

# OMG-Avatar: One-shot Multi-LOD Gaussian Head Avatar

Jianqiang Ren, Lin Liu, Steven Hoi  
Tongyi Lab, Alibaba Group

## Abstract

We propose **OMG-Avatar**, a novel *One-shot* method that leverages a **Multi-LOD (Level-of-Detail) Gaussian representation** for animatable 3D head reconstruction from a single image in 0.2s. Our method enables LOD head avatar modeling using a unified model that accommodates diverse hardware capabilities and inference speed requirements. To capture both global and local facial characteristics, we employ a transformer-based architecture for global feature extraction and projection-based sampling for local feature acquisition. These features are effectively fused under the guidance of a depth buffer, ensuring occlusion plausibility. We further introduce a coarse-to-fine learning paradigm to support Level-of-Detail functionality and enhance the perception of hierarchical details. To address the limitations of 3DMMs in modeling non-head regions such as the shoulders, we introduce a multi-region decomposition scheme in which the head and shoulders are predicted separately and then integrated through cross-region combination. Extensive experiments demonstrate that **OMG-Avatar** outperforms state-of-the-art methods in reconstruction quality, reenactment performance, and computational efficiency.

## 1. Introduction

Reconstructing an animatable 3D head avatar from a single image is a crucial and rapidly evolving research area in computer vision and graphics. This technology has great potential for applications across various domains, including the game and video production industries, virtual meetings, and the emerging Metaverse. To facilitate the widespread adoption of this technology, several key features are essential: high-efficiency reconstruction and inference, rich facial details, and precise controllability over expressions and head poses. In recent years, numerous methods have been developed to tackle this task, which can be divided into 2D-based and 3D-based approaches.

Early 2D-based methods [19, 44, 52] predict deformation flows to warp the latent features of a source portrait and employ GANs (Generative Adversarial Networks) [18] to synthesize the reenacted output. With the rise of la-

tent diffusion models [43], recent approaches [24, 49, 70] have adopted cross-attention mechanisms conditioned on driving signals, achieving superior image quality and better appearance preservation. However, both GAN-based and diffusion-based methods require substantial computational resources, limiting their applicability in real-time scenarios. Furthermore, due to the lack of 3D constraints, these approaches often struggle to maintain multi-view consistency under large pose or viewpoint variations.

In the 3D avatar domain, NeRF and Gaussian Splatting have emerged as prominent approaches due to their high-quality representation and rendering capabilities. Compared to NeRF-based methods [2, 15, 28, 32, 36, 39, 62, 72], 3D Gaussian Splatting has become the prevailing choice owing to its significantly fast rendering speed. Unlike approaches that require extensive per-individual optimization or multi-view inputs [46, 53, 54, 58, 71], recent methods such as [6, 21, 30, 65] generate 3D avatars from a single image, significantly enhancing generalization capabilities. Despite these advancements, key challenges persist in efficiency, computational cost, and compatibility. For instance, GAGAvatar [6] primarily utilizes Gaussian points sampled from the dual-lifting planes of backgrounds, leading to redundancy and inefficiency. LAM [21] employs subdivided vertices as queries in its cross-attention mechanism, and its computational complexity grows exponentially with the number of subdivision levels. Avat3r [30] relies on an additional 3D GAN for 3D lifting to enable single-image reconstruction, which introduces error accumulation. Moreover, they all fail to dynamically adjust computational complexity to accommodate diverse hardware capabilities and inference speed requirements.

To address these challenges, we propose a novel one-shot method for generating multi-LOD 3D Gaussian head avatars using a unified model. Rather than directly performing 2D-to-3D feature mapping on high-resolution meshes, which is computationally expensive, our approach extracts both global and local features from low-resolution meshes and progressively refines them to high-resolution representations through subdivision operations during the training phase, which also enables dynamic multi-LOD rendering at runtime. By leveraging an effective occlusion-aware feature

fusion mechanism, our model delivers superior reconstruction quality while significantly reducing computational cost and the number of Gaussians compared to existing methods. Additionally, to improve representation in non-head regions, we independently model the head and shoulder based on shared features, improving the completeness of the generated avatars.

The main contributions are summarized as follows:

- We present OMG-Avatar, a novel framework that reconstructs animatable Gaussian head avatars from a single image in just 0.2 seconds. It enables dynamic multi-level-of-detail (multi-LOD) rendering while achieving state-of-the-art reconstruction quality and real-time inference at 85 FPS using fewer Gaussians than prior methods.
- We propose a coarse-to-fine learning strategy that progressively refines both transformer-based global features and projection-sampled local features through multi-level subdivisions, followed by an occlusion-aware feature fusion mechanism guided by the depth buffer. These components significantly enhance the fidelity and completeness of the generated avatars.
- Comprehensive experiments on two large-scale datasets demonstrate that OMG-Avatar surpasses state-of-the-art methods in reconstruction accuracy, reenactment fidelity, and computational efficiency.

## 2. Related Work

**2D Talking Head Generation.** Early 2D approaches [5, 51, 63, 73] employ generative adversarial networks (GANs) [18, 23, 25] and incorporate driving expression features for controllable portrait synthesis. Subsequent methods [14, 19, 22, 42, 44, 64] adopt deformation-based frameworks, representing expressions and poses as warping fields to deform the source image. These methods usually struggle with large pose and expression variations due to a lack of 3D awareness. To address this, some [17, 31, 37, 41] integrate 3D Morphable Models (3DMMs) into 2D pipelines, but they still lack support for viewpoint control and free-viewpoint rendering. Recent diffusion-based approaches [8, 49, 56, 70] further improve visual quality and spatio-temporal coherence. However, their high computational cost hinders real-time performance, such as in video conferencing or live chat applications.

**3D Head Avatar Generation.** Traditional 3D head avatars typically rely on 3D Morphable Models (3DMMs) for mesh reconstruction [27, 57], which often fail to capture fine geometric details. In contrast, NeRF-based approaches [1, 3, 15, 16, 20, 28, 29, 39, 40, 50, 68, 69, 72, 74] have significantly improved reconstruction accuracy and detail representation, and several efficient one-shot NeRF methods have also been proposed [7, 32, 35, 59, 62]. However, these NeRF-based methods still face challenges in achieving real-time rendering performance. This limitation is addressed

by 3D Gaussian Splatting [26], which offers faster rendering while maintaining high visual quality. Unlike earlier Gaussian-based methods [46, 53, 58, 71] that rely on extensive individual-specific optimization or multi-view inputs, recent works [6, 21, 30, 65] propose one-shot reconstruction frameworks for head avatars. However, these methods still face challenges such as redundant Gaussian utilization, excessive computational demands, and incomplete head representation, which we aim to address in this paper.

