# TaoAvatar: Real-Time Lifelike Full-Body Talking Avatars for Augmented Reality via 3D Gaussian Splatting

Jianchuan Chen<sup>\*</sup>, Jingchuan Hu<sup>\*</sup>, Gaige Wang<sup>\*</sup>, Zhonghua Jiang Tiansong Zhou, Zhiwen Chen<sup>‡</sup>, Chengfei Lv<sup>†</sup> Alibaba Group, Hangzhou, China



Figure 1. TaoAvatar generates photorealistic, topology-consistent 3D full-body avatars from multi-view sequences. It provides high-quality,

real-time rendering with low storage requirements, compatible across various mobile and AR devices like the Apple Vision Pro.

#### Abstract

Realistic 3D full-body talking avatars hold great potential in AR, with applications ranging from e-commerce live streaming to holographic communication. Despite advances in 3D Gaussian Splatting (3DGS) for lifelike avatar creation, existing methods struggle with fine-grained control of facial expressions and body movements in full-body talking tasks. Additionally, they often lack sufficient details and cannot run in real-time on mobile devices. We present TaoAvatar, a high-fidelity, lightweight, 3DGS-based full-body talking avatar driven by various signals. Our approach starts by creating a personalized clothed human parametric template that binds Gaussians to represent appearances. We then pre-train a StyleUnet-based network to handle complex pose-dependent non-rigid deformation, which can capture high-frequency appearance details but is too resource-intensive for mobile devices. To overcome this, we "bake" the non-rigid deformations into a lightweight MLP-based network using a distillation technique and develop blend shapes to compensate for details. Extensive experiments show that TaoAvatar achieves state-of-the-art rendering quality while running in real-time across various devices, maintaining 90 FPS on high-definition stereo devices such as the Apple Vision Pro.

# **1. Introduction**

Creating lifelike 3D human avatars is a dynamic and rapidly advancing research area in computer vision and graphics, essential for applications in AR/VR, 3D entertainment, and virtual communication. Despite decades of research, achieving realistic and immersive avatars on mobile devices—particularly AR—remains challenging due to computational limitations. Current industry solutions, such as MetaHuman, often rely on high-precision scans and extensive manual effort by artists for 3D modeling, rigging, and motion capture.

In recent years, academic researchers have adopted parametric models [37, 43] for human representation, which are constructed from massive body scans and can capture a performer's shape, expression, and pose from images or videos. These methods [4-6, 27] extend the naked model [37, 43] by adding offsets to represent tailored clothing, but are still struggling with complex geometries and high-frequency details, such as loose skirts and fine hair, due to limitations of topology and texture resolution. With the rise of neural radiance fields (NeRF) [40], many researchers tend to define implicit neural human representations in canonical space, animated by backward linear blend skinning [44, 65] of these parametric models [37, 43]. While these methods can handle arbitrary topology and deliver higher rendering quality, they are highly timeconsuming due to the large amount of sampling required

<sup>\*</sup>Equal contribution. <sup>‡</sup>Project Leader and <sup>†</sup>Corresponding author. Project: https://PixelAI-Team.github.io/TaoAvatar

during volume rendering. Although many efforts [9, 24] have significantly improved rendering efficiency by introducing iNGP [42], the animated results still fail to escape the Uncanny Valley effect when driven. Recently, 3D Gaussian Splatting (3DGS) [29], an explicit and efficient pointbased representation, has gained tremendous attention from researchers for its ability to deliver both high-quality rendering and real-time performance. Compared to implicit representations [40, 54], the explicit point-based representation is more compatible with being integrated with the mesh-based parametric models [37, 43], which can be directly driven by forward linear blend skinning. Many recent methods [20, 34, 46, 48], combining 3DGS and the parametric models [37, 43], have made significant strides toward realizing lifelike 3D human avatars. However, achieving real-time rendering on mobile devices remains challenging, especially for AR glasses, which require high resolution, stereo rendering, and seamless interaction.

We introduce TaoAvatar, a high-fidelity, lightweight, 3DGS-based full-body talking avatar designed for running in real-time on Augmented Reality (AR) devices. Unlike previous approaches [8, 10, 34], which develop implicit parameterized models from scratch, we construct a personalized, clothed parameterized template that preserves the facial expressions and gesture controls inherent to the native SMPLX [43]. At the same time, we bind the Gaussians to the triangles as texture to create a hybrid avatar representation. Inspired by [34], we pre-train a large StyleUnet-based teacher network to learn pose-dependent dynamic deformation maps of Gaussians in 2D space by front and back orthogonal projection. Although the teacher network is powerful enough to capture high-frequency appearance details in various poses, its large number of parameters makes it difficult to run in real-time on AR devices. To achieve high-performance rendering, we bake the dynamic deformation of Gaussians into the non-rigid deformation field of the mesh using a distillation technique, employing a compact MLP-based student network. Meanwhile, we propose two lightweight, learnable blend shapes to compensate for the non-rigid deformation of the Gaussians. After finetuning, our student model achieves high-performance rendering without compromising quality.

To evaluate our approach, we propose a multi-view dataset named **TalkBody4D**, focusing on prevalent fullbody talking scenarios encountered in daily life, which includes rich facial expressions and gestures with synchronous audios different from existing full-body motion datasets [11, 22, 69]. Experimental results demonstrate that our method effectively captures high-frequency appearance details of human avatars while maintaining a lightweight architecture capable of rendering stereo images at 2K resolution on common mobile and AR devices, including the Apple Vision Pro. The resulting full-body avatar is highly expressive, enabling users to animate it with facial expressions, hand gestures, and body poses, thereby highlighting the potential applications of our method. Our contributions can be summarized as follows:

- We introduce **TaoAvatar**, a novel teacher-student framework which yields topologically consistent, well-aligned geometry, further creating a high-fidelity, lightweight, 3DGS-based drivable full-body talking avatar.
- We propose two innovative strategies including non-rigid deformation **baking** and two lightweight blend shapes **compensations** to ensure high-quality rendering and efficient performance on mobile and AR devices, achieving **2K** resolution rendering at **90 FPS** on high-definition stereo devices like the Apple Vision Pro.
- We contribute TalkBody4D, a multi-view dataset to be released, designed for full-body talking scenarios, featuring diverse facial expressions and gestures with synchronous audios. Extensive experiments demonstrate that our approach outperforms other state-of-the-art methods in both quality and performance.

