# 3D Gaussian Parametric Head Model

Yuelang Xu<sup>1</sup><sup>®</sup>, Lizhen Wang<sup>1</sup><sup>®</sup>, Zerong Zheng<sup>2</sup><sup>®</sup>, Zhaoqi Su<sup>1</sup><sup>®</sup>, and Yebin Liu<sup>1</sup><sup>®</sup>

<sup>1</sup> Tsinghua University, Beijing, China
<sup>2</sup> NNKosmos, Hangzhou, China

Abstract. Creating high-fidelity 3D human head avatars is crucial for applications in VR/AR, telepresence, digital human interfaces, and film production. Recent advances have leveraged morphable face models to generate animated head avatars from easily accessible data, representing varying identities and expressions within a low-dimensional parametric space. However, existing methods often struggle with modeling complex appearance details, e.g., hairstyles and accessories, and suffer from low rendering quality and efficiency. This paper introduces a novel approach, 3D Gaussian Parametric Head Model, which employs 3D Gaussians to accurately represent the complexities of the human head, allowing precise control over both identity and expression. Additionally, it enables seamless face portrait interpolation and the reconstruction of detailed head avatars from a single image. Unlike previous methods, the Gaussian model can handle intricate details, enabling realistic representations of varying appearances and complex expressions. Furthermore, this paper presents a well-designed training framework to ensure smooth convergence, providing a guarantee for learning the rich content. Our method achieves high-quality, photo-realistic rendering with real-time efficiency, making it a valuable contribution to the field of parametric head models.

Keywords: 3D Gaussian · Head Avatar · Parametric Model

## 1 Introduction

Creating high-fidelity 3D human head avatars holds significant importance across various fields, including VR/AR, telepresence, digital human interfaces, and film production. The automatic generation of such avatars has been a focal point in computer vision research for many years. Recent methods [12,13,17,38,55,56,61–63,65] can create an animated head avatar through conveniently collected data such as monocular video data or even a picture [22, 26]. Serving as the most fundamental tool in these methods, the 3D morphable models (3DMM) [14, 25], which represent varying identities and expressions within a low-dimensional space, have been proven to be a highly successful avenue in addressing this challenging problem.

Since the traditional parametric 3DMMs are typically limited by the topology of the underlying template mesh and only focus on the face part, some works [15, 16, 28, 59] propose to use implicit Signed Distance Field (SDF) as the



Fig. 1: We utilize hybrid datasets comprising captured multi-view video data and rendered image data from 3D scans for training our model. The trained model can be manipulated using decoupled identity and expression codes to produce a diverse array of high-fidelity head models. When presented with an image, our model can be adjusted to reconstruct the portrait in the image and edit the expression according to any other desired expressions.

geometric representation to model the entire head. Despite their flexibility, these methods fall short in recovering high-frequency geometric and texture details like hairstyles, glasses or accessories. On the other end of the spectrum, Neural Radiance Field (NeRF) [34] based methods [19,64] learn parametric head models by directly synthesizing images, thus eliminating the need of geometry modeling. However, NeRF is built upon volumetric rendering, which involves sampling and integrating points distributed throughout space. Therefore, NeRF-based methods typically suffer from low rendering efficiency and have to trade it off with rendering resolution, thereby greatly reducing rendering quality. Moreover, skipping geometric reconstruction would probably lead to poor 3D consistency.

More recently, 3D Gaussian Splatting (3DGS) [21], which uses explicit Gaussian ellipsoids to represent 3D scenes, has attracted significant attention from the research community. Experiments have verified the superior quality of the rendered results and excellent rendering efficiency compared to previous NeRF-based or surface-based methods even on dynamic scenes [32,48,57,58]. Motivated by this progress, we propose a novel **3D Gaussian Parametric Head Model**, which, for the first time, marries the power of 3DGS with the challenging task of parametric head modeling. Our 3D gaussian parametric head model decouples the control signals of the head into the latent spaces of identity and expression, as is also done in SDF-based face model NPHM [15]. These latent spaces are then mapped to the offsets of the Gaussian positions, which effectively represent the variance of shape and appearance of different identities and expressions. Benefiting from the differentiability of Gaussian splatting, our model can be learned from multi-view video data corpus in an end-to-end manner, without relying on geometry supervision.

Unfortunately, training our 3D Gaussian parametric head model is not quite straightforward, because Gaussian ellipsoids are unstructured and each Gaussian ellipsoid has its own independent learnable attribute. Such a characteristic makes 3DGS powerful in overfitting a specific object or scene, but poses great challenges for generative head modeling. Without proper initialization and regularization, the learned parametric head model may suffer from unstable training or a large number of Gaussian points becoming redundant and noisy, as shown in Fig. 4.

To overcome these challenges, we propose a well-designed two-stage training strategy to ensure smooth convergence of our model training. Specifically, we first roughly train all the networks on a mesh-based guiding model. Subsequently, the network parameters are migrated to the Gaussian model, and all Gaussian points are initialized with the trained mesh geometry to ensure that they are located near the actual surface. Compared to naive initialization with FLAME [25], our initialization strategy leads to a better guess of the positions of Gaussian points, making the subsequent training of the model converge stably and the areas like hairs better recovered. Moreover, we propose to use 3D landmark loss to supervise the deformation of the model learning expressions, which can speed up the convergence and avoid artifacts under exaggerated expressions. Lastly, our method supports training from both 3D head scans and multi-view 2D face datasets, which enhances the versatility and comprehensiveness of facial data collection and model training.

