

# GaussianAvatar: Towards Realistic Human Avatar Modeling from a Single Video via Animatable 3D Gaussians

Liangxiao Hu<sup>†,1</sup>, Hongwen Zhang<sup>2</sup>, Yuxiang Zhang<sup>3</sup>, Boyao Zhou<sup>3</sup>, Boning Liu<sup>3</sup>,  
 Shengping Zhang<sup>\*,1,4</sup>, Liqiang Nie<sup>1</sup>

<sup>1</sup>Harbin Institute of Technology <sup>2</sup>Beijing Normal University

<sup>3</sup>Tsinghua University <sup>4</sup>Peng Cheng Laboratory

{lx.hu, s.zhang}@hit.edu.cn, zhanghongwen@bnu.edu.cn, yx-zl9@mails.tsinghua.edu.cn  
 {bzhou22, liuboning}@mail.tsinghua.edu.cn, nieliqiang@gmail.com




Figure 1. We propose GaussianAvatar, which learns animatable 3D Gaussians to represent detailed human avatars from a single video. Our method maintains a 3D consistent appearance even when animated by out-of-distribution motions.

## Abstract

We present *GaussianAvatar*, an efficient approach to creating realistic human avatars with dynamic 3D appearances from a single video. We start by introducing animatable 3D Gaussians to explicitly represent humans in various poses and clothing styles. Such an explicit and animatable representation can fuse 3D appearances more efficiently and consistently from 2D observations. Our representation is further augmented with dynamic properties to support pose-dependent appearance modeling, where a dynamic appearance network along with an optimizable feature tensor is designed to learn the motion-to-appearance mapping. Moreover, by leveraging the differentiable motion

condition, our method enables a joint optimization of motions and appearances during avatar modeling, which helps to tackle the long-standing issue of inaccurate motion estimation in monocular settings. The efficacy of *GaussianAvatar* is validated on both the public dataset and our collected dataset, demonstrating its superior performances in terms of appearance quality and rendering efficiency. The code and dataset are available at <https://github.com/aipixel/GaussianAvatar>.

## 1. Introduction

Creating a customized human avatar from a single video has great potential for many applications including virtual and augmented reality, the Metaverse, gaming, and movie industries. This task is appealing yet challenging, as the

<sup>†</sup> Work done during an internship at Tsinghua University.

\* Corresponding author.

monocular observations are highly underdetermined for the modeling of a 3D animatable avatar. Meanwhile, the inaccurate body motion estimations and the complex wrinkle deformations also make it extremely difficult to create a realistic avatar from a single video.

The modeling of 3D human avatars from monocular videos involves a fusion process of 2D observations to a 3D consistent human model. For this purpose, existing solutions have leveraged both implicit and explicit representations to create the base model of avatars. Among them, implicit methods [6, 15, 16, 49, 57] define a deformable human NeRF to fuse the image observation from current motion space to a canonical space by inverse skinning. However, the implicit 3D volume is inefficient in representing human surfaces and the inverse skinning also introduces ambiguous correspondences during the fusion process. These issues make it hard for implicit solutions to capture fine-grained details of moving people. As the avatar appearances are concentrated around human surfaces, explicit representations are much more efficient in modeling avatars. Previous attempts [12, 63] have employed differentiable mesh rendering to reconstruct the human surface, but these methods struggle to preserve wrinkle details due to a fixed mesh topology. On the other hand, point-based representations [18, 37, 50, 56, 65] are more effective to handle flexible topology but require millions of points to capture detailed appearances. How to represent humans remains one of the fundamental problems for avatar modeling.

When creating an animatable avatar from a video of moving people, the algorithm is required to learn the relationships between body motions and corresponding appearances. However, the motions estimated from monocular videos are typically faulty, leading to large artifacts in the modeling of dynamic cloth deformations. To address this issue, previous works [1, 15, 16, 42, 45, 49, 57] have attempted to optimize body motions along with the learning of animatable avatar volumes. As the body motions are explicitly represented as parametric meshes, the implicit 3D volume of previous methods makes the optimization indirect and less effective. This issue has even become the main obstacle to achieving high-quality avatar modeling from monocular videos.

To tackle the above issues, we introduce new representations and solutions to achieve high-quality avatar modeling from a single video. The key insight of our solution is to model dynamic human surfaces explicitly and optimize both the motion and appearances jointly in an end-to-end manner. To this end, we propose GaussianAvatar, a method to reconstruct human avatars with dynamic appearances using the 3D Gaussian representation [17]. As an explicit representation, 3D Gaussian can be easily reposed from the canonical space to the motion space via a forward skinning process. Such an animatable 3D Gaussian

representation bypasses the inverse skinning process used in the aforementioned NeRF-based methods and overcomes the previous one-to-many issue [4] during canonicalization. Based on the animatable 3D Gaussian, our method can fuse 3D appearances more consistently from 2D observations to a canonical 3D space. To model dynamic human appearances under different poses, we additionally add pose-dependent properties to 3D Gaussians and incorporate them with the canonical human surfaces. Inspired by previous work [28, 46], we learn a dynamic appearance network on the 2D manifolds of the underlying human mesh to predict dynamic properties of 3D Gaussians. However, due to the strong bias of limited training poses, modeling dynamic appearance solely conditioned on pose information struggles to generalize to novel views and poses. To address this issue, we introduce an optimizable feature tensor to capture a coarse global appearance of human avatars. Subsequently, we incorporate pose-dependent effects on the feature tensor to decode fine-grained details such as wrinkles.

