c}_\Delta, \beta_\Delta) = \mathbf{G}_{\Delta t}(\mathbf{x}, \mathbf{z}_t, \mathbf{v}), \quad (2)$$

where  $\mathbf{x}$  is the query points in the edited space, and  $\mathbf{x}'$  is the deformed point in the space of original avatar under current expression  $\mathbf{e}$ .  $\mathbf{c}_\Delta$  is the modification color and  $\beta_\Delta$  determines the blending weights with the original color.  $\mathbf{z}_g, \mathbf{z}_t$  are the geometry and texture modification latent code respectively, where  $\mathbf{z}_g, \mathbf{z}_t \in R^{1024}$ . They control the generation of modification feature maps in the UV space.  $\mathbf{v}$  is the 3DMM mesh vertices [26] that condition the current expression  $\mathbf{e}$ . Each mesh vertex has a neural feature that is retrieved in the modification feature map using pre-defined

UV mapping. We rasterize the vertex features to the three axis-aligned planes to generate the tri-plane feature. The modification information of the query point  $\mathbf{x}$  is first collected by bilinear interpolation on the tri-plane feature and then decoded by the neural feature decoder [8]. Since the modification is defined as decoupled fields without relying on the original field, our generated modification field can be integrated into various volumetric avatar representations and be animated following the facial expression.

### 3.3. Expression-dependent Modification Learning

To learn the proposed expression-aware modification, we need extensive training data on avatars with a wide range of geometry and texture changes, which is hard to obtain in practice. Following the spirit of learning high-fidelity editing ability from the large-scale generative model [7, 18, 32, 74], we propose a novel expression-aware distillation scheme to deal with insufficient real training data. We leverage the ability of 3DMM-based 3D GAN [50] and 2D face texture editing tools to generate facial editing data, which encompasses a wide range of geometry and texture editing across various expression and viewpoints.

**Geometry Distillation.** We use the teacher 3DMM-based 3D GAN [50]  $\mathbf{G}_n$  to synthesize two volumetric avatars with different geometry (an original avatar  $\mathbf{F}$  and an edited avatar  $\mathbf{F}'$ ) by modifying the 3DMM shape parameter of the original avatar. This provides the paired editing data for our generator to learn how to modify the geometry of the avatar while maintaining consistency across various expression and viewpoints. Specifically, we randomly sample the latent code  $\mathbf{z}_n$  in the latent space of  $\mathbf{G}_n$  as well as 3DMM shape parameter  $\beta$ , expression parameter  $\psi$ , and pose parameter  $\theta$ .  $\beta, \psi$  are sampled from a normal distribution whose absolute mean and standard deviation are within  $[0, 1]$ .  $\theta$  are a group of rotation vectors that have random directions within a unit sphere and magnitude within  $[-6, 6]$  degrees. Then, we sample an edit vector  $\beta_\Delta$  from a uniform distribution

$\mathcal{U}(-3, 3)$  and apply it to the original shape parameter by  $\beta' = \beta_\Delta + \beta$ . These hyperparameters w.r.t. 3DMM coefficients sampling are selected empirically to maintain the shape definition of the human head. Please refer to our supplementary Sec. B.2 for more details on 3DMM sampling. The original avatar  $\mathbf{F}$  and paired edited avatar  $\mathbf{F}'$  are generated by  $\mathbf{F} = \mathbf{G}_n(\mathbf{z}_n, \beta, \psi, \theta)$ ,  $\mathbf{F}' = \mathbf{G}(\mathbf{z}_n, \beta', \psi, \theta)$ . During training, we will apply our modification generator to modify the geometry of  $\mathbf{F}$  such that  $F$  and  $F'$  render the face with the same geometry.

**Texture Distillation.** We distill the capabilities of fine-grained texture editing from 2D face editing algorithms by generating texture-modified avatar  $\mathbf{F}'$  with the teacher 3DMM GAN [50]  $\mathbf{G}_n$ . Specifically, we sample an original avatar  $\mathbf{F} = \mathbf{G}_n(\mathbf{z}_n, \beta, \psi, \theta)$  from the teacher generator and render the image of its positive face. A segmentation-based 2D face texture editing algorithm (SBA) [77] and two makeup transfer algorithms (MTA) [22, 27, 36] are referred to in the distillation. We randomly choose one of them to edit the texture of the rendered face image. For SBA, we define several editable semantic regions of the face. A subset of these regions is selected randomly for texture painting using hues randomly sampled from the HSV color spectrum. For MTA, we randomly choose a makeup image as a reference from the open-sourced makeup dataset [22, 27, 36] and transfer the reference makeup to the rendered face image. The makeup dataset [36] contains complex makeups, such as blushes and makeup jewelry, which allow our generator to learn complicated texture editing patterns. Then, we perform the PTI inversion [43] on the texture-modified face image to lift the 2D texture editing to 3D space and obtain a texture-modified avatar  $\mathbf{F}'$ .

**Modification Learning.** Following the training style of StyleGAN [24], we sample a modification latent code  $\mathbf{z}_{g/t} \in R^{1024}$  in  $\mathcal{Z}$  latent space for each paired editing data. We do not fully sample a 1024-dimensional modification code but sample a reduced code  $\bar{\mathbf{z}}$  from a standard normal distribution and concatenate it with the latent code  $\mathbf{z}_n$  of the original avatar  $\mathbf{F}$  that is sampled from the teacher model, i.e.,  $\mathbf{z}_{g/t} = (\mathbf{z}_n, \bar{\mathbf{z}})$ ,  $\mathbf{z}_n, \bar{\mathbf{z}} \in R^{512}$ . This design is regarded as implicit code guidance that decently integrates knowledge from the teacher model to facilitate the model convergence.  $\mathbf{z}_n$  encodes the facial appearances of avatar  $\mathbf{F}$ , serving as a reference to the superimposition of the modification onto the avatar  $\mathbf{F}$ . Note that during inference, we do not require the concatenation of latent code from the teacher model and directly optimize the full modification code from the edited image using an auto-decoding manner. Our generator generates the modification field following Eq. (2) where  $\mathbf{v}$  is decoded from the 3DMM parameters of the avatar  $\mathbf{F}$  using FLAME model [26]  $\mathbf{E}$ , i.e.,  $\mathbf{v} = \mathbf{E}(\beta, \psi, \theta)$ . To apply the modification field, we feed the deformed query points  $\mathbf{x}'$  to the original avatar  $\mathbf{F}$  to

obtain the density and color, and composite the color with modification color by:

$$\mathbf{c} = (1 - \beta_\Delta) * \mathbf{c}_o + \beta_\Delta * \mathbf{c}_\Delta, \quad (\sigma, \mathbf{c}_o) = \mathbf{F}(\mathbf{x}', \mathbf{d}). \quad (3)$$