# 2. Related work

3D Avatar Representation. Traditional computer graphics techniques [12, 19, 50, 51] use 3D scanning or multiview stereo (MVS) to create full-body meshes, which are then combined with a 3D skeleton to produce driveable templates. However, this approach is both expensive and inefficient. As parametric body models, SMPL [37] and SM-PLX [43] are widely used for their versatility but lack detail in simulating loose clothing. Some works [4, 5, 27] have attempted to add clothing offsets to SMPLX, effectively handling tight clothing but struggling with loose garments. Due to the limited expressiveness of traditional meshes, researchers have combined SMPL with advanced rendering techniques. Some methods [7, 15, 23, 36, 57] use NeRF [40] to implicitly express human bodies, mapping points from the observation space backward to canonical space. Although high-quality, volume rendering is slow and poses challenges for real-time applications. In contrast, 3DGS [29] uses 3D Gaussian distributions for scenes, offering real-time rendering [38, 58], fast training, and the ability to handle complex materials such as hair and translucency. Some methods [34, 46, 48, 56, 63] combining SM-PLX [43] and 3DGS [29], like GaussianAvatar [20], use explicit GS point clouds and forward skinning to efficiently simulate the deformations, offering good rendering quality and speed. This representation presents potential solutions for achieving both real-time performance and computational efficiency.

**Dynamic Nonrigid Deformation Modeling.** Various methods have been developed to simulate non-rigid human deformation realistically. The MLP-based methods [21, 30, 33, 41], like 3DGS-Avatar [46], are noted for lightweight

flexibility but struggle with handling novel poses. Physical simulation methods [59, 68] produce physics-based animations but require unified topology, and complex modeling, and are computationally expensive, which negatively impacts real-time performance. Additionally, some methods [17, 18] use GNNs to learn non-rigid deformations in a datadriven manner. These methods are faster than traditional cloth simulations. Recently, dynamic texture methods are emerging, these methods [26, 32, 36] model complex details in the 2D texture space, while AnimatableGS [34] and Me-shAvatar [10] enhance modeling with front and back projection strategies which avoid texture unwarping.

**3D Drivable Talking Avatars.** Research on drivable digital humans is rapidly advancing for various applications. In 3D Talking Heads, methods [13, 45, 47, 60, 61, 66] excel in capturing and driving facial expressions, enhancing the realism of virtual characters. For drivable bodies, these methods [8, 20, 34, 56, 70] achieve precise restoration of body movements by learning pose-dependent deformation fields. Furthermore, some methods [49, 59] not only support detailed body modeling and animation but also enable head expressions and lip-syncing. Despite enhancing realism, running these methods in real-time on mobile or AR devices is challenging due to their high computational demands, exceeding current edge platform capabilities. Optimizing computational efficiency without compromising high-quality output is a key research focus.

# 3. Method

Leveraging multi-view RGB videos of a performer alongside corresponding SMPLX registrations for each frame, we aim to develop a high-fidelity, lightweight full-body talking avatar capable of high-resolution and high-performance rendering on mobile devices. To achieve this, we first create a hybrid clothed parametric representation by integrating 3D Gaussian Splatting (3DGS) [29] and SMPLX [43], which is more expressive for simulating loose clothing and hair, while maintaining high-quality rendering (Sec. 3.1). Next, we propose a teacher-student framework to reconstruct high-frequency, pose-dependent dynamic details as illustrated in Fig. 2. By leveraging non-rigid deformation baking and incorporating two lightweight blend shape compensations, our method achieves a tradeoff between high quality and high performance (Sec. 3.2).

# 3.1. Hybrid Clothed Parametric Representations

**The Clothed Extension of SMPLX.** To extend SMPLX to complete clothed human geometry with clothes, hair, etc., we choose a frame close to T-pose as a reference, which provides more visible details and less sticky geometry. First, we utilize NeuS2 [54] to reconstruct the geometry of this reference frame and parse out non-body components (e.g., clothes) from it using [53]. However, the body's skeleton

cannot directly drive these components. We estimate the shape and pose of the SMPLX [37] in the current frame and use robust skinning transfer [3] to propagate the body's skinning weights to these non-body components. Finally, we apply inverse skinning to transform these components back to the standard T-pose, resulting in a comprehensive personalized parametric model, referred to as SMPLX++.

Binding Gaussians on Mesh as Texture. After obtaining the complete geometric model, we bind 3D Gaussians to the mesh's triangles as textures, enabling them to move synchronously with the mesh. Inspired by these hybrid representations [45, 48], we define the attributes of the Gaussians in a local coordinate system relative to the triangles. Specifically, for each triangle, we randomly initialize k Gaussians, each of which maintains the local attributes  $\{f, (u, v), \gamma, \mathbf{r}, \mathbf{s}, o, \mathbf{sh}\}$ , where f is the index of the parent triangle, (u, v) is the barycentric coordinates,  $\gamma$  represents the translation along the triangle's normal direction, r and s denote rotation and scaling in the local space, o indicates opacity, and sh consists of spherical harmonic coefficients. Additionally, we define the normal of each Gaussian as  $\mathbf{n}_q = [1, 0, 0]$  within the local space, ensuring consistency with the mesh's normal for subsequent transformations. To deform each Gaussian from the local to the world space, we construct a transformation based on its parent triangle  $F_f(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ , where  $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$  are the vertices of the triangle  $F_f$  in the world space:

$$\mathbf{p} = u \cdot \mathbf{v}_{1} + v \cdot \mathbf{v}_{2} + (1 - u - v) \cdot \mathbf{v}_{3},$$
  

$$\mathbf{R} = [\mathbf{n}, \mathbf{q}, \mathbf{n} \times \mathbf{q}], \mathbf{q} = \frac{(\mathbf{v}_{2} + \mathbf{v}_{3})/2 - \mathbf{v}_{1}}{\|(\mathbf{v}_{2} + \mathbf{v}_{3})/2 - \mathbf{v}_{1}\|},$$
 (1)  

$$e = (\|\mathbf{v}_{1} - \mathbf{v}_{2}\| + \|\mathbf{v}_{2} - \mathbf{v}_{3}\| + \|\mathbf{v}_{1} - \mathbf{v}_{3}\|)/3,$$

where  $\mathbf{n} = \frac{(\mathbf{v}_1 - \mathbf{v}_2) \times (\mathbf{v}_3 - \mathbf{v}_1)}{\|\mathbf{v}_1 - \mathbf{v}_2\| \cdot \|\mathbf{v}_3 - \mathbf{v}_1\|}$  is the normal of the triangle, **p** denotes the surface point on the parent triangle, and *e* denotes the average edge length.