As the proposed animatable 3D Gaussians are differentiable with respect to the motion conditions, it enables a joint optimization of motion and appearances. This merit allows our network to refine the motion along with the avatar modeling process, which helps to tackle the long-standing issue of inaccurate motion estimation in monocular settings. Besides, the refined motion further enhances the accuracy of avatar modeling since the 3D appearance fusion process relies on motion-based skinning. As shown in our experiments, our method is quite robust to initial motion estimation and has the ability to correct the misalignment of motion capture results.

To summarize, our main contributions are as follows:

- We introduce animatable 3D Gaussians for realistic human avatar modeling from a single video. By representing human surfaces explicitly, our method can fuse 3D appearances more consistently and efficiently from 2D observations.
- We augment the animatable 3D Gaussians with dynamic properties to support pose-dependent appearance modeling, where a dynamic appearance network along with an optimizable feature tensor is designed to learn the motion-to-appearance mapping.
- We propose to jointly optimize the motion and appearance during the avatar modeling, enabling our method to correct the misalignment of initial motion and improve the final appearance quality.

## 2. Related Work

**Neural Rendering for Human Reconstruction.** Without the need to define a template mesh for avatar modeling [2, 9, 53], neural rendering has emerged as a potent technique that enables learning avatars directly from im-

ages. Here we briefly review precious work that aims to reconstruct humans using neural rendering.

Due to the high-quality rendering of neural radiance field [29], various efforts [20, 21, 24, 32, 33, 47, 67, 68] have been made to reconstruct the dynamic appearance of moving people. Neural Body [33] associates a latent code to each SMPL [26] vertex to encode the appearance, which is transformed into observation space based on the human pose. Neural Actor [24] learns a deformable radiance field with SMPL as guidance and utilizes a texture map to improve its final rendering quality. TAVA [20] proposes to jointly model the non-rigid warping field and shading effects directly conditioned on the pose vectors. Posevocab [21] designs joint-structured pose embeddings to encode the dynamic appearances under different key poses, such embeddings can better learn joint-related appearance. NeRF-based methods have demonstrated appealing rendering results on human avatar reconstruction, but still struggle to represent human surfaces with the implicit 3D volume. Explicit modeling of human surfaces is a more straightforward way for this task, as the dynamic appearance of humans is mostly reflected on the human surface.

Explicit representations have great potential for human reconstruction. In HF-Avatar [63], meshes are used as the base representation in a coarse-to-fine framework that combines neural texture with dynamic surface deformation for avatar creation. EMA [12] proposes Meshy neural fields to reconstruct human avatars by optimizing the canonical mesh, material, and motion dynamics through inverse rendering in an end-to-end process. PointAvatar [65] employs a deformable point-based representation to separate source color into intrinsic albedo and normal-dependent shading. DVA [35] extends mixtures of volumetric primitives [25] for human avatar modeling. All these attempts have demonstrated the significant potential of explicit representations and their under-exploration. However, meshes are constrained by fixed topologies, whereas point clouds demand a multitude of points to encompass intricate details. 3D Gaussians [17] have showcased their capability in various human tasks [22, 40, 55, 64]. In essence, 3D Gaussians hold promise for human avatar reconstruction and are currently a subject of active research.

**Avatar Modeling from Monocular Videos.** Numerous methods investigate reconstructing humans from single images or monocular videos. Regression-based methods [10, 11, 13, 38, 39, 51, 52, 66] directly recover clothed 3D humans just from a single image. While these methods produce attractive results, they can not recover the dynamic appearance in the reconstruction across the entire sequence. Traditional methods aim to capture human dynamics by tracking individuals in videos using pre-scanned rigged templates [7, 8, 54]. However, the pre-scanning

and manual rigging processes prevent their real-life applications. [1, 5, 30] try to bypass the requirement for predefined human models but face challenges in preserving fine details because of the fixed mesh resolution.

The emergence of neural radiance fields [29] has facilitated the creation of various techniques for reconstructing animatable avatars from monocular videos [3, 14, 16, 41–43, 49]. However, inaccuracies in estimating human motions from monocular videos result in pronounced artifacts. To address the problem, HumanNeRF [49] solves for an update to the inaccurate poses. NeuMan [16] introduces an error-correction network to enable training with erroneous estimates. Vid2Avatar [6] and InstantAvatar [15] also jointly optimize motions by back-propagating the gradient of the image reconstruction loss to the pose parameters. MonoHuman [57] introduces bi-directional constraints to alleviate ambiguous correspondence on novel poses. Attempting to address the inaccurate pose estimation issue with implicit human NeRF models has proven to be inefficient and imprecise. Consequently, we endeavor to address this issue by leveraging an explicit 3D Gaussian representation.