After training on large corpus of multi-view head videos, our parametric Gaussian head model can generate photorealistic images that accurately depict the diverse range of facial appearances, naturally handling complex and exaggerated expressions, while also enabling real-time rendering. Additionally, our method supports single-image fitting and surpasses previous techniques in both reconstruction accuracy and identity consistency. Furthermore, the model resulting from our fitting process allows for the control of various expressions while maintaining naturalness and consistent identity even under exaggerated expressions. The contributions of our method can be summarized as:

- We propose 3D Gaussian Parametric Head Model, a novel parametric head model which utilizes 3D Gaussians as the representation and enables photorealistic rendering quality and real-time rendering speed.
- We propose a well-designed training strategy to ensure that the Gaussian model converges stably while learning rich appearance details and complex expressions efficiently.
- Our 3D Gaussian Parametric Head Model enables the generation of a detailed, high-quality face avatar from a single given image, as well as performing expression and identity editing upon it.

## 2 Related Work

**Parametric Head Models.** Parametric head models are used to represent facial features, expressions, and identities effectively and efficiently. They allow for the creation of realistic human faces with adjustable parameters, making them essential in computer graphics, animation, and virtual reality. Therefore, research in this field has always been a hot topic. Traditional 3D Morphable Models (3DMM) [2,6,14,25,47] are constructed by non-rigidly registering a template mesh with fixed topology to a series of 3D scans. Through this registration process, a 3DMM can be computed using dimensionality reduction techniques such as principal component analysis (PCA). The resulting parametric space captures the variations in facial geometry and appearance across a population. However, while 3DMMs offer a powerful way to represent faces, they do have limitations. These models rely heavily on the correspondence between the 3D scans and the template for accurate fitting and may struggle to represent local surface details like wrinkles or hairstyles that deviate significantly from the template mesh. Recent advances in implicit representation have led to the great development of neural parametric head models. Some methods [15, 16, 49, 59] propose implicit Signed Distance Field (SDF) based head models, which are not constrained by topology thus can recover more complex content like hair compared to previous mesh-based Methods. Other methods [3, 19, 44, 64] propose to use NeRF [34] as the representation of the parametric head models, which can directly synthesize photorealistic images without geometric reconstruction. Cao, et al. [5] use a hybrid representation [30] of mesh and NeRF to train their model on unpublished large-scale light stage data. However, rendering efficiency is typically low in NeRF-based methods, often resulting in a trade-off with rendering resolution.

3D GAN based Head Models. 3D Generative Adversarial Networks (GANs) have revolutionized the field of computer vision, particularly in the domain of human head and face modeling, enabling the generation of face avatars from input images. Traditional methods often require labor-intensive manual work or rely on multi-view images to create 3D models. 3D GANs as a more automated and data-driven approach, which are just trained on single-view 2D images but generate detailed and realistic 3D models of human head [7–9, 11, 18, 35, 52]. Panohead [1] additionally introduces images of hairstyles on the back of characters and trains a full-head generative model. Based on the previous methods, IDE-3D [42] proposes to use semantic map to edit the 3D head model. Next3D [43] and AniFaceGAN [50] extend to use the FLAME model [25] to condition the generated head model, so that the expression and pose of the generated head model can be controlled. AniPortraitGAN [51] further replaces FLAME model with SMPLX model [36] to generate upper body avatars, thus the shoulders and the neck can also be controlled. These 3D GAN-based models primarily leverage the coarse FLAME model for expression control, often leading to a loss of expression details in the generated faces. In contrast, our method directly learns the expression distribution from the dataset, capturing more facial appearance details.

**3D** Gaussians. Recently, 3D Gaussian splatting [21] has shown superior performance compared to NeRF, excelling in both novel view synthesis quality and rendering speed. Several methods have expanded Gaussian representation to dynamic scene reconstruction [32, 48, 57, 58]. For human body avatar modeling, recent approaches [20, 27] propose training a 3D Gaussian avatar animated by SMPL [31] or a skeleton from multi-view videos, surpassing previous methods in rendering quality and efficiency. In the realm of human head avatar modeling, recent techniques [10, 23, 33, 37, 37, 39, 40, 45, 53, 54] also utilize 3D Gaussians to create high-fidelity and efficient head avatars. These approaches center on the creation of a high-fidelity person-specific avatar using data of a single person. In



**Fig. 2:** The pipeline of our method. Our training strategy can be divided into a Guiding Geometry Model for initialization, and a final 3D Gaussian Parametric Head Model. Deformations of each model are further decoupled into identity-related and expression-related deformations. Rendering involves using DMTet to transform the initial model into a mesh and 3D Gaussian Splatting for the Gaussian model. Features from both models are finally upsampled to high-resolution portrait images through a convolutional network  $\Psi$ . During inference, our output exclusively comes from the Gaussian model.

contrast, our method focuses on a versatile prior model that can accommodate varying appearances. Once trained, our model is also capable of person-specific avatar reconstruction by fitting to the input image data provided.

## 3 Method

In this section, we present the 3D Gaussian Parametric Head Model. In contrast to previous mesh-based or NeRF-based models, initializing and training Gaussian-based models pose distinct challenges. This section introduces the dataset and preprocessing, the carefully designed guiding geometry model, the Gaussian Parametric Head Model, and outlines their respective training processes. Additionally, we will also provide the training details and demonstrate how to utilize our method when given a single input image.

### 3.1 Data Preprocessing

We used three datasets for our model training, including a multi-view video dataset NeRSemble [24], and two 3D scans datasets NPHM [15] and FaceV-erse [47]. We do not use the 3D geometry of the scans directly, but render them

into multi-view images and use only the images form the 3 datasets as supervision. In order to better utilize these three different datasets, we need to do preprocessing. First, we resize the images to 512 resolution and adjust the camera parameters. Then, we use BackgroundMattingV2 [29] to extract the foreground characters in the NeRSemble dataset and record the masks. This step is not required for the two synthetic datasets. Next, we use face alignment [4] to detect 2D landmarks in all the images. Through these 2D landmarks, we fit a Basel Face Model (BFM) [14] for each expression of each identity, and record the head pose and 3D landmarks of the BFM. We will use the above processed camera parameters, images, masks, head pose of BFM and 3D landmarks of BFM to train our model.