$$\mathbf{u}_w = \mathbf{p} + \gamma \mathbf{R} \mathbf{n}, \quad \mathbf{r}_w = \mathbf{R} \mathbf{r}, \quad \mathbf{s}_w = e \cdot \mathbf{s}, \quad (2)$$

here,  $\mathbf{u}_w$ ,  $\mathbf{r}_w$ ,  $\mathbf{s}_w$  represent the position, rotation, and scaling in the world spaces after transformation, respectively. Furthermore,  $\mathbf{c}_w = SH(\mathbf{sh}, \mathbf{R}^{-1}\mathbf{d})$  denotes the color along the view direction  $\mathbf{d}$  in the world space.

#### 3.2. Dynamic Mesh-based Gaussian Reconstruction

Learning Dynamic Non-Rigid Gaussian Deformation Maps. Although SMPLX++ can be directly driven by the skeleton, linear blend skinning is insufficient for performing dynamic non-rigid deformation such as clothing folds and swaying motions. Like dynamic Gaussian avatars [26, 34], we utilize a large StyleUnet [52] as the teacher network to capture complex pose-dependent dynamic non-rigid deformation of Gaussians in 2D texture space. Following [34],



Figure 2. **Illustration of our Method.** Our pipeline begins by reconstructing (a) a cloth-extended SMPLX mesh with aligned Gaussian textures. To address complex dynamic non-rigid deformations, (b) we employ a teacher StyleUnet to learn pose-dependent non-rigid maps, which are then baked into a lightweight student MLP to infer non-rigid deformations of our template mesh. For high-fidelity rendering, (c) we introduce learnable Gaussian blend shapes to enhance appearance details.

we rasterize the T-pose mesh, which is colored using posed coordinates and segmentation colors, to obtain the front and back position maps { $\mathcal{P}_f, \mathcal{P}_b$ }. These position maps are then fed into the StyleUnet network along with view directions to generate the non-rigid deformation maps { $\Delta \mathcal{U}_f, \Delta \mathcal{U}_b$ } and other residual Gaussian attribute maps. We add these residuals to the Gaussian's local attributes and transform them into the world space according to Eq. (2). To train the StyleUnet-based teacher network  $\mathcal{T}$ , we adopt the losses including L1, D-SSIM [29], and perceptual loss [64]. Additionally, we introduce a normal loss  $\mathcal{L}_{nor}$ . Consequently, the total reconstruction loss  $\mathcal{L}_{rec}$  is defined as:

$$\mathcal{L}_{rec} = \mathcal{L}_1 + \lambda_{ssim} \mathcal{L}_{ssim} + \lambda_{lpips} \mathcal{L}_{lpips} + \lambda_{nor} \mathcal{L}_{nor}, \quad (3)$$

where  $\lambda_*$  is the loss weights. The normal loss is defined as  $\mathcal{L}_{nor} = ||N_t - N_g||$ , where  $N_t$  is the rendered normal image by the teacher network, and the ground truth normal image  $N_g$  can be obtained from the method [54].

Baking Non-Rigid Gaussian Deformation into Mesh Deformation. Although the StyleUnet-based network  $\mathcal{T}$  can achieve good results, it struggles to run in real-time on mobile devices due to its large number of parameters, especially for AR devices, which require high resolution, stereo rendering, and high performance. Inspired by knowledge distillation, we bake the Gaussian non-rigid deformation of the teacher network into the mesh non-rigid deformation field, which is a compact MLP-based student network S.

$$\Delta \bar{\mathbf{v}}_i = \mathcal{S}\left(\bar{\mathbf{v}}_i, \theta, \mathbf{z}_t\right),\tag{4}$$

where  $\bar{\mathbf{v}}_i$  is the *i*-th vertex coordinate in the canonical space,  $\theta$  is the pose parameter, and  $\mathbf{z}_t$  is a learnable embedding for each frame to compensate for inaccurate pose estimation. To train the student network S with the capability of pose-dependent non-rigid deformation, we directly supervise its output using the Gaussian non-rigid deformation maps  $\{\Delta \mathcal{U}_f, \Delta \mathcal{U}_b\}$  of the teacher network  $\mathcal{T}$ . Specifically, we render the mesh non-rigid deformations  $\{\Delta \mathbf{v}_i\}_{i=1}^N$ to front and back deformation maps  $\{\Delta \mathcal{V}_f, \Delta \mathcal{V}_b\}$  in the canonical space, and then calculate the differences:

$$\mathcal{L}_{non} = \left\| \Delta \mathcal{V}_f - \Delta \mathcal{U}_f \right\| + \left\| \Delta \mathcal{V}_b - \Delta \mathcal{U}_b \right\|.$$
(5)

Furthermore, to mitigate the intersection between clothing and the body, we introduce a semantic loss  $\mathcal{L}_{sem}$ . For each vertex, we assign a semantic label  $\mathbf{e}_i = \mathbf{c}_i + \sin(\tau \bar{\mathbf{v}}_i)$  that integrates the candidate coordinates  $\bar{\mathbf{v}}_i$  and the artificial segmentation color  $\mathbf{c}_i$ , and  $\tau$  is a scale factor. The semantic label for each Gaussian can be obtained by interpolating the semantic label of the three vertices of its parent triangle. Subsequently, we render the Gaussian semantic map  $\mathcal{E}_t$  and the mesh semantic map  $\mathcal{E}_s$  in the world space separately and compute the semantic loss  $\mathcal{L}_{sem}$ :

$$\mathcal{L}_{sem} = \left\| \mathcal{E}_s - \mathcal{E}_t \right\|. \tag{6}$$

In addition, we refine the local attributes by computing the reconstruction loss  $\mathcal{L}_{rec}$ , which takes the rendered image of the teacher network as ground truth. The total loss  $\mathcal{L}_{bak}$  during baking is:

$$\mathcal{L}_{bak} = \mathcal{L}_{rec} + \lambda_{non} \mathcal{L}_{non} + \lambda_{sem} \mathcal{L}_{sem}.$$
 (7)

