

Gaussian Head Avatar: Ultra High-fidelity Head Avatar via Dynamic Gaussians

Yuelang Xu¹, Bengwang Chen¹, Zhe Li¹, Hongwen Zhang², Lizhen Wang¹, Zerong Zheng³, Yebin Liu^{1†}

¹ Department of Automation, Tsinghua University ² Beijing Normal University ³ NNKosmos Technology




Figure 1. Gaussian head avatar achieves ultra high-fidelity image synthesis with controllable expressions at 2K resolution. The above shows different views of the synthesized avatar, and the bottom shows different identities animated by the same expression. 16 views are used during the training.

Abstract

Creating high-fidelity 3D head avatars has always been a research hotspot, but there remains a great challenge under lightweight sparse view setups. In this paper, we propose Gaussian Head Avatar represented by controllable 3D Gaussians for high-fidelity head avatar modeling. We optimize the neutral 3D Gaussians and a fully learned MLP-based deformation field to capture complex expressions. The two parts benefit each other, thereby our method can model fine-grained dynamic details while ensuring expression accuracy. Furthermore, we devise a well-designed geometry-guided initialization strategy based on implicit SDF and Deep Marching Tetrahedra for the stability and convergence of the training procedure. Experiments show our approach outperforms other state-of-the-art sparse-view methods, achieving ultra high-fidelity rendering quality at 2K resolution even under exaggerated expressions. Project page: <https://yuelangx.github.io/gaussianheadavatar>.

†Corresponding author.

1. Introduction

High-fidelity 3D human head avatar modeling is of great significance in many fields, such as VR/AR, telepresence, digital human and film production. Automatically creating high-fidelity avatars has been a research hotspot in computer vision for decades. Although some traditional head avatars [37, 39, 41, 54] can realize high-fidelity animation, they typically require accurate geometries reconstructed and tracked from dense multi-view videos, thus limiting their applications in lightweight settings.

Thanks to the Neural Radiance Fields (NeRF) [43] which show great capability of novel view synthesis in the absence of accurate geometry, recent methods [38, 57] skip the geometry reconstruction and tracking steps but directly learn high-quality NeRF-based head avatars. Other works [42, 48, 71] have verified that NeRF can be applied to either dense or sparse views, which greatly lowers the threshold for head avatar reconstruction. However, it still remains challenging for these NeRF-based approaches to synthesize high-fidelity images at 2K resolutions with pixel-

level details, including wrinkles and eyes.

To overcome this bottleneck and further improve the avatar quality, we introduce 3D Gaussian splatting [26] for 3D head avatar modeling. This is an explicit discrete representation that can be well adapted to the rasterization-based rendering pipeline. It has been verified that the 3D Gaussian representation is capable of rendering complex scenes with low computational consumption. Compared to NeRF, the reconstruction quality of static and dynamic scenes [40, 59, 67] is much better while rendering time cost has been significantly reduced. Motivated by this progress, we propose Gaussian Head Avatar, a novel representation that utilizes 3D Gaussian splatting for ultra high-fidelity head avatar modeling. Although recent 4D Gaussian works [40, 59, 66, 67] have been proposed to reconstruct dynamic scenes, all of them cannot be animated. For modeling the animatable head avatar, it is crucial but still unexplored how to effectively control the deformation of 3D Gaussians and model the dynamic appearances through expression coefficients.

Previous explicit [74] and implicit [2, 73, 76] head avatars usually formulate the facial deformation via linear blend skinning (LBS) using the skinning weights and blendshapes like the FLAME model [32]. However, such an LBS-based formulation fails to represent exaggerated and fine-grained expressions by simple linear operations, limiting the representation ability of the head avatars. Inspired by NeRsemble [28], we propose a fully learnable expression-conditioned deformation field for the 3D head Gaussians, avoiding the limited capability of the LBS-based formulation. Specifically, we input the positions of the 3D Gaussians with expression coefficients into an MLP to directly predict the displacements from the neutral expression to the target one. Similarly, we control the motion of non-face areas, such as the neck, using the head pose as the condition. 3D Gaussian-based representation has the powerful ability to reconstruct high-frequency details, enabling our method to learn accurate deformation fields. In turn, the learned accurate deformation field facilitates the dynamic Gaussian head model to fit more dynamic details. As a result, our method is able to reconstruct finer-grained dynamic details of expressive human heads.

Unfortunately, as a discretized representation, the gradients back-propagated to the 3D Gaussians cannot spread through the whole space. Thus the convergence of training heavily relies on a plausible initialization for both the geometry and the deformation field. However, simply initializing the 3D head Gaussians with a morphable template like FLAME [32] fails to model the long hairstyle and the shoulders. Hence, we further propose an efficient and well-designed geometry-guided initialization strategy. Specifically, instead of starting from stochastic Gaussians or a FLAME model, we initially optimize an implicit signed dis-

tance function (SDF) field along with a color field and a deformation MLP for modeling the basic geometry, color, and the expression-conditioned deformations of the head avatar respectively. The SDF field is converted to a mesh through Deep Marching Tetrahedra (DMTet) [52], with the color and deformation of the vertices predicted by the MLPs. Then we render the mesh and optimize them jointly under the supervision of multi-view RGB images. Finally, we use the mesh with per-vertex features from the SDF field to initialize the 3D Gaussians to lie on the basic head surface while the color and deformation MLPs are carried over to the next stage, ensuring stable training for convergence. The entire initialization process takes only around 10 minutes.