**Hierarchical Gaussian Representation.** The hierarchical representation has been widely adopted for efficiently modeling multi-scale and structured data. HiSplat [47] introduces a hierarchical Gaussian splatting framework for sparse-view reconstruction, employing coarse-grained Gaussians to capture large-scale geometry and fine-grained Gaussians to represent texture details. LODAvatar [13] integrates level-of-detail (LOD) control into Gaussian avatars through hierarchical embedding, achieving a favorable trade-off between visual quality and computational cost. However, it does not support facial animation. Gaussian-Heads [48] leverages a hierarchical representation to model complex facial expressions and head dynamics, but requires over 20 hours of training and relies on a multi-view camera rig. In contrast, our work employs a hierarchical representation to enable both efficient modeling and dynamic LOD rendering at runtime, supporting real-time, animatable avatars without the need for specialized capture setups.

## 3. Method

Fig. 1 illustrates the overall pipeline of our method. Given a source image, we first extract both local and identity features using DINOv2 [38] and estimate the 3D head mesh via a 3DMM modeler. The Hierarchical Projection-based Feature Sampling (HPFS) module then projects the head mesh onto the image plane to sample local features at the corresponding coordinates. Concurrently, the global feature is obtained through cross-attention, where FLAME positional embeddings serve as queries. Subsequently, the Occlusion-Aware Feature Fusion(OAFF) module fuses the global and local features under the guidance of depth buffers, ensuring spatial coherence and occlusion plausibility. Features from non-head regions are further integrated with head-related features to jointly predict Gaussian attributes for splatting rendering. Finally, a neural renderer generates a refined output image based on the coarse splatted feature maps. During training, we progressively subdivide the meshes along with their corresponding features, enabling the network to capture hierarchical details in a coarse-to-fine manner, which enhances both training stability and reconstruction accuracy. The details of each module are explained in the subsequent sections.

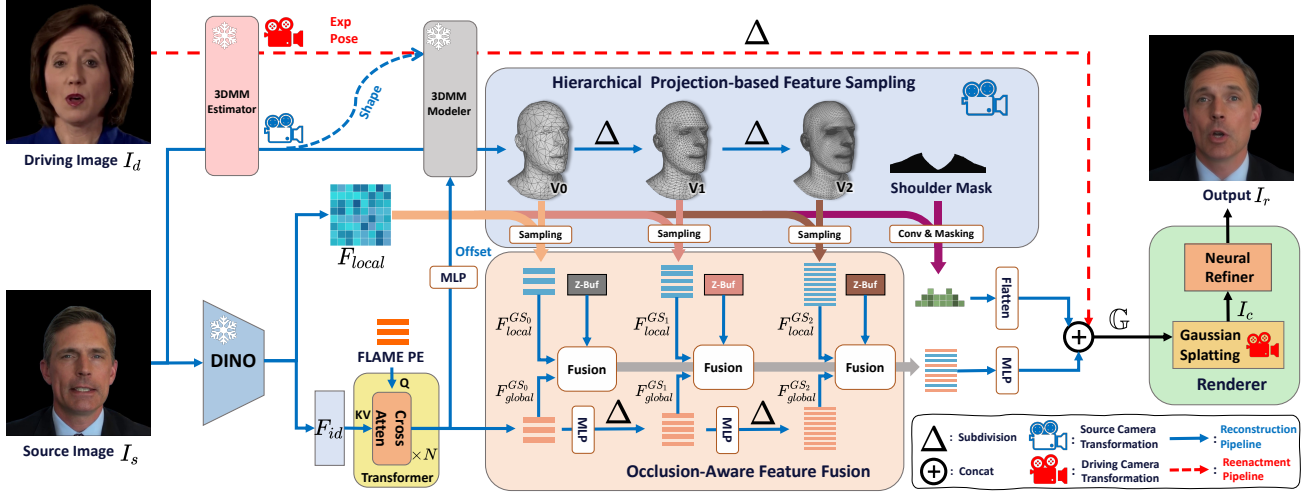


Figure 1. The overall pipeline of OMG-Avatar framework. Our method extracts global features via cross-attention and local details via projection-based sampling, which are fused under the guidance of depth buffers. A coarse-to-fine strategy is proposed to facilitate hierarchical detail perception. The head and shoulder are predicted separately using shared features and then combined for rendering.

### 3.1. Hierarchical Global-Local Feature Extraction

For the task of generating a 3D head model from a single image, the core objective is to establish a 2D-to-3D feature mapping mechanism that transforms image features into 3D spatial features. To incorporate statistical priors on head geometry, we employ the FLAME [31] model as the 3D head representation, which comprises  $N_0 = 5023$  vertices. We leverage DINOv2 [38] to extract both local features  $F_{local}$  and identity features  $F_{id}$  from the source image  $I_s$  following [6]. For the  $F_{id}$ , we assign a learnable positional encoding to each vertex of FLAME as a query, and employ multiple cross-attention blocks to extract global features  $F_{global}^{GS_0}$ . The  $F_{global}^{GS_0}$  is then used to predict vertex offsets via an MLP  $\Phi_{offset}$  to improve the precision of the estimated head mesh  $T_p$  via a 3DMM modeler:

$$T_p(\vec{\beta}, \vec{\theta}, \vec{\psi}) = \bar{T} + B_S(\vec{\beta}; S) + B_P(\vec{\theta}; P) + B_E(\vec{\psi}; E) + \Phi_{offset}(F_{global}^{GS_0}), \quad (1)$$

where  $\bar{T}$  is the template mesh, and  $B_S$ ,  $B_P$  and  $B_E$  represent shape, pose, and expression blendshapes respectively. The initial head vertices  $V_0$  are obtained using a standard skinning function  $W$ :

$$V_0 = W(T_p(\vec{\beta}, \vec{\theta}, \vec{\psi}), \mathbf{J}(\vec{\beta}), \vec{\theta}, \mathcal{W}), \quad (2)$$