#### 3.2 Model Representation

The representation of Gaussian distribution poses challenges due to its unordered and unstructured nature, leading to difficulties in the continuous spread of gradients to neighboring points in space during backpropagation. This often results in convergence failure when Gaussians are randomly initialized. On the other hand, surface-based representations such as mesh are just suitable for rough geometry learning. A direct idea is to utilize an existing 3DMM, such as FLAME [25], as the initial position for the points in 3D Gaussian splatting [21]. However, this coarse initialization still fails to converge the positions of 3D points to the correct locations, as shown in Fig. 4. The network tends to alter the shape of the ellipsoid to achieve a suitable fitting result, leading to inaccurate geometry of the point cloud and blurriness in the rendered image.

To address this problem, a more detailed initialization process is necessary for capturing the diverse head variations using 3D Gaussian splatting. Specifically, we draw inspiration from Gaussian Head Avatar [54] and leverage the implicit signed distance field (SDF) representation to train a guiding geometry model. This guiding geometry model serves as the initial value for the Gaussian model, providing a more effective starting point for the optimization process. We define the initial model as Guiding Geometry Model and the refined model as 3D Gaussian Parametric Head Model.

Guiding Geometry Model. The guiding geometry model receives an identity code  $z^{id}$  and an expression code  $z^{exp}$  as input, producing a mesh with vertices V, faces F, and per-vertex color C that aligns with the specified identity and expression. To achieve this, we use an MLP denoted as  $f_{mean}(\cdot)$  to implicitly model the SDF, which represents the mean geometry:

$$s, \gamma = \boldsymbol{f_{mean}}(x), \tag{1}$$

where s denotes the SDF value,  $\gamma$  denotes the feature from the last layer and x denotes the input position. Then, we convert the implicit SDF through Deep Marching Tetrahedra (DMTet) [41] into an explicit mesh with vertex positions  $V_0$ , per-vertex feature  $\Gamma$  and faces F. Next, we need to transform the mean shape into a neutral-expression shape on condition of the input identity code  $z^{id}$ . To inject identity information into the vertices of the mesh, we first use an injection MLP  $f_{inj}(\cdot)$ , which takes the identity code  $z^{id}$  and the per-vertex feature  $\Gamma$  as input and produces the identity-conditioned per-vertex feature vectors  $H = f_{inj}(z^{id}, \Gamma)$ . Subsequently, utilizing a tiny MLP  $f_{id}(\cdot)$ , we predict the displacement  $\delta V_{id}$  for each vertex. This displacement is used to transform the mean shape into the neutral-expression shape conditioned on the id code  $z^{id}$ .

After completing deformations related to identity, the next step is to capture the deformation induced by facial expressions. We introduce another tiny MLP  $f_{exp}(\cdot)$ . This MLP takes the feature vectors H obtained in the previous step and the expression code  $z^{exp}$  as input, and the output is the displacement  $\delta V_{exp}$ for each vertex. Using this displacement, we update the vertex positions to  $V_{can}$ . Additionally, we feed the same feature vectors H and expression code  $z^{exp}$  to a color MLP,  $f_{col}(\cdot)$ , to predict the 32-channel color C for each vertex. The vertex positions to  $V_{can}$  and 32-channel color C can be described as:

$$V_{can} = V_0 + \boldsymbol{f_{id}}(H) + \boldsymbol{f_{exp}}(H, \boldsymbol{z^{exp}}), \ C = \boldsymbol{f_{col}}(H, \boldsymbol{z^{exp}}).$$
(2)

Finally, we utilize the estimated head pose parameters R and T obtained during data preprocessing to transform the mesh from the canonical space to the world space  $V = R \cdot V_{can} + T$ . After generating the final vertex positions, colors and faces  $\{V, C, F\}$  of the mesh, we render the mesh into a 256-resolution 32-channel feature map  $I_F$  and a mask M through differentiable rasterization with a given camera pose. Subsequently, the feature map is interpreted as a 512-resolution RGB  $I_{hr}$  image through a lightweight convolutional upsampling network  $\Psi(\cdot)$ , as shown in Fig. 2.

3D Gaussian Parametric Head Model. The Gaussian model also takes an identity code  $z^{id}$  and an expression code  $z^{exp}$  as input, producing the positions X, color C, scale S, rotation Q and opacity A of the 3D Gaussians. Similar to the guiding geometry model, we initially maintain an overall mean point cloud, with the mean positions  $X_0$ . However, we no longer generate the per-vertex feature  $\Gamma$  through  $f_{mean}(x)$ . Instead, we directly generate it at once and bind it to the Gaussian points as optimizable variables  $\Gamma_0$ . This is possible since the number of Gaussian points is fixed at this stage. Then we need to transform the mean point cloud into a neutral-expression point cloud, conditioned by the id code  $z^{id}$ . To achieve this, we utilize the same injection MLP  $f_{inj}(\cdot)$  and identity deformation MLP  $f_{id}(\cdot)$  defined in the guiding geometry model, which can generate feature vectors  $H = f_{inj}(z^{id}, \Gamma_0)$  that encode identity information for each point and predict the identity-related displacement of each point. Then, we also need to predict the expression code  $z^{exp}$ -conditioned displacement. The resulting positions  $X_{can}$  and the 32-channel color C of each point, similar to the approach presented in the guiding geometry model, can be described as:

$$X_{can} = X_0 + f_{id}(H) + f_{exp}(H, z^{exp}), \ C = f_{col}(H, z^{exp}).$$
(3)