**Compensating Gaussian Deformation with Lightweight Blend Shapes.** Inspired by pose corrective blend shapes in SMPL [37], we build learnable blend shapes for position and color for each Gaussian. The position blend shapes,  $\mathbf{U} \in \mathbb{R}^{3 \times n}$ , primarily compensate for Gaussian positional adjustments (e.g., hair fluttering, clothing folds) across different poses. Meanwhile, the color blend shapes,  $\mathbf{C} \in \mathbb{R}^{3 \times n}$ , address appearance changes, such as shadow variations caused by self-obscuration. To obtain the corresponding driving coefficients, we employ a mapping network  $\mathcal{H}$  for the head using the expression parameters  $\epsilon$  as input and a mapping network  $\mathcal{B}$  for the body which inputs the body pose parameters  $\theta$ , to get the coefficients  $\mathbf{z}_h \in \mathbb{R}^{n_h}$ for the head and  $\mathbf{z}_b \in \mathbb{R}^{n_b}$  for the body, respectively, enabling independent and fine-grained control over head and body deformations. The compensation of position and color for each Gaussian can be expressed as:

$$\delta \mathbf{u} = BS\left(\mathbf{z}_h \oplus \mathbf{z}_b; \mathbf{U}\right), \delta \mathbf{c} = BS\left(\mathbf{z}_h \oplus \mathbf{z}_b; \mathbf{C}\right),$$
(8)

where  $\oplus$  is a concatenation operation, resulting in the size  $n = n_h + n_b$ . The function  $BS(\cdot, \cdot)$  represents the blend shapes operation mentioned in SMPL [37]. We add these residuals to local attributes:

$$\mathbf{u}_{w} = \mathbf{p} + \mathbf{R} \left( \gamma \mathbf{n} + \delta \mathbf{u} \right),$$
  
$$\mathbf{c}_{w} = SH \left( \mathbf{R}^{-1} \mathbf{d} \right) + \delta \mathbf{c}.$$
 (9)

For the fine-tuning stage on training data, we freeze the nonrigid deformation field network S, optimize the two mapping networks H and B, and two blend shapes U and C. The loss during fine-tuning is the same as that pre-training the teacher network.

After baking mesh non-rigid deformation and applying Gaussian non-rigid compensation, the performance of our student model is significantly enhanced compared to the teacher model, while still maintaining high quality.

# 4. Experiments

### 4.1. Experimental Settings

**Datasets.** To evaluate the performance on the full-body talking task, we introduce a new dataset, **TalkBody4D**, designed for common full-body talking scenarios in everyday life, primarily characterized by rich mouth movements and diverse gestures with synchronous audios. The dataset comprises 4 distinct identities, each wearing 2 different outfits.

Each talking sequence consists of rough 6k frames with 4K image resolution, 60 views with 48 full-body, and 12 faceregion views. To evaluate performance on complex motions and expressions, we select 4 sequences (04, 05, 06, 08) from the **ActorHQ** [22] dataset, supplement with 2 dancing and 2 exaggerated expression sequences. The resolution of our training and testing images is  $1500 \times 2000$  for all experiments. SMPLX registrations for these sequences can be obtained by first initializing SMPLX parameters using existing tools [1, 2, 45], followed by fine-tuning the poses with a photometric loss to improve geometry alignment.

**Implementation Details.** The clothed parametric template has roughly 22k vertices and 45k faces in total, which contains additional 23k faces for clothes, hair, and shoes, compared to the naive SMPLX. We randomly initialize about 220k Gaussians in total to bind with the template as texture, for each triangle containing roughly 4 to 6, and optimize these Gaussians for 10k iterations using the multi-view images of the reference frame. We pre-train our StyleUnetbased[34] teacher network for 600k iterations, and set the loss weights  $\lambda_{ssim} = 0.2, \lambda_{lpips} = 0.01, \lambda_{nor} = 0.02$  and batch size is 1 during the training process. Then we bake the teacher network into a small 5-layer MLP-based student network, which set  $\lambda_{non} = 0.1, \lambda_{sem} = 1.0$  and optimize 30k iterations. Finally, we finetune the mapping networks and blend shapes for Gaussian non-rigid compensation on training data 100k iterations at the batch size of 4. The size of two blend shapes for position and color is n = 28, in which  $n_h = 8$  for the head and  $n_b = 20$  for the body.

#### 4.2. Comparison

We conduct an extensive comparison between our method and recent state-of-the-art full-body avatar methods include GaussianAvatar [20], 3DGS-Avatar [46], MeshAvatar [10], and AnimatableGS [34]. All of these methods learn posedependent non-rigid deformations in the canonical space and then utilize the skeleton-based rigid transformations of SMPL [37] or SMPLX [43] into observation space to model the dynamic human. We uniformly use SMPLX [43] across all these methods to ensure a fair comparison.

**Comparison on Full-body Talking.** We conduct experiments on typical full-body talking scenarios encountered in daily life, which include a variety of mouth movements and gestures. We provide the quantitative comparison on full-body talking as shown in Tab. 1. We also show the qualitative comparison in Fig. 3, and our method can generate more realistic and detailed results, especially in the face region. 3DGS-Avatar [46] directly learns two MLP-based networks to handle the pose-dependent non-rigid deformation of Gaussians but produces very blurry results due to the low-frequency bias of MLPs. GaussianAvatar [20] uses a 2D CNN network by encoding the position map in SMPLX's UV space to generate the non-rigid deformation



Figure 3. Qualitative comparisons on full-body talking tasks. Our method outperforms state-of-the-art methods by producing clearer clothing dynamics and enhanced facial details.

maps of Gaussians. However, it is restricted by the miniclothed topology of SMPLX, struggling to handle loose skirts. MeshAvatar [10] chooses mesh as a human representation, which is insufficient to express the geometry of details, such as hair, and glasses. AnimatableGS [34] can generate high-quality dynamic appearance and complex nonrigid deformation. However, it is restricted by the implicit template of learning from scratch, which makes it difficult to handle fine-grained expressions, like blinking and flashing teeth as shown in the face region in Fig. 3. Therefore we adopt the SMPLX++ model as the template whose face and hands retain the native SMPLX prior. However, due to the amount parameters of the teacher StyleUnet [52], it achieves poor performance and is hard to run on a mobile device in real time. Through baking non-rigid deformation, we can achieve 150+ FPS at the resolution of  $1500 \times 2000$  on Nvidia RTX4090 using only a small MLP-based mesh non-rigid deformation network and two lightweight learnable blend shapes as shown in Tab. 1.