### 3. Method

Our goal is to create a human avatar that enables free-viewpoint rendering and realistic animation using a single video. We first introduce an expressive representation, namely animatable 3D Gaussians, to represent human avatars in Sec. 3.1. With this representation, modeling the dynamic appearance of humans can be regarded as dynamic 3D Gaussian property estimation. Then we build a dynamic appearance network along with an optimizable feature tensor to learn the motion-to-appearance mapping in Sec. 3.2. To mitigate the artifacts caused by inaccurate pose estimation, we adopt a joint motion and appearance optimization approach for refining human poses during training (Sec. 3.3). In the following, we provide a comprehensive explanation of the technical details.

#### 3.1. Animatable 3D Gaussians

Point-based representation [23, 28, 34, 59, 69] has proven its topological flexibility in generating realistic human avatars from scans. Extending this explicit representation to create human avatars from images is a significant endeavor. With this goal in mind, we introduce a novel representation, termed animatable 3D Gaussians, which effectively reconstructs human surfaces with a 3D consistent appearance.

3D Gaussian Splatting [17] is a point-based scene representation that allows high-quality real-time rendering. The scene representation is parameterized by a set of static 3D Gaussians, each of which has the following parameters: 3D center position  $\mathbf{x} \in \mathbb{R}^3$ , color  $\mathbf{c} \in \mathbb{R}^3$ , opacity  $\alpha \in \mathbb{R}$ , 3D rotation in form of quaternion  $\mathbf{q} \in \mathbb{R}^4$  and 3D scaling factor  $\mathbf{s} \in \mathbb{R}^3$ . With these properties, we can generate rendered

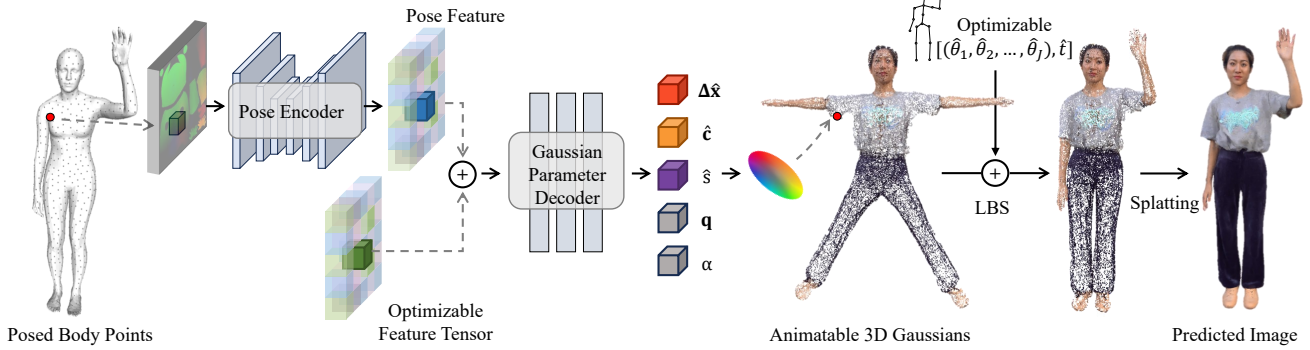


Figure 2. **Overview of GaussianAvatar.** Given a fitted SMPL or SMPL-X model on the current frame, we sample the points on its surface and record their positions on a UV positional map  $I$ , which is then passed to a pose encoder to obtain the pose feature. An optimizable feature tensor is pixel-aligned with the pose feature and learned to capture the coarse appearance of humans. Then the two aligned feature tensors are input into the Gaussian parameter decoder, which predicts each point’s offset  $\Delta\hat{\mathbf{x}}$ , color  $\hat{\mathbf{c}}$ , and scale  $\hat{s}$ . These predictions, along with the fixed rotations  $\mathbf{q}$  and opacity  $\alpha$ , collectively constitute the animatable 3D Gaussians in canonical space.

images from any viewpoint. We refer readers to [17, 27] for the rendering details of 3D Gaussians.

To extend this representation for human avatar modeling, we integrate it with either the SMPL [26] or SMPL-X [31] model as follows:

$$G(\beta, \theta, \mathbf{D}, \mathbf{P}) = \text{Splattting}(W(\mathbf{D}, J(\beta), \theta, \omega), \mathbf{P}), \quad (1)$$

where  $G(\cdot)$  represents a rendered image, and  $\text{Splattting}(\cdot)$  denotes the rendering process of 3D Gaussians from any viewpoint,  $W(\cdot)$  is a standard linear blend skinning function employed for reposing 3D Gaussians,  $\mathbf{D} = T(\beta) + dT$  represents the locations of 3D Gaussians in canonical space, formed by adding corrective point displacements  $dT$  on the template mesh surface  $T(\beta)$ ,  $\mathbf{P}$  denotes the remaining properties of 3D Gaussians, excluding the positions.  $\beta$  and  $\theta$  are the shape and pose parameters,  $J(\beta)$  outputs 3D joint locations. Note that we propagate the skinning weight  $\omega$  from the vertices of the SMPL or SMPL-X model to the nearest 3D Gaussians. With the proposed representation, we can now repose these canonical 3D Gaussians to the motion space for free-view rendering.