The contributions of our method can be summarized as:

- We propose Gaussian Head Avatar, a new head avatar representation that employs controllable dynamic 3D Gaussians to model expressive human head avatars, producing ultra high-fidelity synthesized images at 2K resolutions.
- For modeling high-frequency dynamic details, we employ a fully learned deformation field upon the 3D head Gaussians, which accurately model extremely complex and exaggerated facial expressions.
- We carefully design an efficient initialization strategy that leverages implicit representations to initialize the geometry and deformation, leading to efficient and robust convergence when training the Gaussian Head Avatar.

Benefiting from these contributions, our method surpasses recent state-of-the-art methods under lightweight sparse-view setups on the avatar quality by a large margin.

2. Related Works

3D Head Avatar Reconstruction. Due to the wide application value in the film and digital human industry, 3D head avatar reconstruction from multi-view images has always been a research hotspot. Traditional works [4, 8, 19, 31] reconstruct the scan geometry through multi-view stereo and then register a face mesh template to it. However, such methods usually require heavy computation. With the utilization of deep neural networks, current methods [7, 33, 61, 65] achieve very fast reconstruction, producing even more accurate geometry. Lombardi et al. [37], Bi et al. [5] and Ma et al. [41] represent the full head mesh through a deep neural network and train it with multi-view videos as supervision. However, due to the errors in geometric estimation, mesh-based head avatars typically suffer from texture blur. Therefore, some recent methods [38, 57] utilize NeRF representation [43] to synthesize novel view images without geometry reconstruction, or build NeRF on the head mesh template [39]. Furthermore, the NeRF-based methods are extended to sparse view reconstruction tasks [28, 42, 48, 71] and achieve impressive performance.

Methods which focus on generative model [6, 9, 11, 32, 46, 56, 60] are dedicated to learning general mesh

face templates from large-scale multi-view face images or 3D scans. Recently, implicit SDF-based [68] or NeRF-based [14, 23, 53, 55, 75] methods can learn full-head templates without the limitations of fixed topology, thereby better modeling complex hairstyles and glasses. Cao et al. [14] adopts a hybrid representation of local NeRF built on the mesh surface, which enables high-fidelity rendering and flexible expression control.

3D head avatars reconstruction from monocular videos is also a popular yet challenging research topic. Early methods [12, 13, 15, 24, 25, 45] optimize a morphable mesh to fit the training video. Recent methods [20, 27] leverage neural networks to learn non-rigid deformation upon 3DMM face templates [18, 32], thus can recover more dynamic details. Such methods are not flexible enough to handle complex topologies. Therefore, the latest methods explore to construct head avatar models based on implicit SDF [73], point clouds [74] or NeRF [2, 3, 16, 17, 21, 36, 47, 63, 64, 76].

Point-based Rendering. Point elements as a discrete and unstructured representation can fit geometry with arbitrary topology [69] efficiently. Recent methods [29, 30, 58] open up a differentiable rasterization pipeline, such that the point-based representation is widely used in multi-view reconstruction tasks. Aliev et al. [1] and Ruckert et al. [49] propose to first render the feature map, which is transferred to the images through a convolutional renderer. Xu et al. [62] use neural point cloud associated with neural features to model a NeRF. Recently, 3D Gaussian splatting [26] shows its superior performance, beating NeRF in both novel view synthesis quality and rendering speed. Some approaches [34, 40, 50, 59, 66, 67, 72, 72] extend Gaussian representation to dynamic scene reconstruction. However, these methods can not be migrated to the head avatar reconstruction tasks.

3. Overview

The pipeline of the reconstruction of Gaussian Head Avatar is illustrated in Fig. 2, including the initialization stage and the training stage of Gaussian Head Avatar. Before the beginning of the pipeline, we remove the background [35] of each image and jointly estimate the 3DMM model [18], 3D facial landmarks and the expression coefficients for each frame. In the initialization stage (Sec. 4.3), we reconstruct an SDF-based neutral geometry, and optimize a deformation MLP and a color MLP from the training data as the guidance model. Next, we extract the neutral mesh through DMTet to initialize the neutral Gaussians while the deformation and color MLPs are also inherited from the initialization stage. In the training stage of Gaussian Head Avatar (Sec. 4.2), we deform the neutral Gaussians to the target expression through the dynamic generator given the driving expression coefficients as the condition. Finally, given a camera view, the expressive Gaussians are rendered to a

feature map, which is then fed into the convolutional super resolution network to generate high-resolution avatar images. The whole model is optimized under the supervision of multi-view RGB videos.

4. Method

4.1. Avatar Representation

Generally, the static 3D Gaussians [26] with N points are represented by their positions X , the multi-channel color C , the rotation Q , scale S and opacity A . The rotation Q is represented in the form of quaternion. Subsequently, the Gaussians can be rasterized and rendered to the multi-channel image I given the camera parameters μ . This process can be formulated as:

$$I = \mathcal{R}(X, C, Q, S, A; \mu). \quad (1)$$

Our task is to reconstruct a dynamic head avatar controlled by expression coefficients. Therefore, we formulate the head avatar as dynamic 3D Gaussians conditioned on expressions. To handle the dynamic changes, we input the expression coefficients with head pose to the head avatar model and output the position and other attributes of the Gaussians as above.