**Comparison on Complex Motion and Expression.** Besides full-body talking scenes, our method is also capable of reconstructing complex motions as well as exaggerated expressions as shown in Fig. 4. We provide quantitative comparisons in Tab. 2, our student network achieves competitive results compared to the teacher network, outperforming other methods. Especially for exaggerated expressions, our method significantly outperforms others in reconstructing the face region as shown in Tab. 2. Moreover, our method

	Novel View			Novel Gestures and Expression			Speed
Model	PSNR↑	SSIM↑	LPIPS↓	<b>PSNR</b> ↑	SSIM↑	LPIPS↓	FPS
GaussianAvatar[20]	26.58 (23.57)	.9313 (.8159)	.10577 (.25242)	25.99 (23.15)	.9232 (.8092)	.12265 (.26207)	54
3DGS-Avatar[46]	28.91 (23.95)	.9411 (.8303)	.07984 (.20450)	26.46 (23.32)	.9157 (.8184)	.11804 (.21632)	55
MeshAvatar[10]	28.53 (24.55)	.9360 (.8083)	.09470 (.25572)	27.08 (23.58)	.9229 (.7965)	.10783 (.24947)	22
AnimatableGS[34]	32.50 (26.42)	.9599 (.8587)	.06695 (.19535)	28.05 (23.68)	.9328 (.8142)	.09191 (.22673)	16
Ours (Teacher)	<u>33.45</u> (27.01)	<u>.9649</u> ( <u>.8741</u> )	<b>.04986</b> ( <u>.15613</u> )	<u>28.28</u> (24.28)	<u>.9336</u> (.8291)	<b>.07385</b> ( <u>.18176</u> )	16
Ours (Student)	33.81 (27.80)	<b>.9689</b> ( <b>.8975</b> )	<u>.06437</u> ( <b>.14218</b> )	28.38 (24.99)	.9389 (.8525)	<u>.08874</u> ( <b>.13364</b> )	156

Table 1. Quantitative comparisons on full-body talking task. The results inside the parentheses are evaluated for the face area, and the inference speed is evaluated on Nvidia RTX4090 when rendering images at a resolution of  $1500 \times 2000$ .



Figure 4. Results in challenging scenarios. Our method can obtain high-quality reconstruction even for challenging pose and expressions.

Model	PSNR↑	SSIM↑	LPIPS↓
GaussianAvatar[20]	25.94 (24.33)	.9294 (.8251)	.10478 (.24179)
3DGS-Avatar[46]	30.04 (25.08)	.9403 (.8458)	.08471 (.18044)
MeshAvatar[10]	28.51 (24.94)	.9334 (.8100)	.08846 (.23517)
AnimatableGS[34]	31.81 (26.79)	.9493 (.8608)	.07586 (.19521)
Ours (Teacher)	32.80 (27.40)	<u>.9533</u> ( <u>.8768</u> )	<b>.05581</b> ( <u>.14996</u> )
Ours (Student)	<u>32.72</u> ( <u>27.35</u> )	.9579 (.8836)	<u>.07326</u> ( <b>.13914</b> )

Table 2. Quantitative comparisons about complex motions and expressions reconstruction. The results inside the parentheses are evaluated for face regions with exaggerated expressions.

can obtain high-quality normals as shown in Fig. 4, which can be used for image-based relighting.

**Comparison on Novel Expression and Novel Gesture.** Benefiting from our parametric human model SMPLX++, our full-body avatars can be expression-driven and gesturedriven. We provide quantitative comparisons of self-driven on full-body talking scenes as shown in Tab. 1. As shown in Fig. 5, we provide visualizations of different characters driven by the same skeleton. In addition, our model can accept expression parameter inputs generated from audio by UniTalker [14] as shown in Fig. 5. More demos are provided in the supplementary materials.

### 4.3. Ablation Study

In this subsection, we discuss the key contributions of our method. Additional quantitative results and experiments are available in the supplementary materials.

**The Impact of Template.** The choice of template has a crucial impact on our approach. If we directly use SMPLX as our geometric template, as shown in Tab. 3 (w SMPLX), leading to poor results with numerous artifacts, especially for loose skirts as seen in Fig. 6. After extending the geometry of SMPLX to a complete template with clothes, hair, and shoes, this quality is significantly improved as shown in Tab. 3 (w/o Mesh Non.), which allows our method to reconstruct more challenging clothes.

The Impact of Non-rigid Deformation. Despite having a complete geometry, the rigid transformations of the skeleton are not able to properly represent non-rigid transformations such as skirt swinging and hair flying. We concurrently introduce mesh non-rigid deformation and Gaussian non-rigid deformation in our method. The geometry with-



Figure 5. Novel pose and expression animation. TaoAvatar can be driven by the same skeleton and expression parameters vividly.

out mesh non-rigid deformation is unable to attach to the surface of the character as shown in Fig. 6 (w/o Mesh Non.). The Gaussian non-rigid compensation also plays a crucial role in quantitative metrics. Without the blend shapes for Gaussian non-rigid compensation, the results are prone to produce artifices due to imprecise geometry and variable appearance as shown in Fig. 6 (w/o Gau Non.).

**The Impact of Baking Teacher.** Compared to directly training a small MLP-based network to handle non-rigid deformations as shown in Tab. 3 (w/o Teacher), our proposed multi-stage training strategy is more efficient. With direct supervision of the non-rigid deformation of the teacher network, it is easier for the student network to decouple the geometry deformation and appearance variations.

