

# ASH: Animatable Gaussian Splats for Efficient and Photoreal Human Rendering

Haokai Pang<sup>1,2</sup> † Heming Zhu<sup>1</sup> † Adam Kortylewski<sup>1,3</sup> Christian Theobalt<sup>1,4</sup> Marc Habermann<sup>1,4</sup> ✉

<sup>1</sup> Max Planck Institute for Informatics, Saarland Informatics Campus

<sup>2</sup> ETH Zürich <sup>3</sup> Universität Freiburg

<sup>4</sup> Saarbrücken Research Center for Visual Computing, Interaction and AI

{hpang, hezhu, akortyle, theobalt, mhaberma}@mpi-inf.mpg.de




Figure 1. ASH takes an arbitrary 3D skeletal pose and virtual camera view, which can be controlled by the user, as input, and generates a photorealistic rendering of the human in real time. To achieve this, we propose an efficient and animatable Gaussian representation, which is parameterized on the surface of a deformable template mesh.

## Abstract

*Real-time rendering of photorealistic and controllable human avatars stands as a cornerstone in Computer Vision and Graphics. While recent advances in neural implicit rendering have unlocked unprecedented photorealism for digital avatars, real-time performance has mostly been demonstrated for static scenes only. To address this, we propose ASH, an animatable Gaussian splatting approach for photorealistic rendering of dynamic humans in real time. We parameterize the clothed human as animatable 3D Gaussians, which can be efficiently splatted into image space to generate the final rendering. However, naively learning the Gaussian parameters in 3D space poses a severe challenge in terms of compute. Instead, we attach the Gaussians onto a deformable character model, and learn their parameters in 2D texture space, which allows leveraging efficient 2D convolutional architectures that easily scale with the required number of Gaussians. We benchmark ASH with competing methods on pose-controllable avatars, demonstrating that our method outperforms existing real-time methods by a large margin and shows comparable or even better results than offline methods.*

† Joint first authors.

✉ Corresponding author.

Project page: [vcai.mpi-inf.mpg.de/projects/ash](http://vcai.mpi-inf.mpg.de/projects/ash)

## 1. Introduction

Generating high-fidelity human renderings is a long-standing problem in the field of Computer Graphics and Vision, with a multitude of real-world applications, such as gaming, film production, and AR/VR. Typically, this process is a laborious task, requiring complicated hardware setups and tremendous efforts from skilled artists. To ease the extensive manual efforts, recent advances, including this work, focus on generating photorealistic and controllable human avatars solely from multi-view videos.

Recent works on photorealistic human rendering can be categorized into explicit-based and hybrid methods. Explicit methods represent the human avatar as a deformable template mesh with learned dynamic textures [13, 49]. Although these methods are runtime-efficient and can be seamlessly integrated with the well-established rasterization-based rendering pipeline, the generated rendering often falls short in terms of photorealism and level of detail. Hybrid approaches usually attach a neural radiance field (NeRF) [35] onto a (deformable) human model [14, 29, 42]. Typically, they evaluate the NeRF in an unposed space to model the detailed appearance of clothed humans, and generate color and density values by querying a coordinate-based MLP per ray sample. Although hybrid methods can deliver superior rendering quality through NeRF’s capability to capture delicate appearance details, they are unsuit-

able for real-time applications due to the intensive sampling and MLP evaluations required for volume rendering.

Recently, 3D Gaussian splatting [22] with its impressive rendering quality and real-time capability, has become a promising alternative to NeRFs, which are parameterized with a coordinate-based MLP. However, it originally is only designed for modeling static scenes, which is in stark contrast to our problem setting, i.e., modeling dynamic and animatable human avatars. Thus, one may ask: Can the rendering quality and speed of Gaussian splatting be leveraged to model the skeletal motion-dependent characteristics of clothed humans, and how can pose control be achieved?

To answer this, we propose ASH, a real-time approach for generating photorealistic renderings of animatable human avatars. Given a skeletal motion and a virtual camera view, ASH produces photorealistic renderings of clothed humans with motion-dependent details in real time (see Fig. 1). Importantly, during training, ASH only requires multi-view videos for supervision.

In more detail, our animatable human avatar is parameterized using Gaussian splats. However, naively learning a mapping from skeletal pose to Gaussian parameters in 3D is intractable given the limited compute budget when constraining ourselves to real-time performance. Thus, we propose to attach the Gaussians onto a deformable mesh template of the human. Given the mesh’s uv parameterization, it allows learning the Gaussian parameters efficiently in 2D texture space. Here, each texel covered by a triangle represents a Gaussian. Thus, the number of Gaussians remains *constant*, which is in stark contrast to the original formulation. Similarly, we encode the skeletal motion as pose-dependent normal maps. As a result, learning the mapping from skeletal motion to dynamic and controllable Gaussian parameters simplifies to a 2D-to-2D image translation task, which can be efficiently implemented using 2D convolutional architectures. For supervision, we transform the Gaussians into global space using the deformable template and learned Gaussian displacements, splat the Gaussians following the original formulation, and supervise solely on multi-view videos. Our contributions are:

- We propose a novel method, ASH, that enables real-time and high-quality rendering of animatable clothed human avatars solely learned from multi-view video.
- To this end, we represent the human avatar as animatable and dynamic Gaussian splats, which we attach to a deformable template.
- To efficiently learn such a representation, we phrase the problem as a 2D-to-2D texture translation task effectively circumventing 3D architectures, which do not easily scale to the typically required large number of Gaussians.

Our evaluations and comparisons against state-of-the-art methods on animatable human rendering demonstrate that ASH is a significant step towards real-time, high-fidelity,

and controllable human avatars.