Specifically, we first construct a canonical neutral Gaussian model with expression-independent attributes: $\{\mathbf{X}_0, \mathbf{F}_0, \mathbf{Q}_0, \mathbf{S}_0, \mathbf{A}_0\}$, which are fully optimizable. $\mathbf{X}_0 \in \mathbb{R}^{N \times 3}$ denotes the positions of the Gaussians with a neutral expression in the canonical space. $\mathbf{F}_0 \in \mathbb{R}^{N \times 128}$ denotes the point-wise feature vectors as their intrinsic properties. $\mathbf{Q}_0 \in \mathbb{R}^{N \times 4}$, $\mathbf{S}_0 \in \mathbb{R}^{N \times 3}$ and $\mathbf{A}_0 \in \mathbb{R}^{N \times 1}$ denotes the neutral rotation, scale and opacity respectively. Note that we do not define the neutral color, but directly predict expression-dependent dynamic color from the point-wise feature vectors \mathbf{F}_0 . Then, we construct an MLP-based expression conditioned dynamic generator Φ to generate all the extra dynamic changes to the neutral model. Overall, the whole Gaussian head avatar can be formulated as:

$$\{X, C, Q, S, A\} = \Phi(\mathbf{X}_0, \mathbf{F}_0, \mathbf{Q}_0, \mathbf{S}_0, \mathbf{A}_0; \theta, \beta), \quad (2)$$

with θ denoting expression coefficients and β denoting the head pose. During the training, we optimize all the parameters of the dynamic generator Φ and the neutral Gaussian model $\{\mathbf{X}_0, \mathbf{F}_0, \mathbf{Q}_0, \mathbf{S}_0, \mathbf{A}_0\}$, which are highlighted in bold in the following.

Next, we explain the process of adding expression-related changes to the neutral Gaussian model through the dynamic generator Φ as described in Eqn. 2 in detail.

Positions X' of the Gaussians. Expressions bring about the geometric deformation of the neutral model, which is modeled as the displacements of the Gaussian points.

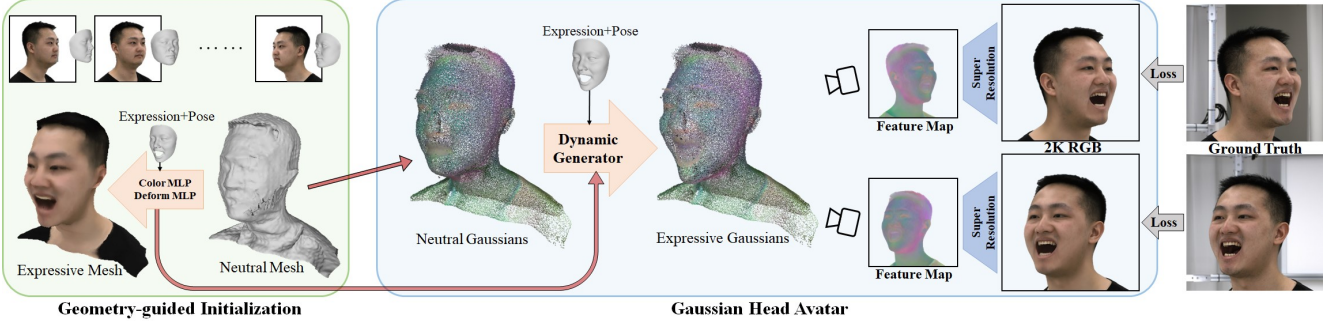


Figure 2. The overview of the Gaussian Head Avatar rendering and reconstruction. We first optimize the guidance model including a neutral mesh, a deformation MLP and a color MLP in the Initialization stage. Then we use them to initialize the neutral Gaussians and the dynamic generator. Finally, 2K RGB images are synthesized through differentiable rendering and the super-resolution network. The Gaussian Head Avatar are trained under the supervision of multi-view RGB videos.

Specifically, we predict the displacements respectively controlled by the expression and the head pose in the canonical space through two different MLPs: $\mathbf{f}_{def}^{exp} \in \Phi$ and $\mathbf{f}_{def}^{pose} \in \Phi$. Then, we add them to the neutral positions.

$$\begin{aligned} X' &= \mathbf{X}_0 + \lambda_{exp}(\mathbf{X}_0) \mathbf{f}_{def}^{exp}(\mathbf{X}_0, \theta) \\ &\quad + \lambda_{pose}(\mathbf{X}_0) \mathbf{f}_{def}^{pose}(\mathbf{X}_0, \beta). \end{aligned} \quad (3)$$

$\lambda_{exp}(\cdot)$ and $\lambda_{pose}(\cdot)$ represent the extent to which the points are affected by the expression or the head pose respectively. Without decoupling but just using the expression coefficients as the global condition [16] which also controls the shoulders and upper body, will produce jittering results during the animation. Here, we assume that the Gaussian points closer to 3D landmarks are more affected by the expression coefficients and less affected by the head pose, while the opposite is true for the Gaussian points far away. Specifically, The 3D landmarks \mathbf{P}_0 of the canonical model are first estimated through the 3DMM model in the data preprocessing and then optimized in the initialization stage 4.3. Then for each Gaussian point, we calculate the above weight $\lambda_{exp}(\cdot)$ and $\lambda_{pose}(\cdot)$ as follows:

$$\lambda_{exp}(x) = \begin{cases} 1, & dist(x, \mathbf{P}_0) < t_1 \\ \frac{t_2 - dist(x, \mathbf{P}_0)}{t_2 - t_1}, & dist(x, \mathbf{P}_0) \in [t_1, t_2] \\ 0, & dist(x, \mathbf{P}_0) > t_2 \end{cases}$$

with $\lambda_{pose}(x) = 1 - \lambda_{exp}(x)$. And $x \in \mathbf{X}_0$ denotes the position of one neutral Gaussian. $dist(x, \mathbf{P}_0)$ denotes the minimum distance from the point x to the 3D landmarks \mathbf{P}_0 . $t_1 = 0.15$ and $t_2 = 0.25$ are predefined hyperparameters when the length of the head is set to approximately 1.