# 5. Application

TaoAvatar offers a lightweight, comprehensive representation for 3D talking body scenarios, easily deployable on various mobile devices using existing tools [25] as demonstrated in Fig. 1. It seamlessly integrates with AI models to enable real-time dialogue interactions. we deployed a 3D digital human agent on the Apple Vision Pro, which interacts with users through an ASR-LLM-TTS pipeline [16, 31,



Figure 6. Ablation Study. Red and Green boxes show artifacts and their improved counterparts.

Model	<b>PSNR</b> ↑	SSIM↑	LPIPS↓	P2S↓	Chamfer↓
w SMPLX	28.47	.9476	.07899	.7690	.9995
w/o Mesh Non.	32.10	.9734	.03814	.4877	.5007
w/o Gau Non.	31.16	.9686	.03932	.2968	.3068
w/o Teacher	32.67	.9751	.03769	.5236	.5359
Full	33.29	.9772	.03464	.2953	.3052

Table 3. **Ablation Study**. We use the mesh reconstructed by NeuS2 [54] as a pseudo-truth for geometry evaluation.

39, 62]. Facial expressions and gestures are dynamically controlled by an Audio2BS model [14], allowing the agent to respond naturally with synchronized speech, expressions, and movements. A live demonstration is available in our supplementary materials.

# 6. Discussion

**Conclusion.** In this work, we present TaoAvatar, a lightweight, lifelike full-body talking avatar solution. We demonstrate how the teacher-student framework captures high-definition facial and body details while ensuring real-time performance on AR devices. TaoAvatar can be driven by diverse signals, including facial expressions, hand gestures, and body poses. Through both quantitative and qualitative evaluations, we showcase the advantages of our approach. Additionally, we validate its practical potential with a real-world application on the Apple Vision Pro.

Limitations and Future Works. TaoAvatar encounters challenges in modeling flexible clothing deformation under exaggerated body poses, which are out-of-distribution of training data. A possible solution is to integrate GNN simulators [17, 18] handling larger hemlines, which are compatible with our approach.

**Potential Social Impact.** TaoAvatar can synthesize lifelike talking digital humans within an augmented reality environment, generating fabricated 3D content or 2D videos. Therefore, responsible use of this technology is essential.

### References

- [1] https://github.com/zju3dv/EasyMocap. 5,13
- [2] https://github.com/ShenhanQian/VHAP. 5, 13
- [3] Rinat Abdrashitov, Kim Raichstat, Jared Monsen, and David Hill. Robust skin weights transfer via weight inpainting. In SIGGRAPH Asia 2023 Technical Communications, SA Technical Communications 2023, Sydney, NSW, Australia, December 12-15, 2023, pages 25:1–25:4. ACM, 2023. 3, 13
- [4] Thiemo Alldieck, Marcus A. Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, pages 8387– 8397. Computer Vision Foundation / IEEE Computer Society, 2018. 1, 2
- [5] Thiemo Alldieck, Marcus A. Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 1175–1186. Computer Vision Foundation / IEEE, 2019. 2
- [6] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019, pages 5419–5429. IEEE, 2019. 1
- [7] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, Xu Jia, and Huchuan Lu. Animatable neural radiance fields from monocular rgb videos, 2021. 2
- [8] Xu Chen, Yufeng Zheng, Michael J. Black, Otmar Hilliges, and Andreas Geiger. SNARF: differentiable forward skinning for animating non-rigid neural implicit shapes. In 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021, pages 11574–11584. IEEE, 2021. 2, 3
- [9] Xu Chen, Tianjian Jiang, Jie Song, Max Rietmann, Andreas Geiger, Michael J Black, and Otmar Hilliges. Fast-snarf: A fast deformer for articulated neural fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10): 11796–11809, 2023. 2
- [10] Yushuo Chen, Zerong Zheng, Zhe Li, Chao Xu, and Yebin Liu. Meshavatar: Learning high-quality triangular human avatars from multi-view videos. *CoRR*, abs/2407.08414, 2024. 2, 3, 5, 6, 7, 13, 14, 16
- [11] Wei Cheng, Ruixiang Chen, Siming Fan, Wanqi Yin, Keyu Chen, Zhongang Cai, Jingbo Wang, Yang Gao, Zhengming Yu, Zhengyu Lin, Daxuan Ren, Lei Yang, Ziwei Liu, Chen Change Loy, Chen Qian, Wayne Wu, Dahua Lin, Bo Dai, and Kwan-Yee Lin. Dna-rendering: A diverse neural actor repository for high-fidelity human-centric rendering. In *IEEE/CVF International Conference on Computer Vision*, *ICCV 2023, Paris, France, October 1-6, 2023*, pages 19925– 19936. IEEE, 2023. 2
- [12] Edilson De Aguiar, Christian Theobalt, Carsten Stoll, and Hans-Peter Seidel. Marker-less deformable mesh tracking

for human shape and motion capture. In 2007 IEEE conference on computer vision and pattern recognition, pages 1–8. IEEE, 2007. 2