Then, we perform volume rendering on the density  $\sigma$  and color  $\mathbf{c}$  using Eq. (1) to render the modified image  $\hat{I}_e$  of avatar  $\mathbf{F}$ . We use the photometric loss to supervise the modified image with the rendered image  $\hat{I}'$  from the edited avatar  $\mathbf{F}'$  under the same camera parameters.

$$\mathcal{L} = \|\hat{I}_e - \hat{I}'\|_2^2. \quad (4)$$

During training, we sample multiple viewpoints and 3DMM expression parameters  $\psi$  for each editing pair  $(\mathbf{F}, \mathbf{F}')$  to enhance the spatial consistency under different expressions.

### 3.4. Avatar Editing with Single Image

In our task, users are allowed to edit a single image with various out-of-box face editing tools, such as Photoshop, drag-based editing [30, 39], text-driven editing [7]. For each editing input, we use the auto-decoding optimization on modification code to lift 2D edits into a 3D expression-aware modification field generated by our model. This field adapts to expression and viewpoint changes and is not tied to the specific avatar representation. The StyleGAN-based generator [24, 43, 55] features a latent space mapping from  $\mathbf{z} \in R^Z$  in  $\mathcal{Z}$  to  $\mathbf{w} \in R^{W_n \times W_d}$  in  $\mathcal{W}$ , where  $\mathbf{w}$  is more influential as it conditions the generator. Therefore, we perform code inversion in  $\mathcal{W}$  space by randomly sampling a modification code  $\mathbf{w}_{g/t}$  during editing. This code conditions a modification field  $\mathbf{G}_{\Delta g/t}(\mathbf{x}, \mathbf{w}_{g/t}, \mathbf{v})$  following Eq. (2). We apply the modification field to the original avatar using Eq. (4). The modified image  $\hat{I}_e$  is rendered following the original avatar’s rendering pipeline and is encouraged to match the user-edited image  $I_e$  by optimizing  $\mathbf{w}_{g/t}$  with the following loss terms:

$$\mathcal{L}_i = \lambda_1 \mathcal{L}_2(\hat{I}_e, I_e) + \lambda_2 \mathcal{L}_{\text{lpips}}(\hat{I}_e, I_e) + \lambda_3 \mathcal{L}_{\text{reg}}(\mathbf{w}, \mathbf{w}_{\text{avg}}), \quad (5)$$

The L2 loss term  $\mathcal{L}_2$  and LPIPS perceptual loss term  $\mathcal{L}_{\text{lpips}}$  encourage the rendered face close to the appearance and structure of the edited face. We set  $\lambda_1 = 1000$ ,  $\lambda_2 = 1$ ,  $\lambda_3 = 1$ . The regularization term  $\mathcal{L}_{\text{reg}}$  applied to the latent code  $\mathbf{w}$  enforces alignment with the distribution of the  $\mathcal{W}$  space, with  $\mathbf{w}_{\text{avg}}$  representing the mean latent code computed from 1000 random samples within the  $\mathcal{W}$  space. Besides, we observe that the L2 loss on the whole face will give an underfitting result for the fine-grained makeup on the eyes, eyebrows and lips. Therefore, we reweight the L2 loss on the facial features by face segmentation mask  $M = \{N_i | i = 1, \dots, m\}$  for texture editing:

$$\mathcal{L}_2 = \sum_{i=1}^m \sum_{r \in N_i} \frac{1}{|N_i|} \|\hat{C}_e(r) - C_e(r)\|_2^2, \quad (6)$$

Methods	Geometry				Texture			
	Roop	PVP	Next3D	Ours	Roop	PVP	Next3D	Ours
Editing preservation $\uparrow$	29.44%	23.33%	5.56%	<b>41.67%</b>	8.33%	10.56%	5.00%	<b>76.11%</b>
Identity preservation $\uparrow$	31.11%	21.11%	5.56%	<b>42.22%</b>	7.78%	12.78%	3.33%	<b>76.11%</b>
Temporal consistency $\uparrow$	30.00%	23.89%	4.44%	<b>41.67%</b>	5.56%	12.78%	1.67%	<b>80.00%</b>
Overall $\uparrow$	29.44%	23.33%	3.89%	<b>43.33%</b>	5.56%	12.78%	2.22%	<b>79.44%</b>
image identity similarity $\uparrow$	0.8373	0.8704	0.8547	<b>0.8845</b>	0.7320	0.8476	0.8500	<b>0.9147</b>

Table 1. We quantitatively compare with the PVP [28], Roop [12], Next3D [50] by user study and image identity similarity [13].

where  $\hat{C}_e(r)$ ,  $C_e(r)$  are the rendered and target color respectively,  $N_i$  are the rays within the  $i$ -th semantic part of the face. Generally, we freeze the weight of the modification generator and only optimize the modification latent code  $w$  to reach a satisfied 3D modification result. When an intense makeup or complicated pattern is painted onto the human face, we will continuously fine-tune the weight of our generator and freeze the latent code  $w$  to achieve more accurate editing results. To animate the edited avatar, users can input the new 3DMM expression parameter to the original avatar and our modification generator simultaneously. In this way, the generated modification field tightly sticks to the original avatar and presents reasonable editing results under different expression and viewpoints.