Color C' of the Gaussians. Modeling the dynamic details typically requires dynamic color that changes with expressions. As we do not pre-define the neutral value in Eqn. 2, the color are directly predict by two color MLPs:

$$\mathbf{f}_{col}^{exp} \in \Phi \text{ and } \mathbf{f}_{col}^{pose} \in \Phi:$$

$$\begin{aligned} C' &= \lambda_{exp}(\mathbf{X}_0) \mathbf{f}_{col}^{exp}(\mathbf{F}_0, \theta) \\ &\quad + \lambda_{pose}(\mathbf{X}_0) \mathbf{f}_{col}^{pose}(\mathbf{F}_0, \beta). \end{aligned} \quad (4)$$

Rotation, Scale and Opacity $\{Q', S', A'\}$ of the Gaussians. These three attributes also dynamic, thereby modeling some detailed expressions-related appearance changes. Here, we just use another two attribute MLPs $\mathbf{f}_{att}^{exp} \in \Phi$ and $\mathbf{f}_{att}^{pose} \in \Phi$ to predict their shift from the neutral value.

$$\begin{aligned} \{Q', S', A'\} &= \{\mathbf{Q}_0, \mathbf{S}_0, \mathbf{A}_0\} \\ &\quad + \lambda_{exp}(\mathbf{X}_0) \mathbf{f}_{att}^{exp}(\mathbf{F}_0, \theta) \\ &\quad + \lambda_{pose}(\mathbf{X}_0) \mathbf{f}_{att}^{pose}(\mathbf{F}_0, \beta). \end{aligned} \quad (5)$$

Finally, we apply rigid rotations and translations $T(\cdot)$ to the Gaussians, transforming them from the canonical space to the world space. Note, the transformation is only implemented for directional variables: $\{X', Q'\}$, while the multi-channel color, the scale and the opacity $\{C', S', A'\}$ are not directional thus remain unchanged.

$$\{X, Q\} = T(\{X', Q'\}, \beta), \quad (6)$$

$$\{C, S, A\} = \{C', S', A'\}. \quad (7)$$

4.2. Training

In this part, we explain the training pipeline of the Gaussian head avatar 4.1 and the loss function. In each iteration, we first generate the expression conditioned 3D Gaussians as Eqn. 2. Then we render a 32-channel image with 512 resolution $I_C \in \mathbb{R}^{512 \times 512 \times 32}$ referring to Eqn. 1. After that we feed the image to a super resolution network Ψ to generate a 2048 resolution RGB image $I_{hr} \in \mathbb{R}^{2048 \times 2048 \times 3}$, such that more details are recovered and noise caused by uneven ambient light or camera chromatic aberration in the training data will be filtered out [49, 64].

During training, we jointly optimize all the learnable parameters mentioned above in bold, including the neutral

Gaussians: $\{\mathbf{X}_0, \mathbf{F}_0, \mathbf{Q}_0, \mathbf{S}_0, \mathbf{A}_0\}$, the dynamic generator: $\{\mathbf{f}_{col}^{exp}, \mathbf{f}_{col}^{pose}, \mathbf{f}_{def}^{exp}, \mathbf{f}_{def}^{pose}, \mathbf{f}_{att}^{exp}, \mathbf{f}_{att}^{pose}\}$, and the super resolution network Ψ . For the loss function, we only use the foreground RGB images I_{gt} as supervision to construct an L1 loss and a VGG perceptual loss [70] between the generated images I_{hr} and the ground truth I_{gt} . Besides, we encourage the first three channels of the 32-channel feature image I_C to be RGB channels, which is ensured by a L1 loss term. The total loss is:

$$\mathcal{L} = \|I_{hr} - I_{gt}\|_1 + \lambda_{vgg} VGG(I_{hr}, I_{gt}) + \lambda_{lr} \|I_{lr} - I_{gt}\|_1, \quad (8)$$

with I_{lr} denoting the first three channels of the 32-channel image I_C . We set the weights $\lambda_{vgg} = 0.1$ and $\lambda_{lr} = 1$.

4.3. Geometry-guided Initialization

Unlike neural networks, the Gaussians act as an unordered and unstructured representation. Random initialization leads to failure to converge while naively using the FLAME model to initialize will significantly reduce the reconstruction quality. In this section, we describe in detail how to optimize a mesh guidance model to provide reliable initialization for the Gaussians in Sec. 4.1.

Mesh Guidance Model. Specifically, we first construct a MLP \mathbf{f}_{sdf} to represent a signed distance field. In addition, this network will also output the corresponding feature vector of each point, which is used for predicting the point color. It can be formulated as:

$$s, \eta = \mathbf{f}_{sdf}(x), \quad (9)$$

with s denotes the SDF value, η denotes the feature vector and x denotes the point position. Then through (DMTet) [51], we can differentially extract a mesh with vertices X , per-vertex feature vectors F and its faces. We also predict the per-vertex 32-channel color as Eqn. 4 by the two color MLPs \mathbf{f}_{col}^{exp} and \mathbf{f}_{col}^{pose} . In parallel, we construct the two deformation MLPs: \mathbf{f}_{def}^{exp} and \mathbf{f}_{def}^{pose} as described in Sec. 4.1 to predict the displacements and add them to the vertex positions. This process is similar to Eqn 3 above, with the Gaussian positions \mathbf{X}_0 replaced by the vertex positions X . Finally we also apply rigid rotations and translations to the deformed mesh, transforming it to the world space and render the deformed mesh into an image I and a mask M through differentiable rasterization [44] according to the camera parameters μ .

Loss Function and Training. Next, we can construct the RGB loss and the silhouette loss to train the guidance model:

$$\mathcal{L}_{RGB} = \|I_{r,g,b} - I_{gt}\|_1, \quad (10)$$

$$\mathcal{L}_{sil} = IOU(M, M_{gt}), \quad (11)$$

with I_{gt} and M_{gt} denote the ground truth RGB image and mask, respectively. $IOU(\cdot)$ denotes Intersection over

Union metrics. Note that only the first three channels R, G, B of the 32-channel image I are supervised by the ground truth RGB images.