Unlike the representations of SDF and DMTet, Gaussians have additional attributes that need to be predicted. Here, we introduce a new MLP to predict Gaussian attributes in the canonical space, including the scale S, rotation  $Q_{can}$ ,

and opacity A. In order to ensure the stability of the generated results, we refrain from directly predicting these values. Instead, we predict their offsets  $\{\delta S, \delta Q, \delta A\}$  relative to the overall mean values  $\{S_0, Q_0, A_0\}$ :

$$\{S, Q_{can}, A\} = \{\boldsymbol{S}_0, \boldsymbol{Q}_0, \boldsymbol{A}_0\} + \boldsymbol{f}_{\boldsymbol{att}}(H, \boldsymbol{z}^{\boldsymbol{exp}}).$$

$$\tag{4}$$

Following this, we utilize the estimated head pose parameters R and T, obtained during data preprocessing, to transform the canonical space variables  $X_{can}$  and  $Q_{can}$  into the world space:  $X = R \cdot X_{can} + T$ ,  $Q = R \cdot Q_{can}$ . For model rendering, we leverage differentiable rendering [21] and neural rendering techniques to generate images. The generated 3D Gaussian parameters, which include  $\{X, C, S, Q, A\}$ , are conditioned by the identity code  $z^{id}$  and expression code  $z^{exp}$ . Finally, we input this feature map into the same upsampling network  $\Psi(\cdot)$  of the guiding geometry model to generate a 512-resolution RGB image.

In the 3D Gaussian Parametric Head Model, we leverage the previously trained guiding geometry model to initialize our variables and networks, rather than initiating them randomly and training from scratch. Specifically, we initialize the Gaussian positions  $X_0$  using the vertex positions of the mean mesh  $V_0$ . Meanwhile, we generate the per-vertex feature  $\Gamma$  from  $f_{mean}(x)$  at the beginning and bind it to the points as an optimizable variable  $\Gamma_0$  as described above. Additionally, all identity codes  $z^{id}$ , expression codes  $z^{exp}$ , and the networks  $\{f_{inj}(\cdot), f_{id}(\cdot), f_{exp}(\cdot), f_{col}(\cdot), \Psi(\cdot)\}$  are directly inherited from the guiding geometry model. Note that, the attribute MLP  $f_{att}(\cdot)$  is a newly introduced network, hence it is initialized randomly. Finally, the overall mean values of the Gaussian attributes  $\{S_0, Q_0, A_0\}$  are initialized following the original 3D Gaussian Splatting [21].

#### 3.3 Loss Functions

To ensure the accurate convergence of the model, we employ various loss functions as constraints, including the basic photometric loss and silhouette loss, to enforce consistency with ground truth of both the rendered high-resolution images  $I_{hr}$  and the rendered masks M:

$$\mathcal{L}_{hr} = ||I_{hr} - I_{gt}||_1, \ \mathcal{L}_{sil} = IOU(M, M_{gt}), \tag{5}$$

with  $I_{gt}$  representing the ground truth RGB images,  $M_{gt}$  representing the ground truth masks. We further encourage the first three channels of the low-resolution feature map  $I_{lr}$  to closely match the ground-truth RGB image  $I_{gt}$  by introducing an  $L_1$  loss:

$$\mathcal{L}_{lr} = ||I_{lr} - I_{gt}||_1.$$
(6)

The geometric deformation caused by expressions is typically complex and cannot be learned through image supervision alone. Therefore, we provide additional coarse supervision for expression deformation learning using 3D landmarks. Specifically, we define the 3D landmarks  $P_0$  in the canonical space, and then predict their displacements and transform them to the world space as P

just like the transformation of the original vertices  $V_0$  above. Then, we construct the landmark loss function:

$$\mathcal{L}_{lmk} = ||\boldsymbol{P} - \boldsymbol{P}_{qt}||_2, \tag{7}$$

with  $P_{gt}$  denoting the ground truth 3D landmarks, which are estimated by fitting a BFM model to the training data during preprocessing.

Moreover, to guarantee the decoupling of identity and expression deformations learned by the model and minimize redundancy, we introduce the following regularization loss function that aims to minimize the magnitude of both deformations:

$$\mathcal{L}_{reg} = ||\delta V_{id}||_2 + ||\delta V_{exp}||_2. \tag{8}$$

During the training of the **Guiding Geometry Model**, we also construct a Laplacian smooth term  $\mathcal{L}_{lap}$  to penalize surface noise or breaks. Overall, the total loss function is formulated as:

$$\mathcal{L} = \mathcal{L}_{hr} + \lambda_{sil}\mathcal{L}_{sil} + \lambda_{lr}\mathcal{L}_{lr} + \lambda_{lmk}\mathcal{L}_{lmk} + \lambda_{reg}\mathcal{L}_{reg} + \lambda_{lap}\mathcal{L}_{lap} \tag{9}$$

with all the  $\lambda$  denoting the weights of each term. In practice, we set  $\lambda_{sil} = 0.1$ ,  $\lambda_{lr} = 0.1$ ,  $\lambda_{lmk} = 0.1$ ,  $\lambda_{reg} = 0.001$  and  $\lambda_{lap} = 100$ . During training, we jointly optimize the bolded variables above:  $\{\boldsymbol{z^{id}}, \boldsymbol{z^{exp}}, \boldsymbol{f_{inj}}(\cdot), \boldsymbol{f_{mean}}(\cdot), \boldsymbol{f_{id}}(\cdot), \boldsymbol{f_{exp}}(\cdot), \boldsymbol{f_{col}}(\cdot), \boldsymbol{\psi}(\cdot), \boldsymbol{P_0}\}$ . Notably, the defined canonical 3D landmarks  $\boldsymbol{P_0}$  are initialized by computing the average of the estimated 3D landmarks from the training dataset.