### 3.2. Dynamic 3D Gaussian Property Estimation

Following the proposed animatable Gaussians, human appearances are determined by the point displacements  $dT$  and properties  $\mathbf{P}$ . Modeling dynamic human appearances can be regarded as estimating these dynamic properties. To model dynamic human appearances under various poses, we introduce a dynamic appearance network along with an optimizable feature tensor to predict these pose-dependent properties of 3D Gaussians. Despite sharing similarities in network structure with [28], we present this framework for a distinct purpose. In [28], the feature tensor serves to decouple the pose-independent human shape from the decoder, and we repurpose it to capture a coarse global appearance




Figure 3. **Effect of iostropy of 3D Gaussians.** (a) Input image, (b)(d) front and back views trained with isotropic 3D Gaussians, (c)(e) front and back views trained with anisotropic 3D Gaussians.

of human avatars. The motivation behind this modification is that directly learning a mapping from human poses to dynamic properties is susceptible to overfitting on the limited training poses. To integrate the global appearance into the feature tensor, we introduce a two-stage training strategy, as discussed in Sec. 3.4.

The dynamic appearance network is designed to learn a mapping from a 2D manifold representing the underlying human shape to the dynamic properties of 3D Gaussians as follows:

$$f_\phi : \mathcal{S}^2 \in \mathbb{R}^3 \rightarrow \mathbb{R}^7. \quad (2)$$

As shown in Fig. 2, the 2D human manifold  $\mathcal{S}^2$  is depicted by a UV positional map  $I \in \mathbb{R}^{H \times W \times 3}$ , where each valid pixel stores the position  $(x, y, z)$  of one point on the posed body surface. The final predictions consist of per point offset  $\Delta\hat{\mathbf{x}} \in \mathbb{R}^3$ , color  $\hat{\mathbf{c}} \in \mathbb{R}^3$ , and scale  $\hat{s} \in \mathbb{R}$  on the canonical surface.

Instead of predicting all properties, we make these slight adjustments to our task based on experimental results. Due to the unbalanced viewpoints in monocular videos,

anisotropic 3D Gaussians are prone to learning an inaccurate 3D shape to fit the most frequently seen view, resulting in poorer performance in side views. As shown in Fig. 3, training our network with isotropic 3D Gaussians yields superior results. We ensure isotropy among all 3D Gaussians by maintaining uniform size across dimensions. Therefore, we introduce a scaler  $\hat{s}$  to represent the scales of 3D Gaussians and set rotations  $\mathbf{q}$  as  $[1, 0, 0, 0]$ . Through experimental observations, we notice that the network tends to learn an opacity value ( $\alpha$ ) of zero on the boundary to correct the human shape. To address this, we fix opacity  $\alpha = 1$  to keep all 3D Gaussians visible, enforcing the network to predict accurate positions of the 3D Gaussians.

The dynamic appearance network consists of two parts: a pose encoder and a Gaussian parameter decoder. The pose encoder takes the UV positional map of posed body points as input to generate a pose-conditioned feature tensor  $O \in \mathbb{R}^{H \times W \times C}$ . We then integrate the pixel-aligned optimized feature tensor  $F \in \mathbb{R}^{H \times W \times C}$  with the pose features before feeding it into the Gaussian parameter decoder to generate final predictions. Following this, we add the predicted offsets to the canonical 3D Gaussians and associate the predicted properties with the corresponding 3D Gaussians. With the estimated poses, we can then repose the canonical 3D Gaussians to the motion space for rendering.

### 3.3. Joint Motion and Appearance Optimization

However, owing to the imprecise estimation of human poses  $\theta = (\theta_1, \theta_2, \dots, \theta_J)$  and translations  $\mathbf{t}$  from monocular videos, the reposed 3D Gaussians in motion space are inaccurately represented and may lead to unsatisfactory rendering outcomes. To address this, we propose to jointly optimize human motions and appearances. To optimize human motions with image loss, we solve for an update  $(\Delta\theta, \Delta\mathbf{t})$  to the estimated body poses and translations as follows:

$$\hat{\Theta} = (\theta + \Delta\theta, \mathbf{t} + \Delta\mathbf{t}). \quad (3)$$

We modify  $\theta$  in Eq. 1 using  $\hat{\Theta}$  to render the proposed animatable 3D Gaussians differentiable with respect to the motion conditions. Different from previous work [6, 16, 49] that jointly optimizes human poses via inverse skinning, we optimize the updates in a forward skinning process, which benefits both motion and appearance optimization.