We also use the estimated 3D facial landmarks P_{gt} as described in Sec. 3 to provide rough guidance for the expression deformation MLP. Specifically, we input the neutral 3D landmarks \mathbf{P}_0 into the expression deformation MLP to predict the expression conditioned landmarks P :

$$P = \mathbf{P}_0 + \mathbf{f}_{def}^{exp}(\mathbf{P}_0, \theta). \quad (12)$$

Then we construct the loss function with 3D facial landmarks P_{gt} as the supervision:

$$\mathcal{L}_{def} = \|P - P_{gt}\|_2, \quad (13)$$

Besides, we introduce three constraints: (1) a regular term \mathcal{L}_{offset} to punish all non-zero displacements to prevent the two deformation MLPs from learning a global constant offset [63], (2) a regular term \mathcal{L}_{lmk} to limit the SDF value at the 3D landmarks to be close to zero, such that the landmarks are located on the surface of the mesh, (3) a Laplacian term \mathcal{L}_{lap} for maintaining the extracted mesh smooth to a certain extent. Overall, the total loss function is formulated as:

$$\mathcal{L} = \mathcal{L}_{RGB} + \lambda_{sil} \mathcal{L}_{sil} + \lambda_{def} \mathcal{L}_{def} + \lambda_{offset} \mathcal{L}_{offset} + \lambda_{lmk} \mathcal{L}_{lmk} + \lambda_{lap} \mathcal{L}_{lap}, \quad (14)$$

with λ denoting the weights of each term, which are set as follows: $\lambda_{sil} = 0.1$, $\lambda_{def} = 1$, $\lambda_{offset} = 0.01$, $\lambda_{lmk} = 0.1$ and $\lambda_{lap} = 100$. We jointly optimize the MLPs mentioned above with the neutral 3D landmarks \mathbf{P}_0 jointly until all MLPs are converged.

Parameters Transfer. Finally, we use the roughly trained mesh guidance model to initialize the Gaussian model. Specifically, we extract the neutral mesh with vertices X and per-vertex features F through DMTet, and directly assign their values to the neutral positions $\mathbf{X}_0 = X$ and the per-vertex feature vectors $\mathbf{F}_0 = F$ of the neutral Gaussians respectively. Then, we retain all the four optimized MLPs: $\{\mathbf{f}_{col}^{exp}, \mathbf{f}_{col}^{pose}, \mathbf{f}_{def}^{exp}, \mathbf{f}_{def}^{pose}\}$ for the Gaussian model. For the other neutral attributes: rotation, scale and opacity, we adopt the original initialization strategy in Gaussian Splatting [26]. And the parameters of the two attribute MLPs: $\{\mathbf{f}_{att}^{exp}, \mathbf{f}_{att}^{pose}\}$ and the super resolution network Ψ are just randomly initialized.

5. Experiments

5.1. Implementation Details

In the experiment we use 12 sets of data, 10 of which are from NeRSemble [28], and the other 2 are multi-view video data from HAvatar [71]. For the 10 identities from NeRSemble, each set contains 2500 to 3000 frames, 16




Figure 3. Qualitative comparisons of different methods on self reenactment task. From left to right: NeRFBlendShape [17], NeRFace [16], HAvatar [71] and Ours. Our method can reconstruct details like beards, teeth, etc. with high quality.

cameras are distributed about 120 degrees in front, and simultaneously capture 2K resolution video. For each identity, We use the sequences marked with "FREE" as the evaluation data, and the rest as the training data. For the 2 identities from HAvatar, each set contains 3000 frames, 8 cameras are distributed about 120 degrees in front, and 4K resolution videos are collected simultaneously. Later, we crop the face area and resize to 2K resolution.

For data preprocessing, we first remove the background [35] and extract 68 2D facial landmarks [10] for all the images. Then, for each frame, we use multi-view images to estimate the corresponding 3D landmarks, expression coefficients and head pose by fitting the Basel Face Model (BFM) [18] to the extracted 2D landmarks. Note that we define the 3D landmarks as the usual 68 landmarks with vertices indexed as multiples of 100 in the BFM vertices.

5.2. Results and Comparisons

Self Reenactment. In this section, we first compare our method with existing SOTA methods in qualitative experiments on self reenactment task. Specifically, NeRFace [16] uses a deep MLP to fit an expression conditioned dynamic NeRF. The current SOTA method HAvatar [71] introduces 3DMM template prior and uses a deep convolutional net-

work to generate a human head NeRF represented by three planes from a mesh template with expression. Note, HAvatar leverages the GAN framework using the adversarial loss function to force the network to generate details that are not view-consistent. For a fair comparison, we remove this part and use VGG perceptual loss as Sec. 4.2 instead.

Qualitative results on self reenactment task are shown in the Fig. 3. Our method can accurately reconstruct pixel-level high-frequency details such as beards, teeth, and hair. Besides, our method can achieve expression transfer more accurately, such as eye movements in the figure.