where  $W$  rotates the morphed mesh  $T_p$  around joints  $\mathbf{J}$  and smooths it using blendweights  $\mathcal{W}$ . We project  $V_0$  into the image space to obtain the corresponding pixel coordinates for each vertex and perform bilinear sampling on  $F_{local}$  to extract per-vertex features, denoted as  $F_{local}^{GS_0}$ . As noted in LAM [21], the original number of  $V_0$  is insufficient for detailed modeling, so we introduce a coarse-to-fine strategy

that progressively subdivides the mesh  $V_k$  and its associated global features  $F_{global}^{GS_k}$  during training:

$$F_{global}^{GS_{k+1}}, V_{k+1} = \Delta(\Phi_k(F_{global}^{GS_k}), V_k), \quad 0 \leq k \leq K, \quad (3)$$

where  $\Delta$  denotes the mesh subdivision operation,  $\Phi_k$  is an MLP network, and  $k$  indicates the subdivision level. With the refined vertices  $V_k$ , the corresponding local feature can be obtained as:

$$F_{local}^{GS_k} = \text{Sampling}(P(V_k), F_{local}), \quad 0 \leq k \leq K, \quad (4)$$

where  $P$  is the camera projective transformation (perspective camera with a focal length of 12). As  $k$  increases, the resolution of head meshes and features are continuously refined. Fig. 2 shows the subdivided vertices. To balance quality and efficiency, we set the maximum subdivision level to  $K = 2$ , resulting in 79, 936 vertices. Notably, unlike LAM which performs costly cross-attention across all 80K vertices, we compute cross-attention at the initial level with  $N_0 = 5K$  vertices. High-resolution geometry is incrementally refined through efficient subdivision and sampling, significantly reducing computational and memory costs while maintaining reconstruction quality.

### 3.2. Occlusion-Aware Feature Fusion

Once hierarchical global and local features are extracted, we propose an occlusion-aware fusion strategy guided by the depth buffer to ensure robust feature integration. During the rasterization of the 3D mesh, occlusion culling is applied to invisible vertices. Consequently, local features sampled via projection are accurate for visible vertices but may be ambiguous for occluded ones, as their corresponding 2D image features are absent. On the other hand, the global features

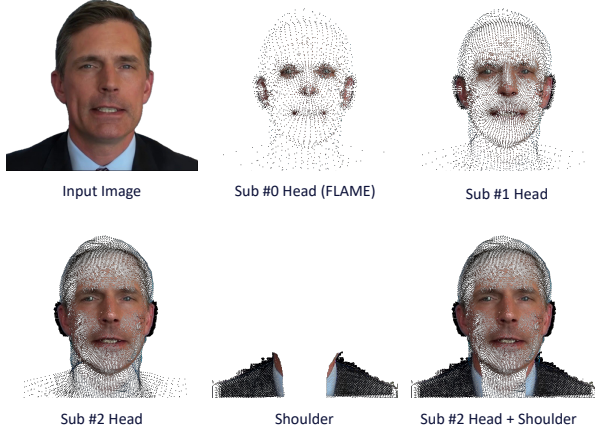


Figure 2. Since the original FLAME model lacks shoulder regions and sufficient geometric representation for high-resolution LOD requirements, we progressively subdivide the head mesh during training and enhance it with predicted shoulder geometry through cross-region combination. Sub  $\#i$  indicates that the head mesh has been subdivided  $i$  times.

contain high-level semantic information from the input image, allowing them to infer plausible representations for occluded regions but lacking high-frequency details. Building upon this observation, we leverage the depth buffer to identify and retain only high-confidence local features from visible vertices, selectively fusing them with global features. Specifically, when each vertex is projected into the image space, its depth value is compared with the depth buffer at the corresponding pixel location. Based on this comparison, a binary visibility mask  $M^{GS_k} \in \{0, 1\}^{N_k}$  is constructed. Formally, for each vertex  $v_i \in V_k$ , let  $z_i$  be its depth in camera space, and  $\hat{z}_i$  be the depth value recorded in the depth buffer. The mask  $M^{GS_k} \in \{0, 1\}^{N_k}$  is defined as:

$$M_i^{GS_k} = \begin{cases} 1, & \text{if } z_i = \hat{z}_i \\ 0, & \text{if } z_i > \hat{z}_i \end{cases}, \quad \forall i = 1, 2, \dots, N_k, \quad (5)$$

Here,  $M_i^{GS_k} = 1$  indicates that the vertex  $v_i$  is visible, while  $M_i^{GS_k} = 0$  identifies occluded vertices. The final fused head feature  $F_h^{GS_k}$  is computed as:

$$F_h^{GS_k} = F_{global}^{GS_k} + F_{local}^{GS_k} \odot M^{GS_k}, \quad 0 \leq k \leq K. \quad (6)$$

By effectively combining the strengths of global and local features, we achieve a more robust representation of the head vertices. A set of MLPs  $\Phi$  are then employed to regress the Gaussian attributes for each head vertex, including color  $c_h$ , opacity  $o_h$ , scale  $s_h$ , and rotation  $r_h$ :

$$c_h, o_h, s_h, r_h = \Phi_{c,o,s,r}(F_h^{GS_K}). \quad (7)$$

The positions  $p_h$  are directly derived from the subdivided mesh  $V_K$ , resulting in a full set of Gaussian parameters  $\mathbb{H}$

that describe the head model:

$$\mathbb{H} = \{c_h, o_h, s_h, r_h, p_h = V_K\}. \quad (8)$$

### 3.3. Multi-region Modeling and Integration

The FLAME model lacks sufficient vertex coverage in the shoulder region, resulting in coarse and blurry reconstructions in prior methods. To address this limitation, inspired by [53], we perform image segmentation on the source image to obtain a shoulder mask  $M_s$ . The extracted local features  $F_{local}$  are passed through a convolutional neural network to generate a feature plane, where separate channels encode the Gaussian rendering attributes. The shoulder-relevant region is then isolated using the mask, obtaining the corresponding parameters as follows:

$$c_s, o_s, s_s, r_s, O_s = \text{Flatten}(\text{Conv}(F_{local}^{GS}) \odot M_s), \quad (9)$$

where  $O_s$  represents the offset for shoulder points. For position estimation, we generate the 3D shoulder points  $\hat{p}_s$  on an image-aligned plane in world space based on the given camera transformation and feature plane resolution. These points are paired with their corresponding direction vectors  $n_s$ . The final shoulder points are then calculated as:

$$\mathbb{S} = \{c_s, o_s, s_s, r_s, p_s = \hat{p}_s + O_s \cdot n_s\}. \quad (10)$$