### 3.4. Training Strategy

In this section, we outline our approach to training the network with inaccurate human motions. We conduct a two-stage optimization process using different loss functions. In the first stage, we aim to fuse the sequential appearances to the optimizable feature tensor and conduct motion optimization to get accurate poses for the dynamic appearance network. In this stage, we optimize the framework without incorporating any pose-dependent information by excluding

the training of the pose encoder. Specifically, we utilize the following loss functions to train our network:

$$\mathcal{L}_{stage_1} = \lambda_{rbg} \mathcal{L}_{rbg} + \lambda_{ssim} \mathcal{L}_{ssim} + \lambda_{lpipe} \mathcal{L}_{lpipe} + \lambda_f \mathcal{L}_f + \lambda_{offset} \mathcal{L}_{offset} + \lambda_{scale} \mathcal{L}_{scale}, \quad (4)$$

where  $\mathcal{L}_{rbg}$ ,  $\mathcal{L}_{ssim}$ , and  $\mathcal{L}_{lpipe}$  are the L1 loss, SSIM loss [48], and LPIPS loss [61], respectively.  $\mathcal{L}_f$ ,  $\mathcal{L}_{offset}$ ,  $\mathcal{L}_{scale}$  calculate the L2-norm of the feature map, predicted offsets and scales, respectively. We set  $\lambda_{rbg} = 0.8$ ,  $\lambda_{ssim} = 0.2$ ,  $\lambda_{lpipe} = 0.2$ ,  $\lambda_f = 1$ ,  $\lambda_{offset} = 10$ ,  $\lambda_{scale} = 1$ .

After the first stage of training, we obtain more accurate human motions and an optimized feature tensor  $F$ . The optimized feature tensor  $F$  captures a coarse appearance of human avatars. In the second stage, we incorporate the pose features encoded by the pose encoder with the trained feature tensor  $F$ . We replace  $\mathcal{L}_f$  with the L2-norm loss  $\mathcal{L}_p$ , which plays the same role as  $\mathcal{L}_f$  in regularizing the limited pose space. By penalizing the pose-dependent features, we can eliminate the strong bias of limited training poses and thus generalize to unseen viewpoints and poses.

## 4. Experiments

### 4.1. Datasets and Metrics

**People-Snapshot Dataset.** This dataset [1] comprises videos of individuals rotating in front of a stationary camera. To ensure a fair quantitative comparison, we follow the same evaluation protocol outlined in InstantAvatar [15].

**NeuMan Dataset.** To assess even more challenging scenarios, we employ outdoor collections from the NeuMan dataset [16]. These videos are recorded using a mobile phone for moving individuals. Specifically, we select four sequences (bike, citron, seattle, jogging) that exhibit most body regions and contain minimal blurry images. We initialize the estimated poses with an off-the-shelf method [44], which is also utilized in [15].

**DynVideo Dataset.** With the limited cloth deformation presented in the above two datasets, we propose the Dynvideo dataset to capture dynamic human appearance. To this end, we utilize a mobile phone to record videos of a character performing various movements, especially rotation, in front of the device. Each of these videos takes about one minute and provides a comprehensive and detailed representation of human movement. We also provide the corresponding SMPL parameter sequences for all videos, obtained by our proposed method. This dataset serves as a valuable resource for evaluating reconstruction quality, with a particular emphasis on dynamic appearances.

**Evaluation Metrics.** We consider three metrics: PSNR, SSIM [48], and LPIPS [61] to access the reconstruction quality on three datasets.

Method	male-3-casual			male-4-casual			female-3-casual			female-4-casual		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
HumanNeRF [49]	26.90	0.9605	0.0181	25.50	0.9397	0.0357	24.46	0.9516	0.0269	27.07	0.9615	0.0152
InstantAvatar [15]	29.53	0.9716	0.0155	27.67	0.9626	0.0307	27.66	0.9709	0.0210	29.11	0.9683	0.0167
Baseline	27.71	0.9713	0.0218	25.09	0.9614	0.0306	24.80	0.9609	0.0310	26.16	0.9610	0.0227
Baseline + Opt.	<b>30.98</b>	<b>0.9791</b>	<b>0.0143</b>	28.77	0.9753	0.0230	<b>29.59</b>	<b>0.9768</b>	<b>0.0220</b>	30.83	0.9768	0.0145
Baseline + Opt. + Dyn.	<b>30.98</b>	0.9790	0.0145	<b>28.78</b>	<b>0.9755</b>	<b>0.0228</b>	29.55	0.9762	0.0225	<b>30.84</b>	<b>0.9771</b>	<b>0.0140</b>

Table 1. **Quantitative evaluation on the People-Snapshot [1] dataset.** Opt. denotes the motion optimization and Dyn. refers to dynamic appearance modeling.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
HumanNeRF [49]	27.06	0.9669	0.0192
InstantAvatar [15]	28.47	0.9715	0.0277
Baseline	27.06	0.9692	0.0183
Baseline + Opt.	29.93	0.9794	0.0126
Baseline + Opt. + Dyn.	<b>29.94</b>	<b>0.9795</b>	<b>0.0124</b>

Table 2. **Quantitative evaluation on NeuMan [16] dataset.**

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
HumanNeRF [49]	24.39	0.9349	0.0349
InstantAvatar [15]	20.31	0.9183	0.1046
Baseline	23.55	0.9340	0.0354
Baseline + Opt.	26.16	0.9522	0.0237
Baseline + Opt. + Dyn.	<b>28.58</b>	<b>0.9616</b>	<b>0.0169</b>

Table 3. **Quantitative evaluation on DynVideo dataset.**




Figure 4. **Motion optimization results.** (a)(d) Original image, (b)(e) our optimized SMPL, (c)(f) ROMP [44] estimates.