Next, we conduct a quantitative evaluation for the four methods on 5 identities and 6 cameras using the evaluation split. The evaluation metrics include: Peak Signal-to-Noise Ratio (PSNR), Structure Similarity Index (SSIM), Learned Perceptual Image Patch Similarity (LPIPS) [70] and Fréchet Inception Distance (FID) [22]. Note, we calculate FID by comparing the distribution of all the training images and all the rendered images. As the task mainly focuses on the reconstruction of the head, we use face parsing¹ to remove the body parts in the image to eliminate their impact in the experiment. As shown in Tab. 1, our method demonstrates a slight improvement in PSNR and SSIM compared with

¹<https://github.com/zllrunning/face-parsing.PyTorch>

Method	PSNR \uparrow	SSIM \uparrow	LPIPS (512) \downarrow	LPIPS (2K) \downarrow	FID (2K) \downarrow
NeRFBlendShape	25.91	0.836	0.123	0.229	54.80
NeRFace	27.14	0.849	0.147	0.234	65.11
HAvatar	27.19	0.883	0.064	0.209	31.06
Ours (w/o SR)	27.82	0.887	0.080	0.202	45.50
Ours	27.70	0.883	0.056	0.098	18.50

Table 1. Quantitative evaluation results of NeRFBlendShape [17], NeRFace [16], HAvatar [71], our method without super resolution and our full method on self reenactment task.




Figure 4. Qualitative comparisons of different methods on cross-identity reenactment task. From left to right: NeRFBlendShape [17], NeRFace [16], HAvatar [71] and Ours. Our method synthesizes high-fidelity images while ensuring the accuracy of expression transfer.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NeRFBlendShape	25.43	0.812	0.148
NeRFace	26.65	0.825	0.151
HAvatar	27.13	0.880	0.65
Ours	27.58	0.882	0.059

Table 2. Quantitative evaluation results of the other SOTA methods and our method on 3D consistency.

previous methods, and a significant improvement in LPIPS and FID, which means that our method can generate more high-frequency details.

Cross-Identity Reenactment. We also qualitatively compare our method with the above SOTA methods on cross-identity reenactment task. As shown in the results, our method is able to synthesize higher-fidelity images with more accurate expression transfer and richer emotions.

Novel View Synthesis. In this section, we show the results of novel view synthesis as shown in Fig. 5 shown. In this case, we use video data from 8 views for training and render the image at a new viewpoint. Next, we quantitatively evaluate the 3D consistency of our method and compare it with other SOTA methods mentioned above. Specif-

ically, we select the 5 identities as above and use the video data from 8 cameras for training while the other 8 hold-out cameras for evaluation. The evaluation metrics include PSNR, SSIM and LPIPS (512 resolution). The results are shown in Tab. 2. Our method outperforms other methods in 3D consistency.

5.3. Ablation Study

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
FLAME-Init	28.73	0.875	0.123
Mesh-Deform	28.83	0.874	0.116
Ours	28.94	0.876	0.108

Table 3. Quantitative evaluation results of the other two ablation baselines and ours on self reenactment task.

Ablation on Initialization Strategies. In order to verify the effectiveness of our geometry-guided initialization strategy 4.3, we compare it with the strategy to use the FLAME model for initialization (FLAME-Init). Specifically, after fitting a FLAME model through multi-view data, we first subdivide the FLAME mesh 4 times and use the neutral vertices as the positions of the neutral Gaussians. Then, the




Figure 5. Novel view synthesis results of our method. We use 8-view synchronized videos for training the avatar.




Figure 6. Ablation study on the initialization strategies: FLAME-initialization and our geometry-guided initialization. Our strategy ensures the hair strands away from the head are well reconstructed.

expression deformation MLP is optimized to learn the displacement of FLAME vertices. We set the per-vertex feature to zeros, while randomly initialize the parameters of the expression color MLP. The initialization of other variables is the same as our strategy. Qualitative results are shown in the Fig. 6. Due to the lack of initialization for the hair and shoulders in FLAME-Init, the points to model these parts are offset from nearby vertices, which leads to sparseness of the Gaussians, resulting in blurring.

Ablation on Deformation Modeling Approaches. We compare our fully learned deformation field with the previous mesh-based deformation (Mesh-Deform). Specifically, we migrate the method in INSTA [76] for controlling the NeRF deformation to our Gaussians. First we fit a 3DMM mesh template. Then, for each Gaussian point, find the closest face on the mesh, and calculate the deformation gradient to estimate the displacement. Qualitative results are shown in the Fig. 7. For some expressions that cannot be captured well by the 3DMM mesh template, our method can learn accurate deformation, thereby achieving the modeling of complex expressions. Quantitative results are shown in the Fig. 3. Our method outperforms both the two ablation baselines on PSNR, SSIM and LPIPS metrics.




Figure 7. Ablation study on the deformation modeling: mesh LBS-based deformation and our fully learned deformation. Our approach can learn complex and exaggerated expressions.

6. Discussion and Conclusion

Ethical Considerations. Our method is capable of creating artificial portrait videos, which have the potential to disseminate misinformation, influence public perceptions, and erode confidence in media sources.

Limitation. For the tongue and teeth inside the mouth or long hair, blurring is sometimes produced in our method due to the lack of tracking methods.

Conclusion. In this paper, we propose Gaussian Head Avatar, a novel representation for head avatar reconstruction, which leverages dynamic 3D Gaussians controlled by a fully learned expression deformation. Experiments demonstrate our method can synthesize ultra high-fidelity images while modeling exaggerated expressions. In addition, we propose a well-designed minute-level initialization strategy to ensure the training convergence. We believe our Gaussian Head Avatar will become the mainstream direction for head avatar reconstruction in the future.

Acknowledgement. This paper is supported by National Key R&D Program of China (2022YFF0902200), the NSFC project No.62125107, the Beijing Municipal Science & Technology Z231100005923030.