Finally, we concatenate the head and shoulder parameter sets along the attribute dimensions to form a complete Gaussian parameter set  $\mathbb{G}$  that covers both head and shoulder regions:

$$\mathbb{G} = \mathbb{H} \oplus \mathbb{S}. \quad (11)$$

### 3.4. Reenactment and Rendering

After reconstructing the Gaussian head avatar  $\mathbb{G}$ , our framework enables efficient reenactment, allowing the model to mimic facial expressions and head movements observed in a target video. As shown in Fig. 1, given a driving image  $I_d$ , the expression and pose parameters are extracted using a 3DMM estimator. These driving parameters are then combined with the identity-related parameters derived from  $I_s$  to generate novel FLAME vertices. The FLAME vertices are further refined via subdivisions, resulting in the final head positions  $p_h$  used for Gaussian rendering. Notably, the reenactment process only needs to update the positional component  $p_h$  in the Gaussian parameter set  $\mathbb{G}$  (the red dashed line in Fig. 1), enabling real-time rendering of avatar animations. To enhance the expressiveness of the Gaussian representation, we adopt a rendering pipeline inspired by [6]. Specifically, instead of rendering RGB values, we predict multi-channel feature maps, where the first three channels encode a coarse RGB image  $I_c$ . The feature maps are subsequently refined by a UNet-based neural refiner to produce the final high-quality output  $I_r$ .

### 3.5. Learning Strategy

The training process is conducted on a large-scale human video dataset in a self-supervised manner. For each video, we randomly sample two frames and assign one as the source image  $I_s$  and the other as the driving image  $I_d$ . The objective is to train the network to generate an output image that closely resembles the driving image in both appearance and motion.

To achieve this, we employ a multi-component loss function, combining L2 loss, SSIM loss, and perceptual loss, applied to both the coarse and refined output images. Additionally, to constrain the displacement of FLAME model vertices, we impose a regularization term on the vertex offsets. The total loss is defined as:

$$\mathcal{L} = \lambda_1 \mathcal{L}_2(I_d, I_c \& I_r) + \lambda_2 \mathcal{L}_{\text{SSIM}}(I_d, I_c \& I_r) + \lambda_3 \mathcal{L}_{\text{percep}}(I_d, I_c \& I_r) + \lambda_4 \mathcal{L}_{\text{reg}}, \quad (12)$$

where  $\mathcal{L}(A, B \& C)$  denotes  $\mathcal{L}(A, B) + \mathcal{L}(A, C)$ ,  $\mathcal{L}_2$ ,  $\mathcal{L}_{\text{SSIM}}$ , and  $\mathcal{L}_{\text{percep}}$  are computed on both the  $I_c$  and refined  $I_r$ , and  $\mathcal{L}_{\text{reg}} = \|\text{offset}\|_2$  penalizes large vertex displacements.

## 4. Experiment

### 4.1. Datasets and Settings

**Datasets.** Our model is trained on the VFHQ [55] dataset, which contains video clips from a variety of interview scenarios. To ensure temporal diversity, we uniformly sample frames from each video following previous works [6, 21], leading to a total of 766, 263 frames across 15, 204 video clips. All images are cropped to focus on the head region, resized to  $512 \times 512$  pixels for consistency, and further processed with camera pose tracking, FLAME parameter estimation, and background removal as described in prior works [6, 7]. For evaluation, we adopt the official test split of the VFHQ dataset, comprising 2, 500 frames from 50 videos. The first frame of each video is used as the source image, while the remaining frames serve as driving and target images for reenactment. Additionally, we evaluate our method on the HDTF dataset [67] using the model trained solely on VFHQ to demonstrate its generalization capability. We process the HDTF dataset using the standard test split introduced in Ma et al. [36] and Li et al. [33], which includes 19 video sequences.

**Implementation Details.** Our transformer network for extracting global features consists of two decoder layers, each with 8 attention heads. The dimension of the FLAME positional encoding is set to 256. Rather than relying on an external semantic segmentation model, we derive the shoulder region mask by calculating the difference between the portrait mask at the bottom of the image and the depth buffer mask in a straightforward manner. The average number of Gaussian points in the shoulder region is 9K. During

training, the weights of the DINOv2 and 3DMM estimator modules are frozen. We train the entire model on a single NVIDIA A100 GPU for 6 epochs using the Adam optimizer with a learning rate of  $1 \times 10^{-4}$  and a batch size of 8. The subdivide levels are gradually increased based on the training stage. We set the loss parameters  $\lambda_1 = 10$ ,  $\lambda_2 = 1$ , and  $\lambda_3 = \lambda_4 = 0.1$ . More details are provided in the supplemental material.

**Evaluation Metrics.** To comprehensively evaluate the performance of both self- and cross-identity reenactment, we employ a multi-faceted assessment framework that incorporates a variety of quantitative metrics. For self-reenactment scenarios where ground-truth data is available, we assess the quality of generated images using three widely adopted objective measures: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [66]. These metrics provide reliable comparisons between synthesized outputs and reference ground-truth images. To evaluate identity preservation, we compute the cosine distance between face recognition features extracted from the source and reenacted images, following the methodology proposed in Deng et al. [9]. For assessing the accuracy of expression and pose transfer, we utilize a 3D Morphable Model (3DMM) estimator [10] to calculate the Average Expression Distance (AED) and Average Pose Distance (APD). Additionally, we use facial landmark detection [4] to measure the Average Keypoint Distance (AKD), which provides further insight into the precision of motion control during animation. In the case of cross-identity reenactment, where ground-truth data is not available, we adopt an evaluation protocol based on Consistency of Identity Similarity (CSIM), AED, and APD—metrics aligned with those used in recent studies [6, 21]. This ensures comparability across different methods and enables meaningful analysis.

### 4.2. Baseline methods

We conduct a comprehensive comparison between our method and state-of-the-art 3D avatar approaches, including ROME [27], StyleHeat [61], OTAvatar [36], HideNeRF [32], GOHA [33], CVTHead [35], GPAvatar [7], Real3DPortrait [60], Portrait4D [11], Portrait4D-v2 [12], GAGAvatar [6], and LAM [21]. For each baseline, we utilize its official implementation to produce results. For our model, the level-of-detail (LOD) can be controlled at inference time, Sub # $i$  indicates that the head mesh has been subdivided  $i$  times.