## 4.2. Experimental Settings

**Implementation details.** We employ a U-Net [36] for extracting pose-dependent features, and the Gaussian parameter decoder is implemented as an 8-layer multilayer perceptron (MLP). We sample approximately 200,000 points on the SMPL mesh surface. The entire framework is trained on a single NVIDIA RTX 3090 GPU, with training times ranging from 0.5 to 6 hours.

**Baseline.** To showcase the efficacy of the proposed mod-

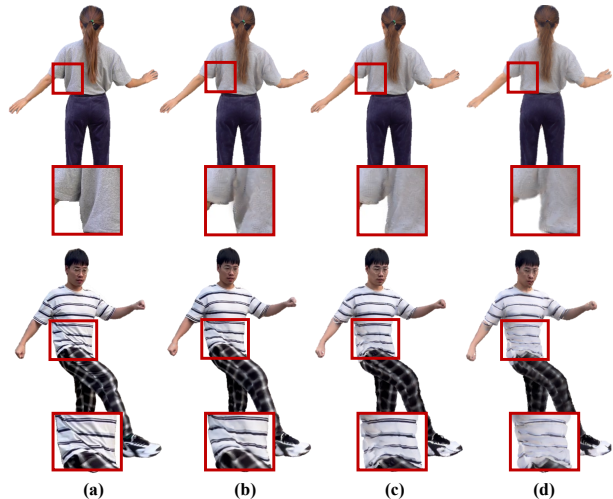


Figure 5. **Qualitative ablation studies.** (a) Ground truth, (b) baseline + Opt. + Dyn., (c) baseline + Opt., (d) baseline.

ules, our method can be partitioned into three components: baseline, motion optimization, and dynamic appearance modeling. As our baseline, we suspend the motion optimization process and exclude the pose-dependent information from the pose encoder.

**Methods for Comparison.** We compare our method against (1) HumanNeRF [49], which implicitly represents a human avatar with a canonical appearance neural field and a motion field; (2) InstantAvatar [15], which achieves fast avatar modeling by using several acceleration strategies for neural fields.

## 4.3. Comparisons with the State of the Art

We report the numeric evaluation results on three datasets. As shown in Table 1, 2, and 3, our proposed method outperforms all baselines on all metrics for recovering more details of dynamic appearance and correcting the artifacts caused by initialized poses. To demonstrate the qualitative evaluation of these datasets, we also visualize the novel view synthesis results on the test splits. As implicit representations, HumanNeRF and InstantAvatar are prone to generate ghosting effects at boundary areas, as shown in Fig. 6. InstantAvatar lacks the capability to model pose-



Figure 6. **Qualitative comparison of novel view synthesis.** We compare the novel view synthesis quality on the People-Snapshot dataset (first row), NeuMan dataset (second row), and DynVideo dataset (last two rows).

dependent deformations, leading to challenges in effectively handling the dynamic appearance of moving people.

To showcase the robustness of our avatar modeling, we collect more challenging poses [19] captured by the monocular camera and evaluate our method and other methods for avatar animation and novel view synthesis. As ground truth

is unavailable for these out-of-distribution poses, we illustrate qualitative results of novel pose and view synthesis in Fig. 1 and Fig. 7. Our method generates realistic animation results with these challenging poses, demonstrating a consistent 3D appearance in novel views with respect to other methods.




Figure 7. **Animation results on out-of-distribution motions.** We compare the animation results on the People-Snapshot dataset (first row), NeuMan dataset (second row), and DynVideo dataset (last row).

#### 4.4. Ablation Studies

In this section, we conduct ablation studies to validate each component of our methods. As shown in Table 1, 2, and 3, our proposed motion optimization module dramatically improves over the baseline on all metrics, demonstrating its effectiveness in modeling human avatars. To highlight the effectiveness of our method in human motion optimization, we illustrate the initialized poses obtained by ROMP [44] and the optimized ones by our method on NeuMan and DynVideo datasets in Fig. 4. Experiments demonstrate that joint motion optimization is capable of correcting inaccurate motion estimation, even for the side and back views. Furthermore, our approach readily extends to enhance the accuracy of existing motion capture methods [58, 60, 62].

We also notice that the dynamic appearance modeling achieves superior results on NeuMan and DynVideo datasets in Table 2 and 3 and slightly favorable outcomes in Table 1, considering that People-Snapshot has limited pose variation with respect to the other two datasets. In Fig. 5, we further illustrate the visible improvements of each component of our method. The motion optimization scheme improves the global quality of the baseline for better pose estimation. The dynamic appearance modeling additionally preserves pose-dependent details such as cloth wrinkles on the human surface.