During the training stage of the **3D** Gaussian Parametric Head Model, we also calculate the perceptual loss [60] to encourage the model to learn more high-frequency details  $\mathcal{L}_{vgg} = VGG(I_{hr}, I_{gt})$ . Similar to training the guiding geometry model, we enforce the first three channels of the feature map to be RGB channels as Eqn. 6, introduce landmarks guidance terms as Eqn. 7 and the regular term for the displacement of points as Eqn. 8. Consequently, the overall loss function can be formulated as:

$$\mathcal{L} = \mathcal{L}_{hr} + \lambda_{vgg} \mathcal{L}_{vgg} + \lambda_{lr} \mathcal{L}_{lr} + \lambda_{lmk} \mathcal{L}_{lmk} + \lambda_{reg} \mathcal{L}_{reg}$$
(10)

with the weights  $\lambda_{vgg} = 0.1$ ,  $\lambda_{lr} = 0.1$ ,  $\lambda_{lmk} = 0.1$  and  $\lambda_{reg} = 0.001$ . In this training stage, we also jointly optimize all the bolded variables and networks mentioned above, including the overall mean positions and attributes of the Gaussians and the 3D landmarks:  $\{\boldsymbol{z^{id}}, \boldsymbol{z^{exp}}, \boldsymbol{f_{inj}}(\cdot), \boldsymbol{f_{id}}(\cdot), \boldsymbol{f_{exp}}(\cdot), \boldsymbol{f_{col}}(\cdot), \boldsymbol{f_{att}}(\cdot), \boldsymbol{\psi}(\cdot), \boldsymbol{X_0}, \boldsymbol{\Gamma_0}, \boldsymbol{S_0}, \boldsymbol{Q_0}, \boldsymbol{A_0}, \boldsymbol{P_0}\}.$ 

#### 3.4 Inference Details

**Image-based Fitting.** When a single RGB portrait image is input, we first align the image according to the processing rules of the training set. Subsequently, we employ gradient descent to fit the image rendered by the 3D Gaussian Parametric Head Model to this input image using the photometric loss  $\mathcal{L}_{lr}$ 

and  $\mathcal{L}_{hr}$  defined in Eqn. 10. This process helps regress the identity code  $z^{id}$ and expression code  $z^{exp}$ . We just optimize for 200 iterations with learning rate  $1 \times 10^{-3}$  for both latent codes. Following this, we fix the latent codes  $z^{id}$  and  $z^{exp}$ , such that the variables H,  $X_{can}$  are also fixed. We further optimize the color MLP  $f_{col}(\cdot)$  and the canonical positions  $X_{can}$  which represent the geometry of the current specific subject, using the same loss function. In this step, we only optimize for 100 iterations with learning rate  $1 \times 10^{-4}$  for both  $f_{col}(\cdot)$ and  $X_{can}$ . This optimization process aims to add some details that cannot be recovered by the trained model itself, ultimately resulting in the reconstructed head model. The entire process has a total of 300 iterations and takes only 30 seconds.

**Expression Editing.** Given a source portrait image providing the subject whose expression is to be edited and a target portrait image providing the target expression. We first obtain the head model of the source subject through optimization as the above-mentioned Image-based Fitting strategy. Then for the target portrait image, we also obtain the head model and corresponding expression code in the same way. Finally, we input the target expression code to the head model of the source subject, so that the expression of the source subject can be edited to the target one.

## 4 Experiments

#### 4.1 Datasets

**NeRSemble** dataset contains over 260 different identities, and collects 72fps multi-view videos from 16 synchronized cameras for each identity. The total length of the videos of a single identity is approximately 6000-11000 frames. In the experiment, we selected 140 of the identities for training and the rest for evaluation. For each identity video, we selected about 150 frames from all 16 views as training data.

**NPHM** dataset contains 5200 3D human head scans. These scans come from 255 different identities, each with about 20 different expressions. We selected approximately 1600 scans of 80 identities for training. Since our method utilizes 2D images as training supervision, we render each scan from 80 different views to generate synthetic image data and record the camera parameters and the masks.

**FaceVerse** dataset is an East Asian human head scan dataset. It contains 2310 scans from 110 different identities, and each identity contains 21 expressions. We selected 1620 scans data of 80 identities for training. Similarly, for each scan, we render multi-view synthetic image data from 80 different views and record the camera parameters and the masks.

## 4.2 Evaluation

**Disentanglement.** We tested the performance of the 3D Gaussian Parametric Model under the control of different identity codes and different expression codes.



Fig. 3: We generate the head models with randomly sampled identity codes and expression codes as condition. Each row corresponds to the same identity code, and each column corresponds to the same expression code.



Fig. 4: We compared our initialization strategy with using the vertices of FLAME model. The left side shows the rendered image, and the right side shows the positions of the Gaussian points.

We randomly sampled 2 identity codes and 5 expression codes to generate 10 head models. Each horizontal row corresponds to the same identity code, and each column corresponds to the same expression code, as shown in Fig. 3. It can be observed that our model performs well in identity consistency and expression consistency, and the two components are fully disentangled.

Ablation on Initialization. To evaluate the effectiveness of our initialization strategy with guiding geometry model outlined in Section 3, we compare it against a FLAME-based initialization strategy. To use FLAME model for the initialization, we first fit a FLAME model to overall mean 3D landmarks which are estimated during data preprocessing. Then, we sample 100,000 points near the surface of the FLAME mesh as an initialization of the mean Gaussian positions  $X_0$ . For the per-vertex features bound to each point  $\Gamma$ , we just set them to zero. And for all the networks  $\{f_{inj}(\cdot), f_{id}(\cdot), f_{exp}(\cdot), f_{col}(\cdot), \Psi(\cdot)\}$  and  $f_{att}(\cdot)$  are randomly initialized as there is no available prior. The initialization process for the Gaussian attributes  $\{S_0, Q_0, A_0\}$  remains the same as in our strategy.