### 4.3. Qualitative Evaluation

We compare our method against baseline approaches on the VFHQ and HDTF datasets, with qualitative results shown in Fig. 3. As can be seen from the figure, GPAvatar [7] suffers from inaccurate expression tracking and exhibits no-



Figure 3. Cross-identity reenactment results on VFHQ and HDTF datasets.

ticeable blurring in the neck region—a limitation also observed in GAGAvatar [6] (see rows 4 and 5). Although Real3DPortrait [60], Portrait4D [11], and Portrait4Dv2 [12] achieve high visual fidelity, they introduce severe misalignment artifacts under certain poses, manifesting as prominent cracking near the neck and chin (row 1). Additionally, LAM [21] exhibits unnatural shoulder tilting (rows 1 and 3) alongside noticeable artifacts around the mouth and teeth (rows 2 and 6). In contrast, our method achieves superior visual quality compared to existing approaches while using significantly fewer Gaussian points. Notably, even our low-resolution variant (Sub #1 with  $\sim 29\text{K}$  Gaussian points), shown in the last column of Fig. 3, maintains comparable visual fidelity, making it particularly well-suited for deployment in high-speed applications or on hardware with limited computational resources.

#### 4.4. Quantitative Evaluation

We report the quantitative results in Tab. 1 and Tab. 2. Our method (Sub #2) outperforms existing approaches across all reconstruction metrics (PSNR, SSIM, and LPIPS), as

well as identity, expression, and pose consistency. Remarkably, our low-resolution LOD Sub #1 surpasses LAM (80K Gaussians) and GAGAvatar (180K Gaussians) on both datasets using only 29K Gaussian points, demonstrating the effectiveness of our hierarchical feature extraction and fusion strategy. We further report the inference efficiency in Tab. 3. Our method achieves an inference speed of 85 FPS on an A100 GPU and 126 FPS on the consumer-grade RTX 4090 GPU, using the native PyTorch framework and the official implementation of 3D Gaussian Splatting. Compared to existing neural-rendering-based methods, our approach attains the highest inference speed. Moreover, our method outperforms LAM (280 FPS on A100 GPU without neural rendering) in terms of geometric details and dynamic textures, achieving an optimal balance between efficiency and visual quality.

#### 4.5. Ablation Studies

**Subdivision Times.** To evaluate the effect of subdivision levels, we compare the results of our model with varying subdivision levels. The visual results are presented in

Table 1. Quantitative results on the VFHQ dataset. The **first**, **second**, and **third** best-performing methods are highlighted. The Sub # indicates the subdivision level for inference.

Method	Self Reenactment							Cross Reenactment		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CSIM $\uparrow$	AED $\downarrow$	APD $\downarrow$	AKD $\downarrow$	CSIM $\uparrow$	AED $\downarrow$	APD $\downarrow$
StyleHeat	19.95	0.726	0.211	0.537	0.199	0.385	7.659	0.407	0.279	0.551
ROME	19.96	0.786	0.192	0.701	0.138	0.186	4.986	0.530	0.259	0.277
OTAvatar	17.65	0.563	0.294	0.465	0.234	0.545	18.19	0.364	0.324	0.678
HideNeRF	19.79	0.768	0.180	0.787	0.143	0.361	7.254	0.514	0.277	0.527
GOHA	20.15	0.770	0.149	0.664	0.176	0.173	6.272	0.518	0.274	0.261
CVTHead	18.43	0.706	0.317	0.504	0.186	0.224	5.678	0.374	0.261	0.311
GPAvatar	21.04	0.807	0.150	0.772	0.132	0.189	4.226	0.564	0.255	0.328
Real3DPortrait	20.88	0.780	0.154	0.801	0.150	0.268	5.971	0.663	0.296	0.411
Portrait4D	20.35	0.741	0.191	0.765	0.144	0.205	4.854	0.596	0.286	0.258
Portrait4D-v2	21.34	0.791	0.144	0.803	0.117	0.187	3.749	0.656	0.286	0.273
GAGAvatar	21.83	0.818	0.122	0.816	0.111	0.135	3.349	0.633	0.253	0.247
LAM	22.65	0.829	0.109	0.822	0.102	0.134	2.059	0.651	0.250	0.356
Ours (Sub #2)	22.72	0.831	0.091	0.869	0.088	0.111	2.045	0.660	0.235	0.257
Ours (Sub #1)	22.68	0.830	0.094	0.858	0.089	0.112	2.055	0.644	0.233	0.260
Ours (Sub #0)	22.18	0.817	0.102	0.855	0.134	0.142	2.790	0.616	0.254	0.279

Table 2. Quantitative results on the HDTF dataset.

Method	Self Reenactment							Cross Reenactment		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CSIM $\uparrow$	AED $\downarrow$	APD $\downarrow$	AKD $\downarrow$	CSIM $\uparrow$	AED $\downarrow$	APD $\downarrow$
StyleHeat	21.41	0.785	0.155	0.657	0.158	0.162	4.585	0.632	0.271	0.239
ROME	20.51	0.803	0.145	0.738	0.133	0.123	4.763	0.726	0.268	0.191
OTAvatar	20.52	0.696	0.166	0.662	0.180	0.170	8.295	0.643	0.292	0.222
HideNeRF	21.08	0.811	0.117	0.858	0.120	0.247	5.837	0.843	0.276	0.288
GOHA	21.31	0.807	0.113	0.725	0.162	0.117	6.332	0.735	0.277	0.136
CVTHead	20.08	0.762	0.179	0.608	0.169	0.138	4.585	0.591	0.242	0.203
GPAvatar	23.06	0.855	0.104	0.855	0.114	0.135	3.293	0.842	0.268	0.219
Real3DPortrait	22.82	0.835	0.103	0.851	0.138	0.137	4.640	0.903	0.299	0.238
Portrait4D	20.81	0.786	0.137	0.810	0.134	0.131	4.151	0.793	0.291	0.240
Portrait4D-v2	22.87	0.860	0.105	0.860	0.111	0.111	3.292	0.857	0.262	0.183
GAGAvatar	23.13	0.863	0.103	0.862	0.110	0.111	2.985	0.851	0.231	0.181
LAM	23.43	0.873	0.097	0.865	0.101	0.093	1.965	0.849	0.230	0.229
Ours (Sub #2)	24.14	0.875	0.061	0.943	0.080	0.064	1.806	0.886	0.226	0.155
Ours (Sub #1)	24.06	0.874	0.063	0.937	0.081	0.066	1.834	0.881	0.227	0.155
Ours (Sub #0)	23.85	0.868	0.067	0.942	0.121	0.085	2.381	0.872	0.246	0.156

Fig. 4. Our observations indicate that higher subdivision levels capture more high-frequency details, such as wrinkles, leading to improved reconstruction quality. Quantitative comparisons are provided in Tab. 1 and Tab. 2. We observe that performance saturates after two levels of subdivision ( Fig. 5). We believe this is primarily due to the limited resolution of DINOv2’s feature maps: projection-based sampling beyond 80K vertices fails to recover additional geometric detail from the 88K (296 × 296) feature grid. Higher-resolution features could unlock further improvements. Additionally, increasing the number of sub-

divisions enhances the reconstruction quality at the cost of reduced inference speed, as shown in Tab. 3. Notably, even our low-resolution LOD demonstrates competitive performance compared to existing methods, highlighting the effectiveness of our framework.