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *European Conference on Computer Vision*, pages 696–712, 2020. [3](#)
- [2] ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. Rignernf: Fully controllable neural 3d portraits. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#), [3](#)
- [3] ShahRukh Athar, Zhixin Shu, and Dimitris Samaras. Flamein-nerf: Neural control of radiance fields for free view face animation. In *IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–8, 2023. [3](#)
- [4] Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. High-quality single-shot capture of facial geometry. *ACM Trans. Graph.*, 29(4), 2010. [2](#)
- [5] Sai Bi, Stephen Lombardi, Shunsuke Saito, Tomas Simon, Shih-En Wei, Kevyn Mcphail, Ravi Ramamoorthi, Yaser Sheikh, and Jason Saragih. Deep relightable appearance models for animatable faces. *ACM Trans. Graph.*, 40(4), 2021. [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. [2](#)
- [7] Timo Bolkart, Tianye Li, and Michael J. Black. Instant multi-view head capture through learnable registration. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 768–779, 2023. [2](#)
- [8] Derek Bradley, Wolfgang Heidrich, Tiberiu Popa, and Alla Sheffer. High resolution passive facial performance capture. 29(4), 2010. [2](#)
- [9] Alan Brunton, Timo Bolkart, and Stefanie Wuhrer. Multilinear wavelets: A statistical shape space for human faces. In *Proceedings of the Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. [2](#)
- [10] Adrian Bulat and Georgios Tzimiropoulos. 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. [6](#)
- [11] Chen Cao, Yanlin Weng, Shun Zhou, Y. Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. In *IEEE Transactions on Visualization and Computer Graphics*, pages 413–425, 2014. [2](#)
- [12] Chen Cao, Derek Bradley, Kun Zhou, and Thabo Beeler. Real-time high-fidelity facial performance capture. *ACM Trans. Graph.*, 34(4), 2015. [3](#)
- [13] Chen Cao, Hongzhi Wu, Yanlin Weng, Tianjia Shao, and Kun Zhou. Real-time facial animation with image-based dynamic avatars. *ACM Trans. Graph.*, 35(4), 2016. [3](#)
- [14] Chen Cao, Tomas Simon, Jin Kyu Kim, Gabe Schwartz, Michael Zollhofer, Shun-Suke Saito, Stephen Lombardi, Shih-En Wei, Danielle Belko, Shoou-I Yu, Yaser Sheikh, and Jason Saragih. Authentic volumetric avatars from a phone scan. *ACM Trans. Graph.*, 41(4), 2022. [3](#)
- [15] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019. [3](#)
- [16] 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 (CVPR)*, pages 8645–8654, 2021. [3](#), [4](#), [6](#), [7](#)
- [17] 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 (Proceedings of SIGGRAPH Asia)*, 41(6), 2022. [3](#), [6](#), [7](#)
- [18] Thomas Gerig, Andreas Forster, Clemens Blumer, Bernhard Egger, Marcel Lüthi, Sandro Schönborn, and Thomas Vetter. Morphable face models - an open framework. pages 75–82, 2017. [3](#), [6](#)
- [19] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. *ACM Trans. Graph.*, 30(6):1–10, 2011. [2](#)
- [20] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular rgb videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18632–18643, 2022. [3](#)
- [21] Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5764–5774, 2021. [3](#)
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Neural Information Processing Systems*, 2017. [6](#)
- [23] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20374–20384, 2022. [3](#)
- [24] Liwen Hu, Shunsuke Saito, Lingyu Wei, Koki Nagano, Jaewoo Seo, Jens Fursund, Iman Sadeghi, Carrie Sun, Yen-Chun Chen, and Hao Li. Avatar digitization from a single image for real-time rendering. *ACM Trans. Graph.*, 36(6), 2017. [3](#)
- [25] Alexandru Eugen Ichim, Sofien Bouaziz, and Mark Pauly. Dynamic 3d avatar creation from hand-held video input. *ACM Trans. Graph.*, 34(4), 2015. [3](#)
- [26] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. [2](#), [3](#), [5](#)
- [27] Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. Realistic one-shot mesh-based head avatars. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [3](#)