We show the visualization results in Figure 4, with the Gaussian model rendering image on the left and the Gaussian positions displayed as point clouds



Fig. 5: The comparison of the different representations with super resolution.

on the right. Our initialization strategy using the guiding geometry model can ensure that all the Gaussian points fall evenly on the actual surface of the model, thereby ensuring reconstruction quality. When using the FLAME model for the initialization, a large number of points wander inside or outside the actual surface of the model, causing noise or redundancy and leading the model to lose some high-frequency information and making it difficult to fully converge. We also perform a quantitative evaluation of different initialization strategies on the rendered images, as shown in Table 1, which shows that our method leads to better rendering results.

Method	$\mathrm{PSNR}\uparrow$	SSIM $\uparrow$	$\mathrm{LPIPS}\downarrow$
FLAME Initialization	25.7	0.82	0.109
Our Initialization	<b>28.0</b>	0.84	0.085

 Table 1: Quantitative evaluation results of our initialization strategy and naive

 FLAME initialization strategy.

Ablation on Representation and Super Resolution. We conduct the ablation study for the guiding mesh model, the Gaussian model and the superresolution network (abbreviated as SR) as shown in the Fig. 5. The corresponding PSNR metrics are: Mesh (15.7), Mesh+SR (17.3), Gaussian (27.0), Gaussian+SR (29.3). Compared to mesh, utilizing 3D Gaussian as the representation brings significant improvements (+12), while the super resolution module adds some details, generating more realistic results.

#### 4.3 Applications

Image-based Fitting. In this section, we demonstrate the capability of our 3D Gaussian Parametric Model for single-image fitting using the fitting strategy detailed in Section 3.4. We compare our model with similar works: HeadNeRF [19], MoFaNeRF [64], and PanoHead [1]. In addition to evaluating the above methods on our evaluation dataset, we also conduct comparisons using cases from MEAD [46] dataset (the first two rows). The qualitative results are presented in Figure 6. Our model exhibits reconstruction accuracy while maintaining excellent 3D consistency and identity preservation. HeadNeRF's fitting results often suffer from missing hair, and they remove the body and neck. MoFaNeRF, trained solely on the FaceScape dataset where all subjects wear hats, struggles to fit hair. As a GAN-based model, PanoHead can achieve highly accurate reproductions



Fig. 6: We compare our method with other SOTA methods on the task of single image fitting. The far left is the input image, and to the right are Our method, HeadNeRF [19], MoFaNeRF [64] and PanoHead [1]. Our model significantly outperforms other methods in reconstruction quality and 3D consistency.

from the input view. However, due to overfitting, the results from side views reveal poor 3D consistency and identity preservation.

In addition to qualitative evaluations, we also conducted quantitative evaluations on 60 images using three metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Face Distance (FD). Here, we provide a brief explanation of the Face Distance (FD). To compute the FD metric, we utilized a face recognition tool <sup>3</sup> to encode two images containing faces into 128dimensional vectors. Subsequently, we calculated the distance between these two vectors to reflect the similarity of the two faces. In our experiments, FD serves as an indicator of identity consistency. The results are shown in Table 2. Our model demonstrates optimal performance in both fitting accuracy and identity consistency.

**Expression Editing.** Our 3D Gaussian Parametric Head Model possesses the capability for expression editing. Upon completing the fitting process on a portrait image, we can animate the model by applying different expression codes.

<sup>&</sup>lt;sup>3</sup> https://github.com/ageitgey/face recognition

Method	$\mathrm{PSNR}\uparrow$	SSIM $\uparrow$	$FD\downarrow$
HeadNeRF	28.9	0.84	0.37
MoFaNeRF	28.6	0.82	0.37
PanoHead	29.1	0.86	0.41
Ours	30.3	0.86	0.35

**Table 2:** Quantitative evaluation results on the task of single image fitting. We compare our method with other 3 SOTA methods: HeadNeRF [19], MoFaNeRF [64], PanoHead [1].



Fig. 7: We perform expression editing on the head model reconstructed from the input image. Our model is able to handle very exaggerated expressions with superior identity consistency.

The detailed pipeline is outlined in Section 3.4. An example is illustrated in Figure 7. Our model can generate images depicting the corresponding expressions of the input subject based on a reference expression (as seen in the lower left corner of each image in the figure). It performs admirably even with exaggerated expressions, producing natural and realistic results.

### 5 Discussion

**Ethical Considerations.** Our technique can generate artificial portrait videos, posing a significant risk of spreading misinformation, shaping public opinions, and undermining trust in media outlets. These consequences could have profound negative effects on society. Therefore, it is crucial to explore methods that effectively differentiate between genuine and manipulated content.

**Limitation.** Our 3D Gaussian Parametric Head Model takes a step forward in the characterization of parametric head models. However, due to the limited amount of training data, the generalization ability of the model is still insufficient. In some cases where the illumination is significantly different from the training set, the reconstruction results are not good.

**Conclusion.** In this paper, we propose the 3D Gaussian Parametric Head Model, a novel framework for parametric head model. This model leverages the power of 3D Gaussians, enabling realistic rendering quality and real-time speed. Our well-designed training strategy ensured stable convergence while enabling the model to learn appearance details and expressions. Besides, our model allows for creating detailed, high-quality face avatars from a single input image, and also enables editing for expressions and identity. We believe our model represents a significant advancement in the field of parametric head model.

## Acknowledgements

The work is supported by the National Science Foundation of China (NSFC) under Grant Number 62125107 and the Postdoctoral Fellowship Program of China Postdoctoral Science Foundation under Grant Number GZC20231304.