**Local-Global and Multi-Region Feature Fusion.** We conduct an ablation study by removing the global-local feature fusion in the OAFF module and using only global features. As shown in Fig. 4 and Tab. 4, The absence of sampled local features significantly degrades identity consistency, while the lack of global features introduces artifacts in dynamic

Table 3. The reenactment speed of neural-rendering-based methods measured in FPS, averaged over 100 frames. Driving parameters estimation time is excluded as they can be precomputed.

Methods	A100 GPU								RTX 4090 GPU		
	HideNeRF	CVTHead	Real3D	P4D-v2	GAGAvatar	Avat3r	LAM	Ours (Sub #2)	Sub #2	Sub #1	Sub #0
FPS	9.73	18.09	4.55	9.62	67.12	53	280	85.94	126.44	148.04	152.57



Figure 4. Our local-global feature fusion (OAFF) and multi-region fusion strategy significantly improve identity consistency and completeness in non-head regions. The neural refiner further boosts visual fidelity, especially for dynamic facial expressions.

Table 4. Ablation results on the VFHQ dataset.

Methods	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CSIM $\uparrow$
w/o Global Feature	20.85	0.796	0.121	0.429
w/o Local Feature	21.21	0.802	0.128	0.429
w/o Refiner	21.42	0.809	0.115	0.842
w/o Shoulder	22.42	0.828	0.099	0.867
Ours	<b>22.72</b>	<b>0.831</b>	<b>0.091</b>	<b>0.869</b>

regions—particularly the eyes and mouth—that are inconsistent with the source image. Furthermore, when the shoulder region is excluded during rendering, the results display an incomplete and blurry appearance of the shoulder, as illustrated in Fig. 4.

**Neural Rendering.** We evaluate the effectiveness of the neural refiner module. Tab. 4 and Fig. 4 demonstrate that the neural refiner enhances fine details such as teeth and plays a key role in capturing expression-dependent features, including forehead wrinkles during eyebrow raising.

## 5. Conclusion

In this paper, we present a novel multi-LOD framework for Gaussian head avatar reconstruction in a feed-forward manner. Our method enables dynamic level-of-detail (LOD) rendering at runtime, offering flexibility to accommodate varying device capabilities and inference speed requirements. Our model exhibits superior reconstruction and reenactment performance with significantly reduced computational cost. This is achieved by an efficient multi-level global and local feature extraction and a coarse-to-fine

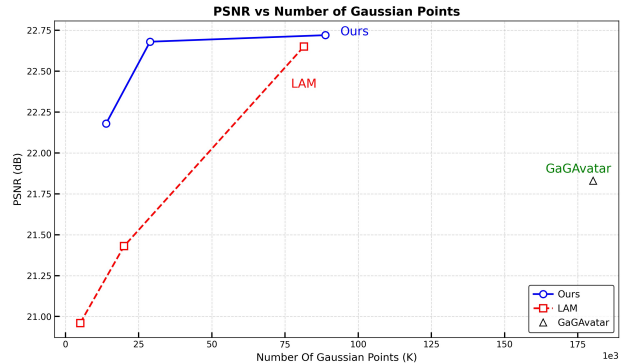


Figure 5. The correlation between Gaussian count and reconstruction performance on the VFHQ dataset.




Figure 6. Novel view results of OMG-Avatar. Since our model is trained on monocular interview videos, it achieves robust performance for head rotations within  $\pm 60^\circ$ , but exhibits noticeable artifacts beyond this range.

refinement strategy, as well as an occlusion-aware fusion mechanism. Moreover, our multi-region modeling scheme effectively enhances the visual fidelity of shoulder areas. Extensive experiments on two public datasets demonstrate that our approach outperforms state-of-the-art methods in terms of reconstruction quality, reenactment performance, and computational efficiency.

**Limitations and Future Work.** Despite achieving strong results, our approach has two main limitations. First, the 3D Gaussian head model relies on the FLAME prior and accurate 3D morphable model (3DMM) tracking. However, FLAME does not capture fine facial dynamics—such as tongue motion, hair deformation, limiting the expressiveness of the generated avatars. Second, training solely on monocular videos reduces robustness to large viewpoint changes ( $> 60^\circ$ , see the last column in Fig. 6). To address these issues, we plan to incorporate multi-view datasets for training, which will enhance spatial understanding and improve robustness across varying viewpoints.