## 4. Experiments

In this section, we evaluate our avatar editing capability from a single perspective. One major difference with static NeRF editing is that we focus on showing how the edits are correctly lifted to 3D avatars under various expression and camera viewpoints.

### 4.1. Dataset and Baselines

**Datasets.** We use a total of 19 neural implicit head avatars from three methods, i.e., 7 from INSTA [76], 8 from NeRFBlendShape [16], 4 from Next3D [50], and show editing results on them. For INSTA [76] and NeRFBlendShape [16], we use the human head data (i.e., a monocular video of a head) provided by their methods to reconstruct the volumetric avatar using their respective representations. For Next3D [50], we random sample its latent space to generate volumetric avatars and perform editing on them. The evaluation datasets exhibit a substantial variation in identities, encompassing a diverse range of races, ages, and genders.

**Baselines.** We pick several baseline methods that can support single-view-based avatar editing. Roop [12] is a face-swapping method that can swap the human face in a video from a single reference view. To compare with Roop, we generate videos of the original avatar rendered in driving signals and single edited frames, and perform face swap. PVP [28] learns a personalized avatar image generator from a monocular video by fine-tuning the latent space of StyleGAN [24], and performs GAN-inversion style optimization [41] to edit the shape and appearance of the avatar. Next3D [50] is a 3DMM-based 3D GAN. To make a fair comparison on the same input data (i.e., a monocular video of avatar and an edited image), we perform GAN inver-

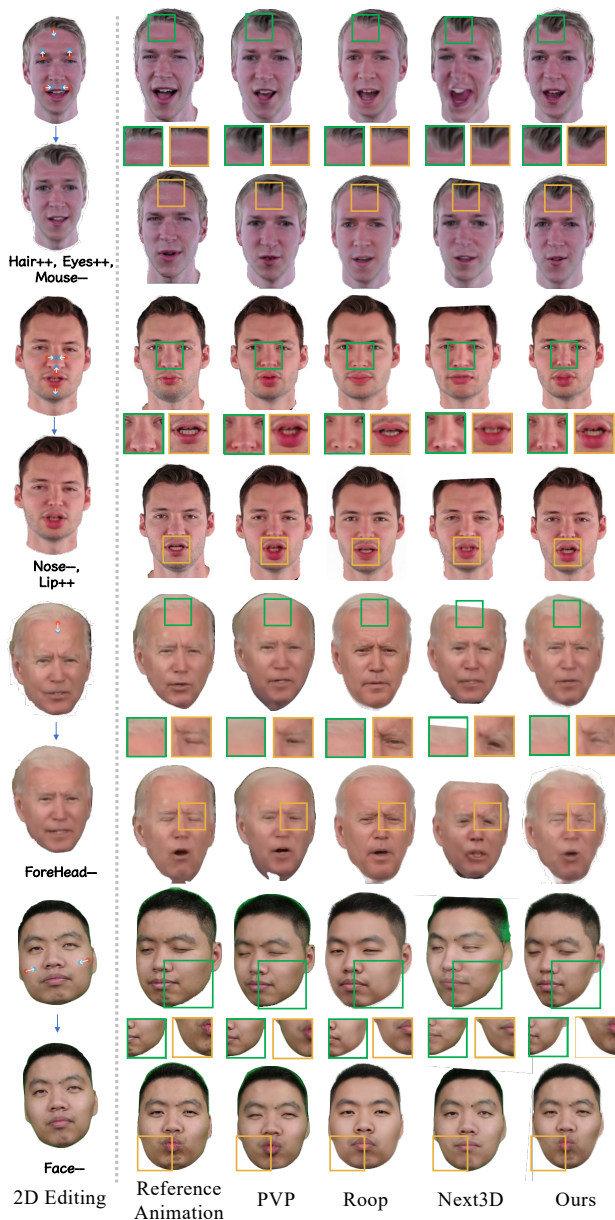
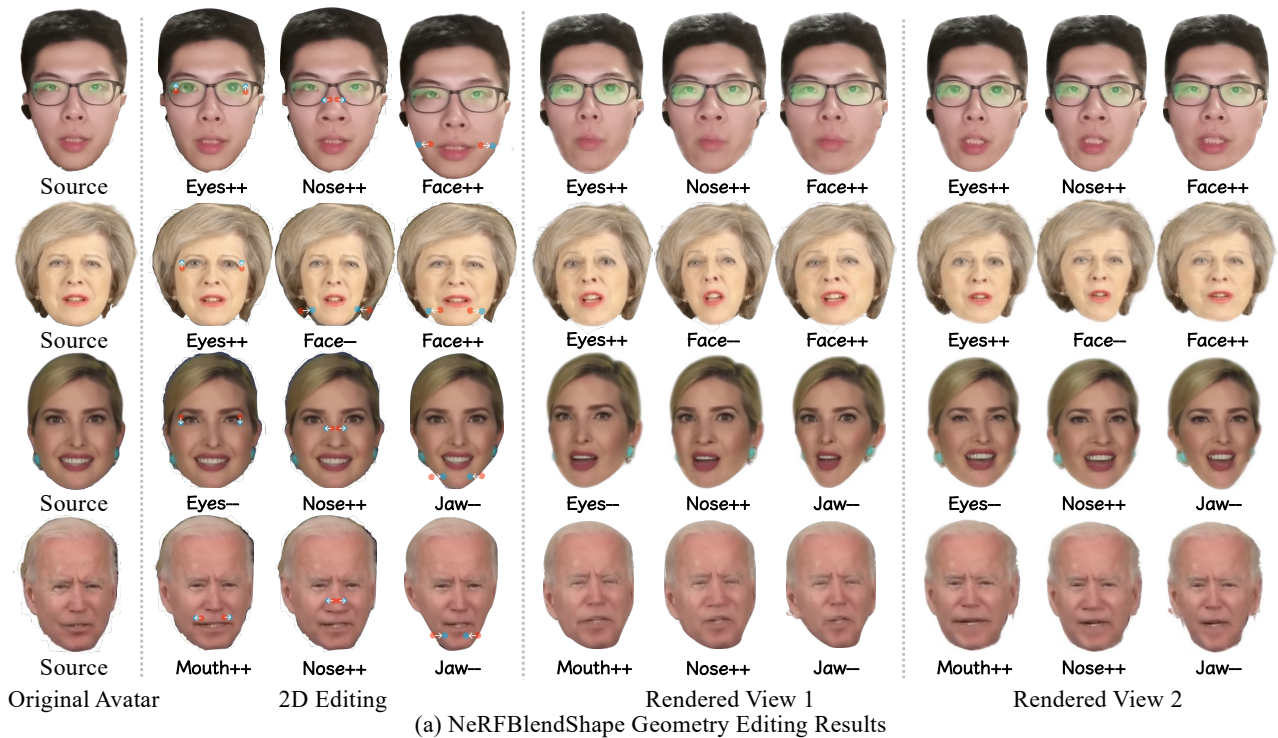