- [13] Helisa Dhamo, Yinyu Nie, Arthur Moreau, Jifei Song, Richard Shaw, Yiren Zhou, and Eduardo Pérez-Pellitero. Headgas: Real-time animatable head avatars via 3d gaussian splatting. In *ECCV* (2), pages 459–476. Springer, 2024. 3
- [14] Xiangyu Fan, Jiaqi Li, Zhiqian Lin, Weiye Xiao, and Lei Yang. Unitalker: Scaling up audio-driven 3d facial animation through A unified model. *CoRR*, abs/2408.00762, 2024. 7, 8, 14
- [15] Yao Feng, Weiyang Liu, Timo Bolkart, Jinlong Yang, Marc Pollefeys, and Michael J. Black. Learning disentangled avatars with hybrid 3d representations. *CoRR*, abs/2309.06441, 2023. 2
- [16] Zhifu Gao, Shiliang Zhang, Ian McLoughlin, and Zhijie Yan. Paraformer: Fast and accurate parallel transformer for nonautoregressive end-to-end speech recognition. In 23rd Annual Conference of the International Speech Communication Association, Interspeech 2022, Incheon, Korea, September 18-22, 2022, pages 2063–2067. ISCA, 2022. 8, 14
- [17] Artur Grigorev, Michael J. Black, and Otmar Hilliges. HOOD: hierarchical graphs for generalized modelling of clothing dynamics. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 16965–16974. IEEE, 2023. 3, 8
- [18] Artur Grigorev, Giorgio Becherini, Michael J. Black, Otmar Hilliges, and Bernhard Thomaszewski. Contourcraft: Learning to resolve intersections in neural multi-garment simulations. In ACM SIGGRAPH 2024 Conference Papers, SIG-GRAPH 2024, Denver, CO, USA, 27 July 2024- 1 August 2024, page 81. ACM, 2024. 3, 8
- [19] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. ACM Transactions on Graphics, 40(4), 2021. 2
- [20] Liangxiao Hu, Hongwen Zhang, Yuxiang Zhang, Boyao Zhou, Boning Liu, Shengping Zhang, and Liqiang Nie. Gaussianavatar: Towards realistic human avatar modeling from a single video via animatable 3d gaussians. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 634–644. IEEE, 2024. 2, 3, 5, 7, 14, 16
- Shoukang Hu, Tao Hu, and Ziwei Liu. Gauhuman: Articulated gaussian splatting from monocular human videos. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22,* 2024, pages 20418–20431. IEEE, 2024. 2
- [22] Mustafa Isik, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Nießner. Humanrf: High-fidelity neural radiance fields for humans in motion. ACM Trans. Graph., 42(4):160:1–160:12, 2023. 2, 5
- [23] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. Selfrecon: Self reconstruction your digital avatar from monocular video. In *IEEE/CVF Conference on Computer Vision and*

Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022, pages 5595–5605. IEEE, 2022. 2

- [24] Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 16922– 16932, 2023. 2
- [25] Xiaotang Jiang, Huan Wang, Yiliu Chen, Ziqi Wu, Lichuan Wang, Bin Zou, Yafeng Yang, Zongyang Cui, Yu Cai, Tianhang Yu, Chengfei Lyu, and Zhihua Wu. MNN: A universal and efficient inference engine. In *Proceedings of the Third Conference on Machine Learning and Systems, MLSys 2020, Austin, TX, USA, March 2-4, 2020.* mlsys.org, 2020. 8
- [26] Yujiao Jiang, Qingmin Liao, Xiaoyu Li, Li Ma, Qi Zhang, Chaopeng Zhang, Zongqing Lu, and Ying Shan. UV gaussians: Joint learning of mesh deformation and gaussian textures for human avatar modeling. *CoRR*, abs/2403.11589, 2024. 3
- [27] Yujiao Jiang, Qingmin Liao, Zhaolong Wang, Xiangru Lin, Zongqing Lu, Yuxi Zhao, Hanqing Wei, Jingrui Ye, Yu Zhang, and Zhijing Shao. Smplx-lite: A realistic and drivable avatar benchmark with rich geometry and texture annotations. In *IEEE International Conference on Multimedia* and Expo, ICME 2024, Niagara Falls, ON, Canada, July 15-19, 2024, pages 1–6. IEEE, 2024. 1, 2
- [28] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June* 16-22, 2024, pages 5322–5332. IEEE, 2024. 16
- [29] 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, 4
- [30] Muhammed Kocabas, Jen-Hao Rick Chang, James Gabriel, Oncel Tuzel, and Anurag Ranjan. HUGS: human gaussian splats. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June* 16-22, 2024, pages 505–515. IEEE, 2024. 2
- [31] Jungil Kong, Jihoon Park, Beomjeong Kim, Jeongmin Kim, Dohee Kong, and Sangjin Kim. VITS2: improving quality and efficiency of single-stage text-to-speech with adversarial learning and architecture design. In 24th Annual Conference of the International Speech Communication Association, Interspeech 2023, Dublin, Ireland, August 20-24, 2023, pages 4374–4378. ISCA, 2023. 8, 14
- [32] Youngjoong Kwon, Baole Fang, Yixing Lu, Haoye Dong, Cheng Zhang, Francisco Vicente Carrasco, Albert Mosella-Montoro, Jianjin Xu, Shingo Takagi, Daeil Kim, Aayush Prakash, and Fernando De la Torre. Generalizable human gaussians for sparse view synthesis. In ECCV (78), pages 451–468. Springer, 2024. 3
- [33] Jiahui Lei, Yufu Wang, Georgios Pavlakos, Lingjie Liu, and Kostas Daniilidis. GART: gaussian articulated template models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June* 16-22, 2024, pages 19876–19887. IEEE, 2024. 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR* 2024, Seattle, WA, USA, June 16-22, 2024, pages 19711– 19722. IEEE, 2024. 2, 3, 5, 6, 7, 13, 14, 15, 16
- [35] Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. GS-IR: 3d gaussian splatting for inverse rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22,* 2024, pages 21644–21653. IEEE, 2024. 16
- [36] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: neural free-view synthesis of human actors with pose control. *ACM Trans. Graph.*, 40(6):219:1–219:16, 2021. 2, 3
- [37] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multiperson linear model. ACM transactions on graphics (TOG), 34(6):1–16, 2015. 1, 2, 3, 5
- [38] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *3DV*, 2024. 2
- [39] Chengfei Lv, Chaoyue Niu, Renjie Gu, Xiaotang Jiang, Zhaode Wang, Bin Liu, Ziqi Wu, Qiulin Yao, Congyu Huang, Panos Huang, Tao Huang, Hui Shu, Jinde Song, Bin Zou, Peng Lan, Guohuan Xu, Fei Wu, Shaojie Tang, Fan Wu, and Guihai Chen. Walle: An end-to-end, generalpurpose, and large-scale production system for device-cloud collaborative machine learning. In *16th USENIX Symposium on Operating Systems Design and Implementation, OSDI* 2022, Carlsbad, CA, USA, July 11-13, 2022, pages 249–265. USENIX Association, 2022. 8
- [40] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 2, 13, 16
- [41] Arthur Moreau, Jifei Song, Helisa Dhamo, Richard Shaw, Yiren Zhou, and Eduardo Pérez-Pellitero. Human gaussian splatting: Real-time rendering of animatable avatars. In *CVPR*, pages 788–798. IEEE, 2024. 2
- [42] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. ACM transactions on graphics (TOG), 41(4):1–15, 2022. 2
- [43] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf.* on Computer Vision and Pattern Recognition (CVPR), pages 10975–10985, 2019. 1, 2, 3, 5, 14
- [44] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14314–14323, 2021. 1
- [45] Shenhan Qian, Tobias Kirschstein, Liam Schoneveld, Davide Davoli, Simon Giebenhain, and Matthias Nießner. Gaus-

sianavatars: Photorealistic head avatars with rigged 3d gaussians. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 20299–20309. IEEE, 2024. 3, 5