## References

- [1] ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. Rignerrf: Fully controllable neural 3d portraits. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 20364–20373, 2022. 2
- [2] Yunpeng Bai, Yanbo Fan, Xuan Wang, Yong Zhang, Jingxiang Sun, Chun Yuan, and Ying Shan. High-fidelity facial avatar reconstruction from monocular video with generative priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4541–4551, 2023. 1
- [3] Ziqian Bai, Feitong Tan, Zeng Huang, Kripasindhu Sarkar, Danhang Tang, Di Qiu, Abhimitra Meka, Ruofei Du, Mingsong Dou, Sergio Orts-Escolano, et al. Learning personalized high quality volumetric head avatars from monocular rgb videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16890–16900, 2023. 2
- [4] 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 *Proceedings of the IEEE international conference on computer vision*, pages 1021–1030, 2017. 5
- [5] Egor Burkov, Igor Pasechnik, Artur Grigorev, and Victor Lempitsky. Neural head reenactment with latent pose descriptors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13786–13795, 2020. 2
- [6] Xuangeng Chu and Tatsuya Harada. Generalizable and animatable gaussian head avatar. *Advances in Neural Information Processing Systems*, 37:57642–57670, 2024. 1, 2, 3, 4, 5, 6
- [7] Xuangeng Chu, Yu Li, Ailing Zeng, Tianyu Yang, Lijian Lin, Yunfei Liu, and Tatsuya Harada. Gpavatar: Generalizable and precise head avatar from image (s). *arXiv preprint arXiv:2401.10215*, 2024. 2, 5
- [8] Jiahao Cui, Hui Li, Yao Yao, Hao Zhu, Hanlin Shang, Kaihui Cheng, Hang Zhou, Siyu Zhu, and Jingdong Wang. Hallo2: Long-duration and high-resolution audio-driven portrait image animation. *arXiv preprint arXiv:2410.07718*, 2024. 2
- [9] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. 5
- [10] 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 workshops*, pages 0–0, 2019. 5
- [11] Yu Deng, Duomin Wang, Xiaohang Ren, Xingyu Chen, and Baoyuan Wang. Portrait4d: Learning one-shot 4d head avatar synthesis using synthetic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7119–7130, 2024. 5, 6
- [12] Yu Deng, Duomin Wang, and Baoyuan Wang. Portrait4d-v2: Pseudo multi-view data creates better 4d head synthesizer. In *European Conference on Computer Vision*, pages 316–333. Springer, 2024. 5, 6
- [13] Xiaonuo Dongye, Hanzhi Guo, Le Luo, Haiyan Jiang, Yihua Bao, Zeyu Tian, and Dongdong Weng. Lodavatar: Hierarchical embedding and adaptive levels of detail with gaussian splatting for enhanced human avatars. *arXiv preprint arXiv:2410.20789*, 2024. 2
- [14] Nikita Drobyshev, Jenya Chelishev, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. Megaportraits: One-shot megapixel neural head avatars. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2663–2671, 2022. 2
- [15] Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8649–8658, 2021. 1, 2
- [16] Xuan Gao, Chenglai Zhong, Jun Xiang, Yang Hong, Yudong Guo, and Juyong Zhang. Reconstructing personalized semantic facial nerf models from monocular video. *ACM Transactions on Graphics (TOG)*, 41(6):1–12, 2022. 2
- [17] Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Luthi, Sandro Schönborn, and Thomas Vetter. Morphable face models—an open framework. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 75–82. IEEE, 2018. 2
- [18] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 1, 2
- [19] Jianzhu Guo, Dingyun Zhang, Xiaoqiang Liu, Zhizhou Zhong, Yuan Zhang, Pengfei Wan, and Di Zhang. Liveportrait: Efficient portrait animation with stitching and retargeting control. *arXiv preprint arXiv:2407.03168*, 2024. 1, 2
- [20] 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*, pages 5784–5794, 2021. 2
- [21] Yisheng He, Xiaodong Gu, Xiaodan Ye, Chao Xu, Zhengyi Zhao, Yuan Dong, Weihao Yuan, Zilong Dong, and Liefeng Bo. Lam: Large avatar model for one-shot animatable gaussian head. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–13, 2025. 1, 2, 3, 5, 6
- [22] Fa-Ting Hong, Longhao Zhang, Li Shen, and Dan Xu. Depth-aware generative adversarial network for talking head video generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3397–3406, 2022. 2
- [23] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 2

- [24] Jianwen Jiang, Chao Liang, Jiaqi Yang, Gaojie Lin, Tianyun Zhong, and Yanbo Zheng. Loopy: Taming audio-driven portrait avatar with long-term motion dependency. *arXiv preprint arXiv:2409.02634*, 2024. 1
- [25] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 2
- [26] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2
- [27] Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. Realistic one-shot mesh-based head avatars. In *European Conference on Computer Vision*, pages 345–362. Springer, 2022. 2, 5
- [28] Taekyung Ki, Dongchan Min, and Gyeongsu Chae. Learning to generate conditional tri-plane for 3d-aware expression controllable portrait animation. In *European Conference on Computer Vision*, pages 476–493. Springer, 2024. 1, 2
- [29] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Transactions on Graphics (TOG)*, 42(4):1–14, 2023. 2
- [30] Tobias Kirschstein, Javier Romero, Artem Sevastopolsky, Matthias Nießner, and Shunsuke Saito. Avat3r: Large animatable gaussian reconstruction model for high-fidelity 3d head avatars. *arXiv preprint arXiv:2502.20220*, 2025. 1, 2
- [31] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017. 2, 3
- [32] Weichuang Li, Longhao Zhang, Dong Wang, Bin Zhao, Zhigang Wang, Mulin Chen, Bang Zhang, Zhongjian Wang, Liefeng Bo, and Xuelong Li. One-shot high-fidelity talking-head synthesis with deformable neural radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17969–17978, 2023. 1, 2, 5
- [33] Xueting Li, Shalini De Mello, Sifei Liu, Koki Nagano, Umar Iqbal, and Jan Kautz. Generalizable one-shot 3d neural head avatar. *Advances in Neural Information Processing Systems*, 36:47239–47250, 2023. 5
- [34] C LOOP. Smooth subdivision surfaces based on triangles. *Master’s thesis, University of Utah, Department of Mathematics*, 1987. 2
- [35] Haoyu Ma, Tong Zhang, Shanlin Sun, Xiangyi Yan, Kun Han, and Xiaohui Xie. Cvthead: One-shot controllable head avatar with vertex-feature transformer. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6131–6141, 2024. 2, 5
- [36] Zhiyuan Ma, Xiangyu Zhu, Guo-Jun Qi, Zhen Lei, and Lei Zhang. Otavatar: One-shot talking face avatar with controllable tri-plane rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16901–16910, 2023. 1, 5, 2
- [37] Marco Nef. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. 2
- [38] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 3
- [39] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5865–5874, 2021. 1, 2
- [40] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021. 2
- [41] 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. Ieee, 2009. 2
- [42] Yurui Ren, Ge Li, Yuanqi Chen, Thomas H Li, and Shan Liu. Pirenderer: Controllable portrait image generation via semantic neural rendering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 13759–13768, 2021. 2
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1
- [44] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in neural information processing systems*, 32, 2019. 1, 2
- [45] Matthew Tancik, Ben Mildenhall, and Ren Ng. Stegastamp: Invisible hyperlinks in physical photographs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2117–2126, 2020. 1
- [46] Jiapeng Tang, Davide Davoli, Tobias Kirschstein, Liam Schoneveld, and Matthias Niessner. Gaf: Gaussian avatar reconstruction from monocular videos via multi-view diffusion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5546–5558, 2025. 1, 2
- [47] Shengji Tang, Weicai Ye, Peng Ye, Weihao Lin, Yang Zhou, Tao Chen, and Wanli Ouyang. Hisplat: Hierarchical 3d gaussian splatting for generalizable sparse-view reconstruction. *arXiv preprint arXiv:2410.06245*, 2024. 2
- [48] Kartik Teotia, Hyeonwoo Kim, Pablo Garrido, Marc Habermann, Mohamed Elgharib, and Christian Theobalt. Gaussianheads: End-to-end learning of drivable gaussian head avatars from coarse-to-fine representations. *ACM Transactions on Graphics (TOG)*, 43(6):1–12, 2024. 2
- [49] Linrui Tian, Qi Wang, Bang Zhang, and Liefeng Bo. Emo: Emote portrait alive generating expressive portrait videos with audio2video diffusion model under weak conditions. In *European Conference on Computer Vision*, pages 244–260, 2024. 1, 2