Figure 3. We compare geometry editing with PVP [28], Roop [12], Next3D [50] on INSTA [76] and NeRFBlendshape [16] avatars. The "Reference Animation" denotes the image of the original avatar under the same expression with the rendered edited view.

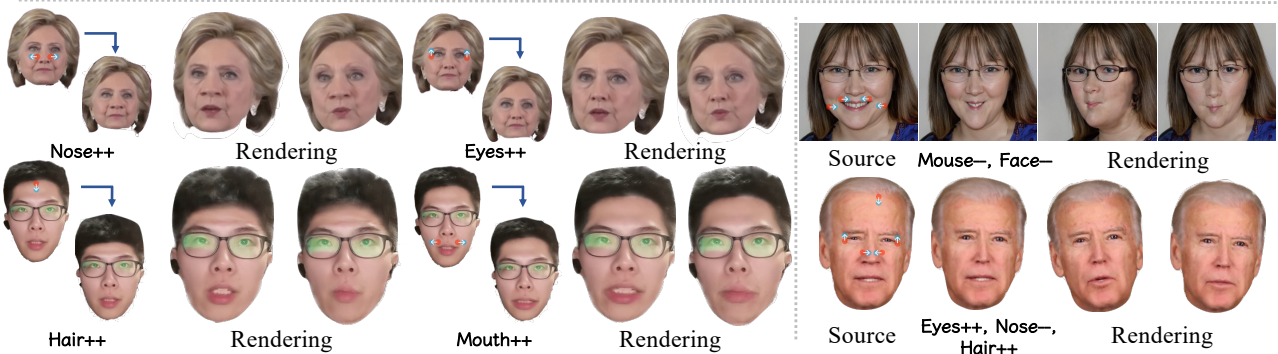
sion [43] with Next3D twice. First, we fine-tune Next3D with the input video to make it learn the original geometry and texture of the avatar. Second, we fine-tune Next3D on the edited image based on the weights of the code and generator from the first fine-tuning.

### 4.2. Qualitative Comparison

**Geometry Editing.** We first compare our method with baselines on geometry editing, e.g., changing the size of the eyes or the contour of the cheek. We use 2D editing tools, e.g., Photoshop and DragGAN [39], to modify the shape of various facial features. Figure 3 shows qualitative results on avatars from INSTA [76] (the top two) and



(a) NeRFBlendShape Geometry Editing Results



(b) NeRFBlendShape Geometry Editing Results

(c) INSTA & Next3D Geometry Editing Results

Figure 4. Our geometry editing results with the drag-style 2D editing on INSTA [76], NeRFBlendshape [16], and Next3D [50] avatars.

NeRFBlendshape [16] (the bottom two). Roop [12] fails to handle the fine-grained geometry change, like hairlines, and lips. Next3D [50] is able to successfully update the avatar based on the editing, however, changes the untouched part and causes an obvious identity shift. PVP [28] can make edits while preserving the identity, however, the magnitude of change tends to be smaller than the given image. In contrast, our method produces the desired editing effect from the edited image and preserves the multiview consistency and identity of the original face.

We further show more geometry editing results of our method on avatars in various representations in Fig. 4. Our method supports a convenient way to adjust the size of diverse facial features, such as eyes, mouths, jaw, etc., by editing a single rendered image from the avatars. The edits on 2D images are successfully lifted onto the avatar and rendered across different viewpoints and expression. Please refer to the supplementary Sec. C.3 for detailed visualiza-

tions of modified head geometry.

**Texture Editing.** We then show our capability in texture editing. We utilize Photoshop, an online makeup app WebBeauty [1], and text-driven editing method Instructpix2pix [7] to modify the texture on 2D renderings. We show comparisons with PVP [28], Roop [12] and Next3D [50] on four distinct heads avatars (INSTA [76] for the upper three and NeRFBlendshape [16] for bottom one) in Fig. 5. Roop [12] is ineffective in transferring non-human-face-like texture, thus failing in all examples. PVP [28] only transfers partial or blurry textures, and also causes shifts across the expression and head poses. Next3D [50] successfully uplifts the texture editing in 2D images sharply. However, it still suffers from the identity shift issue in the lower two heads and a blurred pattern in the upper two heads in Fig. 5. In contrast, our method faithfully paints the complicated texture following the edited image and preserves the identity of the original avatar and consis-