#### 5. Conclusion and Discussion

We introduce GaussianAvatar, a human avatar reconstruction method based on the proposed animatable 3D Gaussians from monocular videos. For dynamic human appearance modeling, we leverage a dynamic appearance network along with an optimizable feature tensor to enhance the representation with dynamic properties. Besides, we implement a joint motion and appearance optimization scheme to rectify estimated motion and enhance the overall reconstruction quality. Our method shows the capability to reconstruct avatars with dynamic appearances, enabling realistic animation while maintaining real-time rendering speed.

**Limitation.** Similar to [15, 49, 57], our method may generate artifacts due to inaccurate foreground segmentations in videos and encounter challenges in modeling loose outfits such as dresses.

**Potential Social Impact.** Given our method’s capability to reconstruct a realistic personalized character from a monocular video, it is imperative to exercise caution and consider the potential for technology misuse.

**Acknowledgement.** This work was supported by the NSFC project (Nos. 62272134, 62236003, 62072141, 62301298, and 62125107), Shenzhen College Stability Support Plan (Grant No. GXWD20220817144428005), and the Major Key Project of PCL (PCL2023A10-2).



## References

- [1] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, 2018. 2, 3, 5, 6
- [2] Timur Bagautdinov, Chenglei Wu, Tomas Simon, Fabian Prada, Takaaki Shiratori, Shih-En Wei, Weipeng Xu, Yaser Sheikh, and Jason Saragih. Driving-signal aware full-body avatars. *ACM TOG*, 2021. 2
- [3] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, Xu Jia, and Huchuan Lu. Animatable neural radiance fields from monocular rgb videos. *arXiv preprint arXiv:2106.13629*, 2021. 3
- [4] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021. 2
- [5] Chen Guo, Xu Chen, Jie Song, and Otmar Hilliges. Human performance capture from monocular video in the wild. In *3DV*, 2021. 3
- [6] Chen Guo, Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Vid2avatar: 3d avatar reconstruction from videos in the wild via self-supervised scene decomposition. In *CVPR*, 2023. 2, 3, 5
- [7] Marc Habermann, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, and Christian Theobalt. Livecap: Real-time human performance capture from monocular video. *ACM TOG*, 2019. 3
- [8] Marc Habermann, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular human performance capture using weak supervision. In *CVPR*, 2020. 3
- [9] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM TOG*, 2021. 2
- [10] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. In *NeurIPS*, 2020. 3
- [11] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. Arch++: Animation-ready clothed human reconstruction revisited. In *ICCV*, 2021. 3
- [12] Xiaoke Huang, Yiji Cheng, Yansong Tang, Xiu Li, Jie Zhou, and Jiwen Lu. Efficient meshy neural fields for animatable human avatars. *arXiv preprint arXiv:2303.12965*, 2023. 2, 3
- [13] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 3
- [14] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. Selfrecon: Self reconstruction your digital avatar from monocular video. In *CVPR*, 2022. 3
- [15] Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. In *CVPR*, 2023. 2, 3, 5, 6, 8
- [16] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *ECCV*, 2022. 2, 3, 5, 6
- [17] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 2, 3, 4
- [18] Christoph Lassner and Michael Zollhofer. Pulsar: Efficient sphere-based neural rendering. In *CVPR*, 2021. 2
- [19] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. Learn to dance with aist++: Music conditioned 3d dance generation. In *ICCV*, 2021. 7
- [20] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. In *ECCV*, 2022. 3
- [21] Zhe Li, Zerong Zheng, Yuxiao Liu, Boyao Zhou, and Yebin Liu. Posevocab: Learning joint-structured pose embeddings for human avatar modeling. In *SIGGRAPH*, 2023. 3
- [22] Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. In *CVPR*, 2024. 3
- [23] Siyou Lin, Hongwen Zhang, Zerong Zheng, Ruizhi Shao, and Yebin Liu. Learning implicit templates for point-based clothed human modeling. In *ECCV*, 2022. 3
- [24] 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 TOG*, 2021. 3
- [25] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhofer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Trans. Graph.*, 2021. 3
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM TOG*, 2015. 3, 4
- [27] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *3DV*, 2024. 4
- [28] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J Black. The power of points for modeling humans in clothing. In *ICCV*, 2021. 2, 3, 4
- [29] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 3
- [30] Gyeongsik Moon, Hyeongjin Nam, Takaaki Shiratori, and Kyoung Mu Lee. 3d clothed human reconstruction in the wild. In *ECCV*, 2022. 3
- [31] 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 *CVPR*, 2019. 4
- [32] Sida Peng, Shangzhan Zhang, Zhen Xu, Chen Geng, Boyi Jiang, Hujun Bao, and Xiaowei Zhou. Animatable neural implicit surfaces for creating avatars from videos. In *ICCV*, 2021. 3
- [33] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 3