- [50] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12959–12970, 2021. 2
- [51] Duomin Wang, Yu Deng, Zixin Yin, Heung-Yeung Shum, and Baoyuan Wang. Progressive disentangled representation learning for fine-grained controllable talking head synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17979–17989, 2023. 2
- [52] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10039–10049, 2021. 1
- [53] Tianhao Walter Wu, Jing Yang, Zhilin Guo, Jingyi Wan, Fangcheng Zhong, and Cengiz Oztireli. Gaussian head & shoulders: High fidelity neural upper body avatars with anchor gaussian guided texture warping. In *The Thirteenth International Conference on Learning Representations*, 2024. 1, 2, 4
- [54] Jun Xiang, Xuan Gao, Yudong Guo, and Juyong Zhang. Flashavatar: High-fidelity head avatar with efficient gaussian embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1802–1812, 2024. 1
- [55] Liangbin Xie, Xintao Wang, Honglun Zhang, Chao Dong, and Ying Shan. Vfhq: A high-quality dataset and benchmark for video face super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 657–666, 2022. 5
- [56] Mingwang Xu, Hui Li, Qingkun Su, Hanlin Shang, Liwei Zhang, Ce Liu, Jingdong Wang, Yao Yao, and Siyu Zhu. Hallo: Hierarchical audio-driven visual synthesis for portrait image animation. *arXiv preprint arXiv:2406.08801*, 2024. 2
- [57] Sicheng Xu, Jiaolong Yang, Dong Chen, Fang Wen, Yu Deng, Yunde Jia, and Xin Tong. Deep 3d portrait from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7710–7720, 2020. 2
- [58] Yuelang Xu, Benwang Chen, Zhe Li, Hongwen Zhang, Lizhen Wang, Zerong Zheng, and Yebin Liu. Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1931–1941, 2024. 1, 2
- [59] Songlin Yang, Wei Wang, Yushi Lan, Xiangyu Fan, Bo Peng, Lei Yang, and Jing Dong. Learning dense correspondence for nerf-based face reenactment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6522–6530, 2024. 2
- [60] Zhenhui Ye, Tianyun Zhong, Yi Ren, Jiaqi Yang, Weichuang Li, Jiawei Huang, Ziyue Jiang, Jinzheng He, Rongjie Huang, Jinglin Liu, et al. Real3d-portrait: One-shot realistic 3d talking portrait synthesis. *arXiv preprint arXiv:2401.08503*, 2024. 5, 6
- [61] Fei Yin, Yong Zhang, Xiaodong Cun, Mingdeng Cao, Yanbo Fan, Xuan Wang, Qingyan Bai, Baoyuan Wu, Jue Wang, and Yujiu Yang. Styleheat: One-shot high-resolution editable talking face generation via pre-trained stylegan. In *European conference on computer vision*, pages 85–101. Springer, 2022. 5
- [62] Wangbo Yu, Yanbo Fan, Yong Zhang, Xuan Wang, Fei Yin, Yunpeng Bai, Yan-Pei Cao, Ying Shan, Yang Wu, Zhongqian Sun, et al. Nofa: Nerf-based one-shot facial avatar reconstruction. In *ACM SIGGRAPH 2023 conference proceedings*, pages 1–12, 2023. 1, 2
- [63] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9459–9468, 2019. 2
- [64] Bowen Zhang, Chenyang Qi, Pan Zhang, Bo Zhang, Hsiang-Tao Wu, Dong Chen, Qifeng Chen, Yong Wang, and Fang Wen. Metaportrait: Identity-preserving talking head generation with fast personalized adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22096–22105, 2023. 2
- [65] Dongbin Zhang, Yunfei Liu, Lijian Lin, Ye Zhu, Yang Li, Minghan Qin, Yu Li, and Haoqian Wang. Guava: Generalizable upper body 3d gaussian avatar. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14205–14217, 2025. 1, 2
- [66] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5
- [67] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3661–3670, 2021. 5
- [68] Zicheng Zhang, Ruobing Zheng, Bonan Li, Congying Han, Tianqi Li, Meng Wang, Tiande Guo, Jingdong Chen, Ziwen Liu, and Ming Yang. Learning dynamic tetrahedra for high-quality talking head synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5209–5219, 2024. 2
- [69] 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 Transactions on Graphics*, 43(1):1–16, 2023. 2
- [70] Xiaochen Zhao, Hongyi Xu, Guoxian Song, You Xie, Chenxu Zhang, Xiu Li, Linjie Luo, Jinli Suo, and Yebin Liu. X-nemo: Expressive neural motion reenactment via disentangled latent attention. *arXiv preprint arXiv:2507.23143*, 2025. 1, 2
- [71] Xiaozheng Zheng, Chao Wen, Zhaohu Li, Weiyi Zhang, Zhuo Su, Xu Chang, Yang Zhao, Zheng Lv, Xiaoyuan Zhang, Yongjie Zhang, et al. Headgap: Few-shot 3d head avatar via generalizable gaussian priors. In *2025 International Conference on 3D Vision (3DV)*, pages 946–957. IEEE, 2025. 1, 2
- [72] Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J Black, and Otmar Hilliges. Pointavatar: Deformable

- point-based head avatars from videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21057–21067, 2023. [1](#), [2](#)
- [73] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4176–4186, 2021. [2](#)
- [74] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Instant volumetric head avatars. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4574–4584, 2023. [2](#)