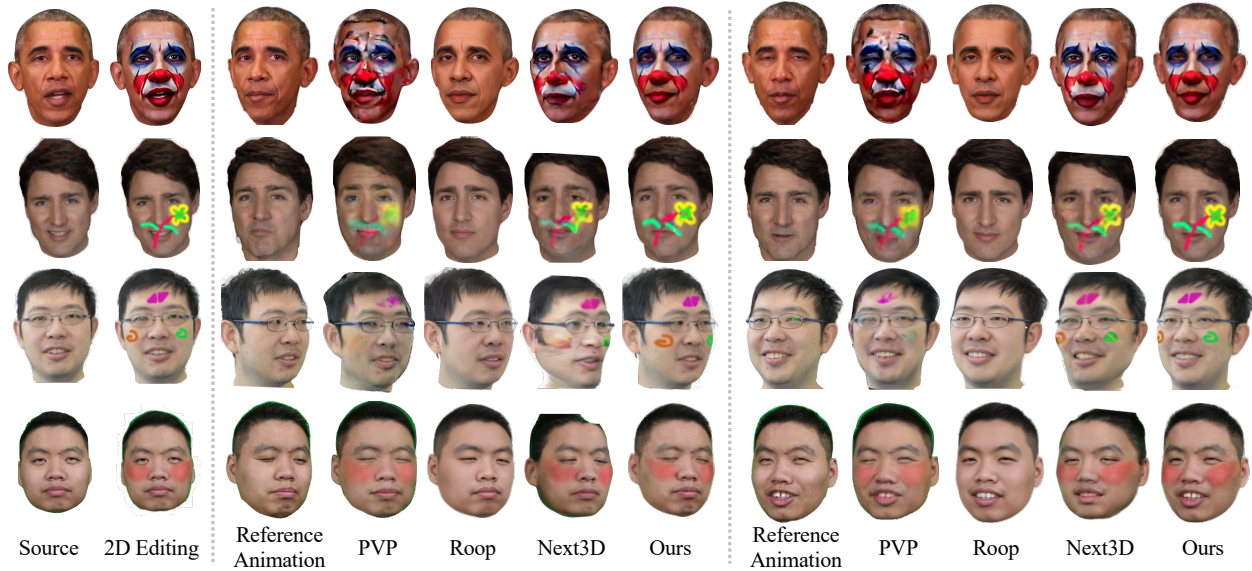


Figure 5. We compare texture editing with PVP [28], Roop [12], Next3D [50] on INSTA [76] and NeRFBlendshape [16] avatars. The “Reference Animation” denotes the image of the original avatar under the same expression with the rendered edited view.

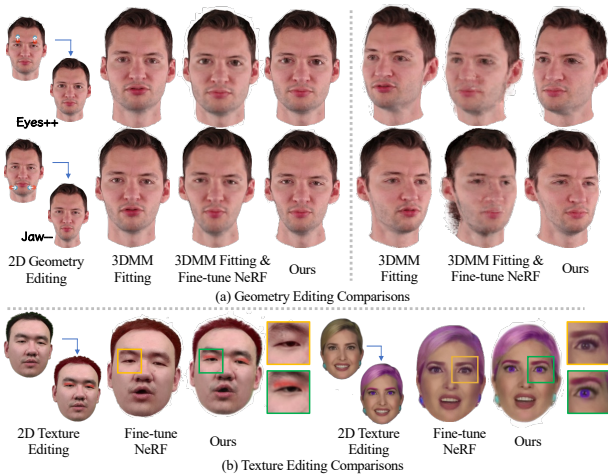


Figure 6. Analysis of our effectiveness with the naive baselines that can accomplish the single-view avatar editing.

tency across multiple viewpoints and expressions.

We show extensive texture editing results in three avatar representations in Fig. 7. Our method supports a wide range of texture editing, including global style transfer (“Add a clown makeup” via text-driven editing), semantic-driven editing (changing the hair color), and free-form sketch (painting on the face). Please refer to the supplementary Sec. C.1/2 for hybrid editing and face reenactment results.

### 4.3. Quantitative Comparison

**User Study.** We conducted a user study to validate our method further quantitatively. Following the evaluation protocol of Avatarstudio [38], users are required to watch the rendered videos of different methods side by side and answer each question by picking up one of the methods. For each group of editing results, we will ask four questions the

same as Avatarstudio [38] on editing preservation, identity preservation, temporal consistency and overall performance in Tab. 1. We collected statistics from 30 participants in 12 groups of edit results. Results are reported in Tab. 1. Our method exhibits the best editing ability while keeping the best consistency across different facial expressions for geometry and texture editing. Moreover, our method performs the best in keeping non-edited parts untouched and making the human heads still recognizable after being edited e.g., identity preservation metric in Tab. 1. Please refer to the supplementary Sec. B.4 for more details.

**Image Identity Similarity Evaluation.** Following the VoLux-GAN [51], we further evaluate the cross-view identity consistency in our edit results. We take 7 groups of geometry editing results and 7 groups of texture editing results, and render the 350 images of edited avatars with different viewpoints and expressions. We calculate the cosine similarities between each rendered image and the single-view 2D edited image and average the similarities on all rendered images as metrics. As reported in Tab. 1, our method outperforms all baselines in recovering desired editing effect and retaining identity consistency.

### 4.4. Native Editing Capability of Avatar

In this section, we analyze the editing capability from the original avatar model and show the necessity of our design.

**Comparison to 3DMM-based Geometry Editing.** One intuitive way to support geometry editing is updating the underlying 3DMM geometry [14]. Here, we investigate this method and verify its capability. Specifically, we run state-of-the-art single image-based 3DMM reconstruction method [75], and update the 3DMM shape parameter of the avatar model with the estimated one. Note that such an ap-