### References

- An, S., Xu, H., Shi, Y., Song, G., Ogras, U.Y., Luo, L.: Panohead: Geometry-aware 3d full-head synthesis in 360deg. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20950–20959 (June 2023)
- Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH 1999). pp. 187–194. ACM Press (1999)
- Bühler, M.C., Sarkar, K., Shah, T., Li, G., Wang, D., Helminger, L., Orts-Escolano, S., Lagun, D., Hilliges, O., Beeler, T., et al.: Preface: A data-driven volumetric prior for few-shot ultra high-resolution face synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3402–3413 (2023)
- Bulat, A., Tzimiropoulos, G.: How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In: International Conference on Computer Vision (2017)
- Cao, C., Simon, T., Kim, J.K., Schwartz, G., Zollhoefer, M., Saito, S.S., Lombardi, S., Wei, S.E., Belko, D., Yu, S.I., Sheikh, Y., Saragih, J.: Authentic volumetric avatars from a phone scan. ACM Trans. Graph. 41(4) (jul 2022)
- Cao, C., Weng, Y., Zhou, S., Tong, Y., Zhou, K.: Facewarehouse: A 3d facial expression database for visual computing. In: IEEE Transactions on Visualization and Computer Graphics. vol. 20, pp. 413–425 (2014)
- Chan, E., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5795–5805 (2020)
- Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., Mello, S.D., Gallo, O., Guibas, L., Tremblay, J., Khamis, S., Karras, T., Wetzstein, G.: Efficient geometryaware 3D generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16102–16112 (2022)
- Chen, X., Deng, Y., Wang, B.: Mimic3d: Thriving 3d-aware gans via 3d-to-2d imitation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Chen, Y., Wang, L., Li, Q., Xiao, H., Zhang, S., Yao, H., Liu, Y.: Monogaussianavatar: Monocular gaussian point-based head avatar. In: ACM SIGGRAPH 2023 Conference Proceedings (2024)
- Deng, Y., Yang, J., Xiang, J., Tong, X.: Gram: Generative radiance manifolds for 3d-aware image generation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 10663–10673 (2021)
- Gafni, G., Thies, J., Zollhofer, M., Niessner, M.: Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8645–8654 (June 2021)

- 16 Xu et al.
- Gao, X., Zhong, C., Xiang, J., Hong, Y., Guo, Y., Zhang, J.: Reconstructing personalized semantic facial nerf models from monocular video. ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 41(6) (2022)
- Gerig, T., Forster, A., Blumer, C., Egger, B., Lüthi, M., Schönborn, S., Vetter, T.: Morphable face models - an open framework. In: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018). pp. 75–82 (2017)
- Giebenhain, S., Kirschstein, T., Georgopoulos, M., Rünz, M., Agapito, L., Nießner, M.: Learning neural parametric head models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Giebenhain, S., Kirschstein, T., Georgopoulos, M., Rünz, M., Agapito, L., Nießner, M.: Mononphm: Dynamic head reconstruction from monocular videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Grassal, P.W., Prinzler, M., Leistner, T., Rother, C., Nießner, M., Thies, J.: Neural head avatars from monocular rgb videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18632–18643 (June 2022)
- Gu, J., Liu, L., Wang, P., Theobalt, C.: Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In: International Conference on Learning Representations (2022)
- Hong, Y., Peng, B., Xiao, H., Liu, L., Zhang, J.: Headnerf: A real-time nerf-based parametric head model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20374–20384 (June 2022)
- Hu, L., Zhang, H., Zhang, Y., Zhou, B., Liu, B., Zhang, S., Nie, L.: Gaussianavatar: Towards realistic human avatar modeling from a single video via animatable 3d gaussians. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics 42(4) (July 2023)
- Khakhulin, T., Sklyarova, V., Lempitsky, V., Zakharov, E.: Realistic one-shot mesh-based head avatars. In: Proceedings of the European Conference on Computer Vision (ECCV) (2022)
- Kirschstein, T., Giebenhain, S., Nießner, M.: Diffusionavatars: Deferred diffusion for high-fidelity 3d head avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Kirschstein, T., Qian, S., Giebenhain, S., Walter, T., Nießner, M.: Nersemble: Multi-view radiance field reconstruction of human heads. ACM Trans. Graph. 42(4) (jul 2023)
- 25. Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a model of facial shape and expression from 4d scans. ACM Trans. Graph. **36**(6) (nov 2017)
- 26. Li, X., De Mello, S., Liu, S., Nagano, K., Iqbal, U., Kautz, J.: Generalizable one-shot neural head avatar. NeurIPS (2023)
- Li, Z., Zheng, Z., Wang, L., Liu, Y.: Animatable gaussians: Learning posedependent gaussian maps for high-fidelity human avatar modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Lin, C.Z., Nagano, K., Kautz, J., Chan, E.R., Iqbal, U., Guibas, L., Wetzstein, G., Khamis, S.: Single-shot implicit morphable faces with consistent texture parameterization. In: ACM SIGGRAPH 2023 Conference Proceedings (2023)