- [34] Sergey Prokudin, Qianli Ma, Maxime Raafat, Julien Valentin, and Siyu Tang. Dynamic point fields. In *ICCV*, 2023. 3
- [35] Edoardo Remelli, Timur Bagautdinov, Shunsuke Saito, Chenglei Wu, Tomas Simon, Shih-En Wei, Kaiwen Guo, Zhe Cao, Fabian Prada, Jason Saragih, et al. Drivable volumetric avatars using texel-aligned features. In *SIGGRAPH*, 2022. 3
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 6
- [37] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM TOG*, 2022. 2
- [38] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 3
- [39] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020. 3
- [40] Ruizhi Shao, Jingxiang Sun, Cheng Peng, Zerong Zheng, Boyao Zhou, Hongwen Zhang, and Yebin Liu. Control4d: Efficient 4d portrait editing with text. In *CVPR*, 2024. 3
- [41] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. In *NeurIPS*, 2021. 3
- [42] Shih-Yang Su, Timur Bagautdinov, and Helge Rhodin. Danbo: Disentangled articulated neural body representations via graph neural networks. In *ECCV*, 2022. 2
- [43] Shih-Yang Su, Timur Bagautdinov, and Helge Rhodin. Npc: Neural point characters from video. In *ICCV*, 2023. 3
- [44] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J Black, and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *ICCV*, 2021. 5, 6, 8
- [45] Gusi Te, Xiu Li, Xiao Li, Jinglu Wang, Wei Hu, and Yan Lu. Neural capture of animatable 3d human from monocular video. In *ECCV*, 2022. 2
- [46] Maria Vakalopoulou, Guillaume Chassagnon, Norbert Bus, Rafael Marini, Evangelia I Zacharaki, M-P Revel, and Nikos Paragios. Atlasnet: Multi-atlas non-linear deep networks for medical image segmentation. In *MICCAI*, 2018. 2
- [47] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdf. In *ECCV*, 2022. 3
- [48] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004. 5
- [49] 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 *CVPR*, 2022. 2, 3, 5, 6, 8
- [50] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *CVPR*, 2020. 2
- [51] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. Icon: Implicit clothed humans obtained from normals. In *CVPR*, 2022. 3
- [52] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. Econ: Explicit clothed humans optimized via normal integration. In *CVPR*, 2023. 3
- [53] Feng Xu, Yebin Liu, Carsten Stoll, James Tompkin, Gaurav Bharaj, Qionghai Dai, Hans-Peter Seidel, Jan Kautz, and Christian Theobalt. Video-based characters: creating new human performances from a multi-view video database. *ACM TOG*, 2011. 2
- [54] Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. Monoperfcap: Human performance capture from monocular video. In *SIGGRAPH*, 2018. 3
- [55] 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 *CVPR*, 2024. 3
- [56] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM TOG*, 2019. 2
- [57] Zhengming Yu, Wei Cheng, Xian Liu, Wayne Wu, and Kwan-Yee Lin. Monohuman: Animatable human neural field from monocular video. In *CVPR*, 2023. 2, 3, 8
- [58] Hongwen Zhang, Yating Tian, Xinchu Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *ICCV*, 2021. 8
- [59] Hongwen Zhang, Siyou Lin, Ruizhi Shao, Yuxiang Zhang, Zerong Zheng, Han Huang, Yandong Guo, and Yebin Liu. Closet: Modeling clothed humans on continuous surface with explicit template decomposition. In *CVPR*, 2023. 3
- [60] Hongwen Zhang, Yating Tian, Yuxiang Zhang, Mengcheng Li, Liang An, Zhenan Sun, and Yebin Liu. Pymaf-x: Towards well-aligned full-body model regression from monocular images. *IEEE T-PAMI*, 2023. 8
- [61] 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
- [62] Yuxiang Zhang, Hongwen Zhang, Liangxiao Hu, Jiajun Zhang, Hongwei Yi, Shengping Zhang, and Yebin Liu. Proxycap: Real-time monocular full-body capture in world space via human-centric proxy-to-motion learning. In *CVPR*, 2024. 8
- [63] Hao Zhao, Jinsong Zhang, Yu-Kun Lai, Zerong Zheng, Yingdi Xie, Yebin Liu, and Kun Li. High-fidelity human avatars from a single rgb camera. In *CVPR*, 2022. 2, 3
- [64] Shunyuan Zheng, Boyao Zhou, Ruizhi Shao, Boning Liu, Shengping Zhang, Liqiang Nie, and Yebin Liu. Gps-gaussian: Generalizable pixel-wise 3d gaussian splatting for real-time human novel view synthesis. In *CVPR*, 2024. 3
- [65] Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J Black, and Otmar Hilliges. Pointavatar: Deformable point-based head avatars from videos. In *CVPR*, 2023. 2, 3
- [66] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE T-PAMI*, 2021. 3

- [67] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *CVPR, 2022*. 3
- [68] Zerong Zheng, Xiaochen Zhao, Hongwen Zhang, Boning Liu, and Yebin Liu. Avatarrex: Real-time expressive full-body avatars. *ACM TOG*, 2023. 3
- [69] Boyao Zhou, Jean-Sébastien Franco, Federica Bogo, Bוגra Tekin, and Edmond Boyer. Reconstructing human body mesh from point clouds by adversarial gp network. In *ACCV, 2020*. 3