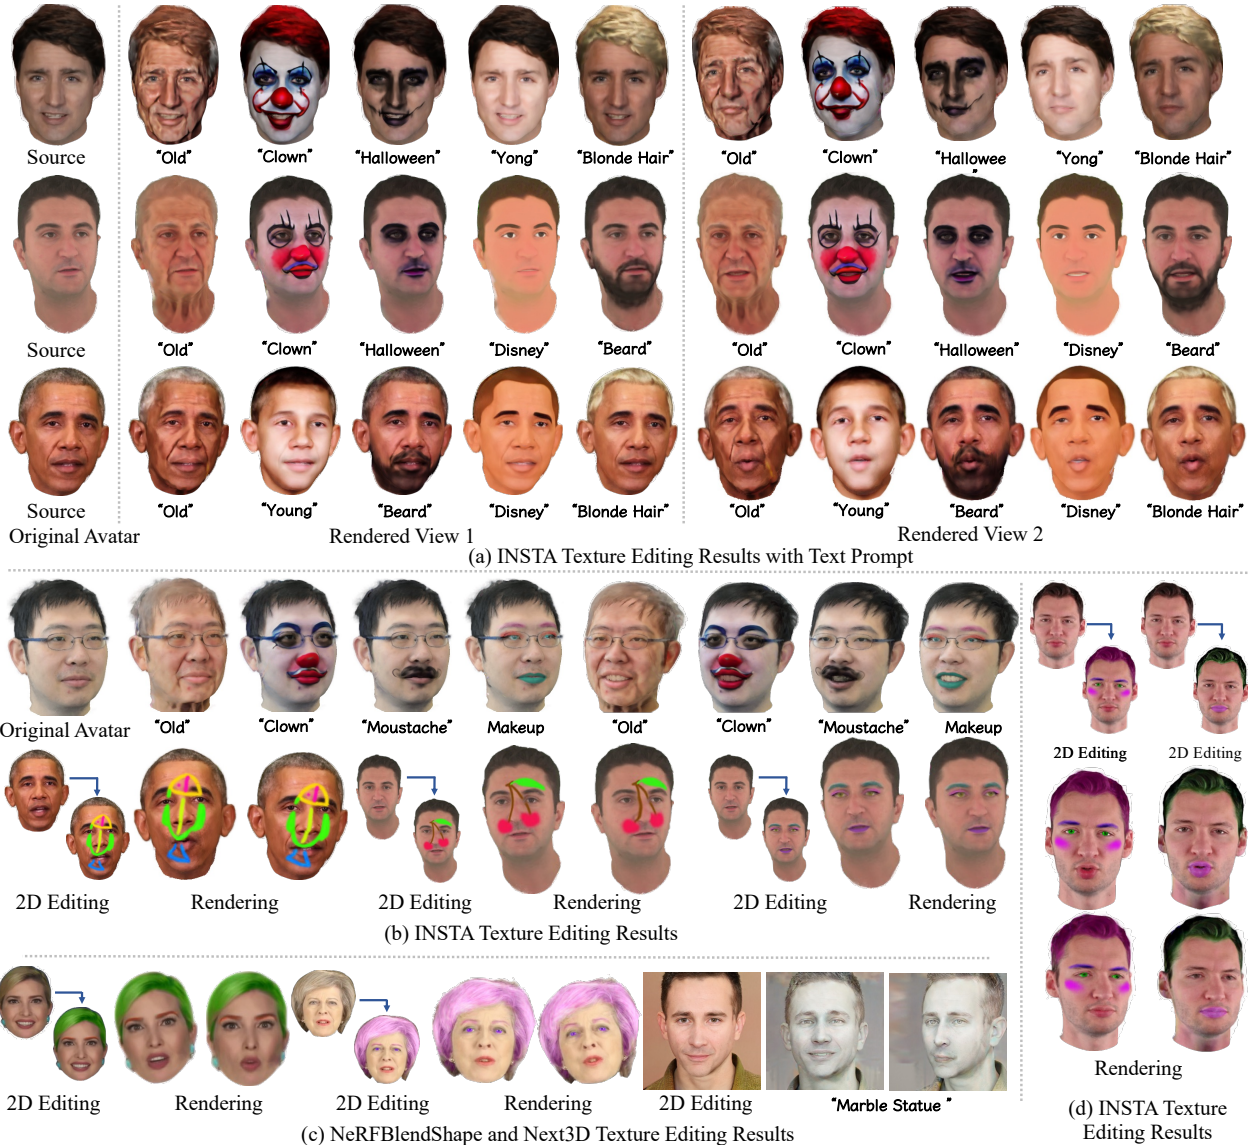


Figure 7. We show our texture editing results using the 2D editing method with text-prompt, pattern painting and makeup drawing on INSTA [76] and NeRFBlendShape [16], Next3D [50] avatars.

proach is only available for those avatars that use 3DMM explicitly, and we take INSTA [76] as an example and show the results in Fig. 6 (a). The 3DMM fitting tends to fail when the edited face is out-of-distribution, so the magnitude of editing may not be correct (e.g. the face shape). Changing the 3DMM parameter could also result in blurry rendering from the pre-trained avatar model, which cannot be trivially fixed even after fine-tuning the rendering decoder.

**Comparison to Fine-tuning for Texture Editing.** Texture editing could be done by fine-tuning the avatar with the one-shot edited image. We test this method on avatars from NeRFBlendShape [16] and show the results in Fig. 6 (b). While large structured changes, e.g., hair color, can be edited, detailed editing is largely ignored. Please refer to

the supplementary Sec. C.4 for more ablation studies.

## 5. Conclusion

We have proposed a novel generic editing approach that allows users to edit various volumetric head avatar representations from a single image, where an expression-aware modification generator lifts the editing to the 3D avatar while maintaining consistency across multiple expression and viewpoints. As a limitation, we cannot add additional objects (e.g., hat) or modify the hairstyle as shown in our supplementary Sec. C.5, which may be improved by learning extra specialized geometry addition and hair modification generators.