- [46] Zhiyin Qian, Shaofei Wang, Marko Mihajlovic, Andreas Geiger, and Siyu Tang. 3dgs-avatar: Animatable avatars via deformable 3d gaussian splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR* 2024, Seattle, WA, USA, June 16-22, 2024, pages 5020– 5030. IEEE, 2024. 2, 5, 7, 14, 16
- [47] Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. Relightable gaussian codec avatars. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 130–141. IEEE, 2024. 3
- [48] Zhijing Shao, Zhaolong Wang, Zhuang Li, Duotun Wang, Xiangru Lin, Yu Zhang, Mingming Fan, and Zeyu Wang. Splattingavatar: Realistic real-time human avatars with mesh-embedded gaussian splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR* 2024, Seattle, WA, USA, June 16-22, 2024, pages 1606– 1616. IEEE, 2024. 2, 3
- [49] Kaiyue Shen, Chen Guo, Manuel Kaufmann, Juan Jose Zarate, Julien Valentin, Jie Song, and Otmar Hilliges. Xavatar: Expressive human avatars. In *IEEE/CVF Conference* on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023, pages 16911– 16921. IEEE, 2023. 3
- [50] Carsten Stoll, Juergen Gall, Edilson De Aguiar, Sebastian Thrun, and Christian Theobalt. Video-based reconstruction of animatable human characters. ACM Transactions on Graphics (TOG), 29(6):1–10, 2010. 2
- [51] Daniel Vlasic, Pieter Peers, Ilya Baran, Paul E. Debevec, Jovan Popovic, Szymon Rusinkiewicz, and Wojciech Matusik. Dynamic shape capture using multi-view photometric stereo. *ACM Trans. Graph.*, 28(5):174, 2009. 2
- [52] Lizhen Wang, Xiaochen Zhao, Jingxiang Sun, Yuxiang Zhang, Hongwen Zhang, Tao Yu, and Yebin Liu. Styleavatar: Real-time photo-realistic portrait avatar from a single video. In ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH 2023, Los Angeles, CA, USA, August 6-10, 2023, pages 67:1–67:10. ACM, 2023. 3, 6, 15
- [53] Wenbo Wang, Hsuan-I Ho, Chen Guo, Boxiang Rong, Artur Grigorev, Jie Song, Juan Jose Zarate, and Otmar Hilliges. 4d-dress: A 4d dataset of real-world human clothing with semantic annotations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 550–560. IEEE, 2024. 3, 13
- [54] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 3272–3283. IEEE, 2023. 2, 3, 4, 8, 13, 15, 16
- [55] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4): 600–612, 2004. 15

- [56] Jing Wen, Xiaoming Zhao, Zhongzheng Ren, Alexander G. Schwing, and Shenlong Wang. Gomavatar: Efficient animatable human modeling from monocular video using gaussians-on-mesh. In *CVPR*, pages 2059–2069. IEEE, 2024. 2, 3
- [57] Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 16189–16199. IEEE, 2022. 2
- [58] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20310– 20320, 2024. 2
- [59] Donglai Xiang, Timur M. Bagautdinov, Tuur Stuyck, Fabian Prada, Javier Romero, Weipeng Xu, Shunsuke Saito, Jingfan Guo, Breannan Smith, Takaaki Shiratori, Yaser Sheikh, Jessica K. Hodgins, and Chenglei Wu. Dressing avatars: Deep photorealistic appearance for physically simulated clothing. *ACM Trans. Graph.*, 41(6):222:1–222:15, 2022. 3
- [60] Jun Xiang, Xuan Gao, Yudong Guo, and Juyong Zhang. Flashavatar: High-fidelity head avatar with efficient gaussian embedding. In *CVPR*, pages 1802–1812. IEEE, 2024. 3
- [61] Yuelang Xu, Bengwang Chen, Zhe Li, Hongwen Zhang, Lizhen Wang, Zerong Zheng, and Yebin Liu. Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 1931–1941. IEEE, 2024. 3
- [62] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report. *CoRR*, abs/2407.10671, 2024. 8, 14
- [63] Ye Yuan, Xueting Li, Yangyi Huang, Shalini De Mello, Koki Nagano, Jan Kautz, and Umar Iqbal. Gavatar: Animatable 3d gaussian avatars with implicit mesh learning. In *CVPR*, pages 896–905. IEEE, 2024. 2
- [64] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, pages 586– 595. Computer Vision Foundation / IEEE Computer Society, 2018. 4, 15

- [65] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Efficiently generated human radiance field from sparse inputs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7743–7753, 2022. 1
- [66] Zhongyuan Zhao, Zhenyu Bao, Qing Li, Guoping Qiu, and Kanglin Liu. Psavatar: A point-based shape model for realtime head avatar animation with 3d gaussian splatting, 2024. 3
- [67] Peng Zheng, Dehong Gao, Deng-Ping Fan, Li Liu, Jorma Laaksonen, Wanli Ouyang, and Nicu Sebe. Bilateral reference for high-resolution dichotomous image segmentation. *CoRR*, abs/2401.03407, 2024. 15
- [68] Yang Zheng, Qingqing Zhao, Guandao Yang, Wang Yifan, Donglai Xiang, Florian Dubost, Dmitry Lagun, Thabo Beeler, Federico Tombari, Leonidas J. Guibas, and Gordon Wetzstein. Physavatar: Learning the physics of dressed 3d avatars from visual observations. *CoRR*, abs/2404.04421, 2024. 3
- [69] Zerong Zheng, Xiaochen Zhao, Hongwen Zhang, Boning Liu, and Yebin Liu. Avatarrex: Real-time expressive fullbody avatars. ACM Trans. Graph., 42(4):158:1–158:19, 2023. 2
- [70] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Instant volumetric head avatars. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 4574–4584. IEEE, 2023. 3