- Lin, S., Ryabtsev, A., Sengupta, S., Curless, B., Seitz, S., Kemelmacher-Shlizerman, I.: Real-time high-resolution background matting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (Jun 2021)
- Lombardi, S., Simon, T., Schwartz, G., Zollhoefer, M., Sheikh, Y., Saragih, J.: Mixture of volumetric primitives for efficient neural rendering. ACM Trans. Graph. 40(4) (jul 2021)
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIGGRAPH Asia) 34(6), 248:1–248:16 (Oct 2015)
- Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In: 3DV (2024)
- Ma, S., Weng, Y., Shao, T., Zhou, K.: 3d gaussian blendshapes for head avatar animation. In: ACM SIGGRAPH 2023 Conference Proceedings (2024)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
- Or-El, R., Luo, X., Shan, M., Shechtman, E., Park, J.J., Kemelmacher-Shlizerman, I.: Stylesdf: High-resolution 3d-consistent image and geometry generation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 13493–13503 (2021)
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3D hands, face, and body from a single image. In: Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 10975–10985 (2019)
- Qian, S., Kirschstein, T., Schoneveld, L., Davoli, D., Giebenhain, S., Nießner, M.: Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Qin, M., Liu, Y., Xu, Y., Zhao, X., Liu, Y., Wang, H.: High-fidelity 3d head avatars reconstruction through spatially-varying expression conditioned neural radiance field. In: AAAI Conference on Artificial Intelligence (2023)
- Saito, S., Schwartz, G., Simon, T., Li, J., Nam, G.: Relightable gaussian codec avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- 40. Shao, Z., Wang, Z., Li, Z., Wang, D., Lin, X., Zhang, Y., Fan, M., Wang, Z.: SplattingAvatar: Realistic Real-Time Human Avatars with Mesh-Embedded Gaussian Splatting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Shen, T., Gao, J., Yin, K., Liu, M.Y., Fidler, S.: Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In: Advances in Neural Information Processing Systems (NeurIPS) (2021)
- Sun, J., Wang, X., Shi, Y., Wang, L., Wang, J., Liu, Y.: Ide-3d: Interactive disentangled editing for high-resolution 3d-aware portrait synthesis. ACM Transactions on Graphics (TOG) 41(6), 1–10 (2022)
- Sun, J., Wang, X., Wang, L., Li, X., Zhang, Y., Zhang, H., Liu, Y.: Next3d: Generative neural texture rasterization for 3d-aware head avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)

- 18 Xu et al.
- 44. Wang, D., Chandran, P., Zoss, G., Bradley, D., Gotardo, P.: Morf: Morphable radiance fields for multiview neural head modeling. In: ACM SIGGRAPH 2022 Conference Proceedings. SIGGRAPH '22, Association for Computing Machinery, New York, NY, USA (2022)
- 45. Wang, J., Xie, J.C., Li, X., Xu, F., Pun, C.M., Gao, H.: Gaussianhead: High-fidelity head avatars with learnable gaussian derivation (2024)
- 46. Wang, K., Wu, Q., Song, L., Yang, Z., Wu, W., Qian, C., He, R., Qiao, Y., Loy, C.C.: Mead: A large-scale audio-visual dataset for emotional talking-face generation. In: Proceedings of the European Conference on Computer Vision (ECCV) (Augest 2020)
- 47. Wang, L., Chen, Z., Yu, T., Ma, C., Li, L., Liu, Y.: Faceverse: a fine-grained and detail-controllable 3d face morphable model from a hybrid dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (Jun 2022)
- Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering (2024)
- 49. Wu, S., Yan, Y., Li, Y., Cheng, Y., Zhu, W., Gao, K., Li, X., Zhai, G.: Ganhead: Towards generative animatable neural head avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 437–447 (2023)
- Wu, Y., Deng, Y., Yang, J., Wei, F., Qifeng, C., Tong, X.: Anifacegan: Animatable 3d-aware face image generation for video avatars. In: Advances in Neural Information Processing Systems (2022)
- Wu, Y., Xu, S., Xiang, J., Wei, F., Chen, Q., Yang, J., Tong, X.: Aniportraitgan: Animatable 3d portrait generation from 2d image collections. In: SIGGRAPH Asia 2023 Conference Proceedings (2023)
- Xiang, J., Yang, J., Deng, Y., Tong, X.: Gram-hd: 3d-consistent image generation at high resolution with generative radiance manifolds. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) pp. 2195–2205 (2022)
- Xiang, J., Gao, X., Guo, Y., Zhang, J.: Flashavatar: High-fidelity head avatar with efficient gaussian embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Xu, Y., Chen, B., Li, Z., Zhang, H., Wang, L., Zheng, Z., Liu, Y.: Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Xu, Y., Wang, L., Zhao, X., Zhang, H., Liu, Y.: Avatarmav: Fast 3d head avatar reconstruction using motion-aware neural voxels. In: ACM SIGGRAPH 2023 Conference Proceedings (2023)
- Xu, Y., Zhang, H., Wang, L., Zhao, X., Han, H., Guojun, Q., Liu, Y.: Latentavatar: Learning latent expression code for expressive neural head avatar. In: ACM SIGGRAPH 2023 Conference Proceedings (2023)
- 57. Yang, Z., Yang, H., Pan, Z., Zhu, X., Zhang, L.: Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting (2023)
- 58. Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction (June 2023)
- Yenamandra, T., Tewari, A., Bernard, F., Seidel, H., Elgharib, M., Cremers, D., Theobalt, C.: i3dmm: Deep implicit 3d morphable model of human heads. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2021)

- Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 586–595 (June 2018)
- Zhao, X., Wang, L., Sun, J., Zhang, H., Suo, J., Liu, Y.: Havatar: High-fidelity head avatar via facial model conditioned neural radiance field. ACM Trans. Graph. (oct 2023)
- Zheng, Y., Abrevaya, V.F., Bühler, M.C., Chen, X., Black, M.J., Hilliges, O.: I m avatar: Implicit morphable head avatars from videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13535–13545 (June 2022)
- Zheng, Y., Yifan, W., Wetzstein, G., Black, M.J., Hilliges, O.: Pointavatar: Deformable point-based head avatars from videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Zhuang, Y., Zhu, H., Sun, X., Cao, X.: Mofanerf: Morphable facial neural radiance field. In: Proceedings of the European Conference on Computer Vision (ECCV) (2022)
- 65. Zielonka, W., Bolkart, T., Thies, J.: Instant volumetric head avatars (June 2023)