**Acknowledgment:** This work was partially supported by the NSFC (No. 62102356) and Ant Group.



## References

- [1] Webar/beauty demo app. [6](#)
- [2] Rameen Abdal, Hsin-Ying Lee, Peihao Zhu, Menglei Chai, Aliaksandr Siarohin, Peter Wonka, and Sergey Tulyakov. 3davatargan: Bridging domains for personalized editable avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4552–4562, 2023. [2](#)
- [3] ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. Rignerf: Fully controllable neural 3d portraits. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 20364–20373, 2022. [1](#)
- [4] Ziqian Bai, Feitong Tan, Zeng Huang, Kripasindhu Sarkar, Danhang Tang, Di Qiu, Abhimitra Meka, Ruofei Du, Mingsong Dou, Sergio Orts-Escolano, et al. Learning personalized high quality volumetric head avatars from monocular rgb videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16890–16900, 2023. [1](#), [2](#), [3](#)
- [5] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20919–20929, 2023. [2](#)
- [6] V Blanz and T Vetter. A morphable model for the synthesis of 3d faces. In *26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH 1999)*, pages 187–194. ACM Press, 1999. [1](#), [3](#)
- [7] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. [2](#), [3](#), [4](#), [6](#)
- [8] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#), [3](#)
- [9] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *Proceedings of the European Conference on Computer Vision*, pages 333–350. Springer, 2022. [2](#)
- [10] Anton Cherepkov, Andrey Voynov, and Artem Babenko. Navigating the gan parameter space for semantic image editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3671–3680, 2021. [2](#)
- [11] deepfakes. faceswap. <https://github.com/deepfakes/faceswap>, 2023. Accessed: 2023-10-10. [1](#), [2](#)
- [12] deepfakes. roop. [SomdevSangwan](#), 2023. Accessed: 2023-10-10. [1](#), [2](#), [5](#), [6](#), [7](#)
- [13] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. [5](#)
- [14] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (TOG)*, 40(4): 1–13, 2021. [7](#)
- [15] Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8649–8658, 2021. [1](#), [3](#)
- [16] Xuan Gao, Chenglai Zhong, Jun Xiang, Yang Hong, Yudong Guo, and Juyong Zhang. Reconstructing personalized semantic facial nerf models from monocular video. *ACM Transactions on Graphics (TOG)*, 41(6):1–12, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [17] Songwei Ge, Taesung Park, Jun-Yan Zhu, and Jia-Bin Huang. Expressive text-to-image generation with rich text. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7545–7556, 2023. [2](#)
- [18] 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](#)
- [19] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *Advances in neural information processing systems*, 33:9841–9850, 2020. [2](#)
- [20] Amir Hertz, Kfir Aberman, and Daniel Cohen-Or. Delta denoising score. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2328–2337, 2023. [2](#)
- [21] Lianghua Huang, Di Chen, Yu Liu, Yujun Shen, Deli Zhao, and Jingren Zhou. Composer: Creative and controllable image synthesis with composable conditions. *arXiv preprint arXiv:2302.09778*, 2023. [2](#)
- [22] Wentao Jiang, Si Liu, Chen Gao, Jie Cao, Ran He, Jiashi Feng, and Shuicheng Yan. Psgan: Pose and expression robust spatial-aware gan for customizable makeup transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5194–5202, 2020. [2](#), [4](#)
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. [2](#)
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. [2](#), [4](#), [5](#)
- [25] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. [2](#)

- [26] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017. 3, 4
- [27] Tingting Li, Ruihe Qian, Chao Dong, Si Liu, Qiong Yan, Wenwu Zhu, and Liang Lin. Beautygan: Instance-level facial makeup transfer with deep generative adversarial network. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 645–653, 2018. 2, 4
- [28] K-E Lin, Alex Trevithick, Keli Cheng, Michel Sarkis, Mohsen Ghafoorian, Ning Bi, Gerhard Reitmayr, and Ravi Ramamoorthi. Pvp: Stylized video prior for editable dynamic portraits using stylegan. In *Computer Graphics Forum*, page e14890. Wiley Online Library, 2023. 1, 2, 5, 6, 7
- [29] Huan Ling, Karsten Kreis, Daiqing Li, Seung Wook Kim, Antonio Torralba, and Sanja Fidler. Editgan: High-precision semantic image editing. *Advances in Neural Information Processing Systems*, 34:16331–16345, 2021. 2
- [30] Pengyang Ling, Lin Chen, Pan Zhang, Huaian Chen, and Yi Jin. Freedrag: Point tracking is not you need for interactive point-based image editing. *arXiv preprint arXiv:2307.04684*, 2023. 2, 4
- [31] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. 2
- [32] Aryan Mikaeili, Or Perel, Mehdi Safaei, Daniel Cohen-Or, and Ali Mahdavi-Amiri. Sked: Sketch-guided text-based 3d editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14607–14619, 2023. 2, 3
- [33] 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*, 65(1):99–106, 2021. 2
- [34] Chong Mou, Xintao Wang, Jiechong Song, Ying Shan, and Jian Zhang. Dragondiffusion: Enabling drag-style manipulation on diffusion models. *arXiv preprint arXiv:2307.02421*, 2023. 2
- [35] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. 2
- [36] Thao Nguyen, Anh Tuan Tran, and Minh Hoai. Lipstick ain't enough: beyond color matching for in-the-wild makeup transfer. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 13305–13314, 2021. 2, 4
- [37] Thu Nguyen-Phuoc, Gabriel Schwartz, Yuting Ye, Stephen Lombardi, and Lei Xiao. Alteredavatar: Stylizing dynamic 3d avatars with fast style adaptation. *arXiv preprint arXiv:2305.19245*, 2023. 2
- [38] Mohit Mendiratta Pan, Mohamed Elgharib, Kartik Teotia, Ayush Tewari, Vladislav Golyanik, Adam Kortylewski, Christian Theobalt, et al. Avatarstudio: Text-driven editing of 3d dynamic human head avatars. *arXiv preprint arXiv:2306.00547*, 2023. 2, 7
- [39] Xingang Pan, Ayush Tewari, Thomas Leimkühler, Lingjie Liu, Abhimitra Meka, and Christian Theobalt. Drag your gan: Interactive point-based manipulation on the generative image manifold. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 2, 4, 5
- [40] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 2
- [41] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2085–2094, 2021. 2, 5
- [42] Ivan Perov, Daiheng Gao, Nikolay Chervoniy, Kunlin Liu, Sugasa Marangonda, Chris Umé, Mr Dpfks, Carl Shift Facenheim, Luis RP, Jian Jiang, et al. Deepfacelab: Integrated, flexible and extensible face-swapping framework. *arXiv preprint arXiv:2005.05535*, 2020. 1, 2
- [43] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on graphics (TOG)*, 42(1):1–13, 2022. 4, 5
- [44] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance Fields without Neural Networks. In *CVPR*, 2022. 2
- [45] Ruizhi Shao, Jingxiang Sun, Cheng Peng, Zerong Zheng, Boyao Zhou, Hongwen Zhang, and Yebin Liu. Control4d: Dynamic portrait editing by learning 4d gan from 2d diffusion-based editor. *arXiv preprint arXiv:2305.20082*, 2023. 2
- [46] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9243–9252, 2020. 2
- [47] Yujun Shi, Chuhui Xue, Jiachun Pan, Wenqing Zhang, Vincent YF Tan, and Song Bai. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. *arXiv preprint arXiv:2306.14435*, 2023. 2
- [48] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. 2
- [49] Jingxiang Sun, Xuan Wang, Yichun Shi, Lizhen Wang, Jue Wang, and Yebin Liu. Ide-3d: Interactive disentangled editing for high-resolution 3d-aware portrait synthesis. *arXiv preprint arXiv:2205.15517*, 2022. 2
- [50] Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. Next3d: Generative neural texture rasterization for 3d-aware head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20991–21002, 2023. 1, 2, 3, 4, 5, 6, 7, 8
- [51] Feitong Tan, Sean Fanello, Abhimitra Meka, Sergio Orts-Escolano, Danhang Tang, Rohit Pandey, Jonathan Taylor,