- [28] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Trans. Graph.*, 42(4), 2023. 2, 5
- [29] Georgios Kopanas, Thomas Leimkühler, Gilles Rainer, Clément Jambon, and George Drettakis. Neural point catacaustics for novel-view synthesis of reflections. *ACM Transactions on Graphics (TOG)*, 41(6):1–15, 2022. 3
- [30] Christoph Lassner and Michael Zollhofer. Pulsar: Efficient sphere-based neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1440–1449, 2021. 3
- [31] Marc Levoy, Kari Pulli, Brian Curless, Szymon Rusinkiewicz, David Koller, Lucas Pereira, Matt Gintzton, Sean Anderson, James Davis, Jeremy Ginsberg, Jonathan Shade, and Duane Fulk. The digital michelangelo project: 3d scanning of large statues. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, page 131–144, USA, 2000. ACM Press/Addison-Wesley Publishing Co. 2
- [32] 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), 2017. 2, 3
- [33] Tianye Li, Shichen Liu, Timo Bolkart, Jiayi Liu, Hao Li, and Yajie Zhao. Topologically consistent multi-view face inference using volumetric sampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3824–3834, 2021. 2
- [34] Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [35] Shanchuan Lin, Andrey Ryabtsev, Soumyadip Sengupta, Brian Curless, Steve Seitz, and Ira Kemelmacher-Shlizerman. Real-time high-resolution background matting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3, 6
- [36] Xian Liu, Yinghao Xu, Qianyi Wu, Hang Zhou, Wayne Wu, and Bolei Zhou. Semantic-aware implicit neural audio-driven video portrait generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 3
- [37] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *ACM Trans. Graph.*, 37(4):68:1–68:13, 2018. 1, 2
- [38] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4):65:1–65:14, 2019. 1, 2
- [39] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhofer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Trans. Graph.*, 40(4), 2021. 1, 2
- [40] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis, 2023. 2, 3
- [41] Shugao Ma, Tomas Simon, Jason Saragih, Dawei Wang, Yuecheng Li, Fernando De La Torre, and Yaser Sheikh. Pixel codec avatars. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 64–73, 2021. 1, 2
- [42] Marko Mihajlovic, Aayush Bansal, Michael Zollhofer, Siyu Tang, and Shunsuke Saito. Keypointnerf: Generalizing image-based volumetric avatars using relative spatial encoding of keypoints. In *European conference on computer vision*, 2022. 1, 2
- [43] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2
- [44] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting Triangular 3D Models, Materials, and Lighting From Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8280–8290, 2022. 5
- [45] Koki Nagano, Jaewoo Seo, Jun Xing, Lingyu Wei, Zimo Li, Shunsuke Saito, Aviral Agarwal, Jens Fursund, and Hao Li. Pagan: Real-time avatars using dynamic textures. *ACM Trans. Graph.*, 37(6), 2018. 3
- [46] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance*, pages 296–301, 2009. 2
- [47] Minghan Qin, Yifan Liu, Yuelang Xu, Xiaochen Zhao, Yebin Liu, and Haoqian Wang. High-fidelity 3d head avatars reconstruction through spatially-varying expression conditioned neural radiance field. In *AAAI Conference on Artificial Intelligence*, 2023. 3
- [48] Amit Raj, Michael Zollhofer, Tomas Simon, Jason Saragih, Shunsuke Saito, James Hays, and Stephen Lombardi. Pva: Pixel-aligned volumetric avatars. In *arXiv:2101.02697*, 2020. 1, 2
- [49] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM Trans. Graph.*, 41(4), 2022. 3, 4
- [50] Ruizhi Shao, Jingxiang Sun, Cheng Peng, Zerong Zheng, Boyao Zhou, Hongwen Zhang, and Yebin Liu. Control4d: Efficient 4d portrait editing with text. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2024. 3
- [51] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 5
- [52] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [53] 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 (CVPR)*, 2023. 3
- [54] Cong Wang, Di Kang, Yanpei Cao, Linchao Bao, Ying Shan, and Song-Hai Zhang. Neural point-based volumetric avatar: Surface-guided neural points for efficient and photorealistic volumetric head avatar. In *ACM SIGGRAPH Asia 2023 Conference Proceedings*, 2023. 1
- [55] Daoye Wang, Prashanth Chandran, Gaspard Zoss, Derek Bradley, and Paulo Gotardo. Morf: Morphable radiance fields for multiview neural head modeling. In *ACM SIGGRAPH 2022 Conference Proceedings*, New York, NY, USA, 2022. Association for Computing Machinery. 3
- [56] Lizhen Wang, Zhiyuan Chen, Tao Yu, Chenguang Ma, Liang Li, and Yebin Liu. 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)*, 2022. 2
- [57] 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 (CVPR)*, pages 5704–5713, 2021. 1, 2
- [58] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. SynSin: End-to-end view synthesis from a single image. In *CVPR*, 2020. 3
- [59] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering, 2023. 2, 3
- [60] Sijing Wu, Yichao Yan, Yunhao Li, Yuhao Cheng, Wenhan Zhu, Ke Gao, Xiaobo Li, and Guangtao Zhai. Ganhead: Towards generative animatable neural head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 437–447, 2023. 2
- [61] Yunze Xiao, Hao Zhu, Haotian Yang, Zhengyu Diao, Xiangju Lu, and Xun Cao. Detailed facial geometry recovery from multi-view images by learning an implicit function. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022. 2
- [62] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5438–5448, 2022. 3
- [63] 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*, 2023. 3, 5
- [64] Yuelang Xu, Hongwen Zhang, Lizhen Wang, Xiaochen Zhao, Huang Han, Qi Guojun, and Yebin Liu. Latentavatar: Learning latent expression code for expressive neural head avatar. In *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. 3, 4
- [65] Kai Yang, Hong Shang, Tianyang Shi, Xinghan Chen, Jingkai Zhou, Zhongqian Sun, and Wei Yang. Asm: Adaptive skinning model for high-quality 3d face modeling supplementary material. 2021. 2
- [66] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction, 2023. 2, 3
- [67] Zeyu Yang, Hongye Yang, Zijie Pan, Xiatian Zhu, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting, 2023. 2, 3
- [68] T Yenamandra, A Tewari, F Bernard, HP Seidel, M Elgharib, D Cremers, and C Theobalt. i3dmm: Deep implicit 3d morphable model of human heads. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [69] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH ASIA)*, 38(6), 2019. 3
- [70] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. 5, 6
- [71] Xiaochen Zhao, Lizhen Wang, Jingxiang Sun, Hongwen Zhang, Jinli Suo, and Yebin Liu. Havatar: High-fidelity head avatar via facial model conditioned neural radiance field. *ACM Trans. Graph.*, 2023. Just Accepted. 1, 2, 5, 6, 7
- [72] Shunyuang Zheng, Boyao Zhou, Ruizhi Shao, Boning Liu, Shengping Zhang, Liqiang Nie, and Yebin Liu. Gps-gaussian: Generalizable pixel-wise 3d gaussian splatting for real-time human novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [73] 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 (CVPR)*, pages 13535–13545, 2022. 2, 3
- [74] 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 (CVPR)*, 2023. 2, 3
- [75] Yiyu Zhuang, Hao Zhu, Xusen Sun, and Xun Cao. Mofanerf: Morphable facial neural radiance field. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 3
- [76] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Instant volumetric head avatars, 2022. 2, 3, 8