- Ping Tan, and Yinda Zhang. Volux-gan: A generative model for 3d face synthesis with hdri relighting. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–9, 2022. 7
- [52] Junshu Tang, Bo Zhang, Binxin Yang, Ting Zhang, Dong Chen, Lizhuang Ma, and Fang Wen. 3dfaceshop: Explicitly controllable 3d-aware portrait generation. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 1, 2
- [53] Ziyang Wang, Timur Bagautdinov, Stephen Lombardi, Tomas Simon, Jason Saragih, Jessica Hodgins, and Michael Zollhofer. Learning compositional radiance fields of dynamic human heads. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5704–5713, 2021. 2
- [54] Yue Wu, Yu Deng, Jiaolong Yang, Fangyun Wei, Qifeng Chen, and Xin Tong. Anifacegan: Animatable 3d-aware face image generation for video avatars. *Advances in Neural Information Processing Systems*, 35:36188–36201, 2022. 1, 2
- [55] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. Gan inversion: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3121–3138, 2022. 4
- [56] Yinghao Xu, Sida Peng, Ceyuan Yang, Yujun Shen, and Bolei Zhou. 3d-aware image synthesis via learning structural and textural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18430–18439, 2022. 2
- [57] Yuelang Xu, Lizhen Wang, Xiaochen Zhao, Hongwen Zhang, and Yebin Liu. Avatarmav: Fast 3d head avatar reconstruction using motion-aware neural voxels. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–10, 2023. 1, 3
- [58] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13779–13788, 2021. 2
- [59] 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*, pages 597–614. Springer, 2022.
- [60] Bangbang Yang, Yinda Zhang, Yijin Li, Zhaopeng Cui, Sean Fanello, Hujun Bao, and Guofeng Zhang. Neural rendering in a room: amodal 3d understanding and free-viewpoint rendering for the closed scene composed of pre-captured objects. *ACM Transactions on Graphics (TOG)*, 41(4):1–10, 2022. 2
- [61] Chenyu Yang, Wanrong He, Yingqing Xu, and Yang Gao. Elegant: Exquisite and locally editable gan for makeup transfer. In *European Conference on Computer Vision*, pages 737–754. Springer, 2022. 2
- [62] Xingrui Yang, Hai Li, Hongjia Zhai, Yuhang Ming, Yuqian Liu, and Guofeng Zhang. Vox-fusion: Dense tracking and mapping with voxel-based neural implicit representation. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 499–507. IEEE, 2022. 2
- [63] Weicai Ye, Shuo Chen, Chong Bao, Hujun Bao, Marc Pollefeys, Zhaopeng Cui, and Guofeng Zhang. Intrinsicnerf: Learning intrinsic neural radiance fields for editable novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 339–351, 2023. 2
- [64] Deheng Zhang, Clara Fernandez-Labrador, and Christopher Schroers. Coarf: Controllable 3d artistic style transfer for radiance fields. In *2024 International Conference on 3D Vision (3DV)*. IEEE, 2024. 2
- [65] Hao Zhang, Yanbo Xu, Tianyuan Dai, Tai Chi-Keung Tang, et al. Fdnerf: Semantics-driven face reconstruction, prompt editing and relighting with diffusion models. *arXiv preprint arXiv:2306.00783*, 2023. 2
- [66] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)*, 40(6):1–18, 2021. 2
- [67] Boming Zhao, Bangbang Yang, Zhenyang Li, Zuoyue Li, Guofeng Zhang, Jiashu Zhao, Dawei Yin, Zhaopeng Cui, and Hujun Bao. Factorized and controllable neural re-rendering of outdoor scene for photo extrapolation. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1455–1464, 2022. 2
- [68] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C Bühler, Xu Chen, Michael J Black, and Otmar Hilliges. Im avatar: Implicit morphable head avatars from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13545–13555, 2022. 1, 2, 3
- [69] Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J Black, and Otmar Hilliges. Pointavatar: Deformable point-based head avatars from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21057–21067, 2023. 1, 2, 3
- [70] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *European conference on computer vision*, pages 592–608. Springer, 2020. 2
- [71] Jiapeng Zhu, Ceyuan Yang, Yujun Shen, Zifan Shi, Bo Dai, Deli Zhao, and Qifeng Chen. Linkgan: Linking gan latents to pixels for controllable image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7656–7666, 2023. 2
- [72] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5104–5113, 2020. 2
- [73] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022. 2

- [74] 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](#)
- [75] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Towards metrical reconstruction of human faces. In *European Conference on Computer Vision*, pages 250–269. Springer, 2022. [7](#)
- [76] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Instant volumetric head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4574–4584, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [77] zllrunning. face-makeup.pytorch. <https://github.com/zllrunning/face-makeup.PyTorch>, 2023. Accessed: 2023-10-10. [4](#)