## 2. Related Work

**Neural Rendering and Scene Representation.** In the last few years, volumetric representations [15, 50, 51] and neural radiance fields (NeRF) [35, 74] have received significant attention due to their ability to generate high-quality geometry and appearance [57, 70]. However, rendering a NeRF is typically slow as it requires querying an MLP for each ray sample during volume rendering. To address this, subsequent research focused on accelerating the inference process of NeRF: Neural Sparse Voxel Fields [29] adopts an octree to prune the ray samples. DVGO [56] models the scene with an explicit density and feature grid. Plenoxels [10] and PlenOctree [73] replace the MLP with a hierarchical 3D grid storing spherical harmonics, achieving an interactive test-time framerate. TensorRF [6] and Instant-NGP [36] accelerated inference with compact scene representations, i.e., decomposed tensors and neural hash grids. 3D Gaussian Splatting [22] encodes the scene with Gaussian splats storing the density and spherical harmonics, which achieves state-of-the-art rendering quality and shows real-time capability. However, all the above methods are tailored for static scenes, and it is non-trivial to extend them for modeling the dynamic appearances of clothed humans.

There are also notable advancements for extending the concept of NeRFs [27, 39, 40, 44, 60, 67] to dynamic scenes. However, most of these works only support playback of the same dynamic sequence under novel views and, therefore, cannot be adopted for user-controlled pose-dependent dynamic appearance of clothed humans.

**Animatable Neural Human Rendering.** Since this work focuses on animatable human rendering, i.e., at test time, the approach solely takes the skeletal motion as input, we do not discuss works on replay [17, 24, 31, 43, 64, 66], reconstruction [1, 2, 28, 33, 34, 47, 59, 68], and image-based free-viewpoint rendering [45, 62]. Here, according to the underlying shape representation and rendering scheme, we can categorize the literature into two streams, i.e., mesh-based methods and hybrid methods.

Mesh-based methods [3, 5, 13, 49, 61, 69, 71] adopt an explicit, motion-controllable template mesh to model the geometry of clothed humans, with texture space for encoding appearance features. Xu et al. [71] first achieved novel motion and pose synthesis by querying and wrapping texture patches from the captured dataset. Casas et al. [5] and Volino et al. [61] proposed an interactive system that model the appearance as a temporally consistent layered representation in textures space. However, the rendering quality is limited due to the coarse geometry proxy. TNA [49] adopts a texture stack for modeling the dynamic humans’ appearances, though it cannot generate motion-dependent appear-

ances. To address this issue, DDC [13] employs differentiable rendering to learn the non-rigid deformations and dynamic texture maps of clothed humans. At test time, DDC can generalize to novel poses and views and produce real-time photorealistic rendering. Our method outperforms DDC in terms of rendering quality by a large margin while maintaining real-time capability.

Although mesh-based methods provide intuitive control through skeletal poses and integrate seamlessly with established rasterization pipelines, their rendering quality is restrained by the resolution of the template mesh. To this end, hybrid methods are introduced, which articulate the implicit fields with the explicit shape proxy, i.e., parametric human body models [20, 32, 38, 41], or person-specific template meshes. A popular line of research [4, 7, 9, 16, 19, 26, 37, 53, 54, 63, 65] introduced deformable human NeRFs that unwrap the posed space to a shared canonicalized space with inverse kinematics. To better model the pose-dependent appearance of humans, recent studies [11, 14, 25, 30, 42, 72, 76] further introduce motion-aware residual deformations in the canonicalized space. Neural Actor [30] and HDHumans [14] are most closely related to our work within this category. Neural Actor utilizes the texture map of the parametric human body mesh as local pose features to infer dynamic appearances. However, it fails to generalize to characters with loose outfits. HDHumans jointly optimizes the neural implicit fields and the explicit template mesh and, thus, is able to handle loose clothing. However, both methods are slow and take roughly 5 seconds to render a single frame. In stark contrast, our proposed method is capable of real-time rendering with a quality that is on par with, or even superior to HDHumans.

### 3. Method

Our goal is to generate motion-controllable, photorealistic renderings of humans learned solely from multi-view RGB videos (Fig. 2). Specifically, ASH takes the skeletal motions and a virtual camera view as input at inference, and produces high-fidelity renderings in real-time ( $\sim 30fps$ ). To this end, we propose to model the dynamic character with 3D Gaussian splats, parametrized as texels in the texture space of a deformable template mesh. This texel-based parameterization of 3D Gaussian splats enables us to model the mapping from skeletal motions to the Gaussian splat parameters as a 2D image-2-image translation task. Next, we will explain ASH from the following aspects: The background and problem settings (Sec. 3.1), modeling animatable Gaussian splats (Sec. 3.2), and our dedicated training strategy tailored for our animatable Gaussians (Sec. 3.3).

#### 3.1. Problem Setting and Background

We assume a segmented multi-view video  $\mathbf{I}_{f,c}$  of an actor, recorded in a studio with  $C$  synchronized and calibrated

cameras, where  $f$  and  $c$  denote the frame and camera IDs, respectively.  $\mathbf{C}_c$  denotes the camera projection matrix. Additionally, each frame  $\mathbf{I}_{f,c}$  is annotated with the 3D skeletal pose  $\theta_f \in \mathbb{R}^D$  using a markerless motion capture system [58]. Here,  $D$  indicates the number of degrees of freedom (DoFs) for the character’s skeleton. The skeletal motion of the subject  $\bar{\theta}_f \in \mathbb{R}^{k \times D}$  is depicted by a sliding window of skeletal poses from frame  $f-k+1$  to frame  $f$ , where the root translation is normalized w.r.t. the  $f$ th frame.

For training, our model takes the skeletal poses  $\theta_f$  and camera parameters  $\mathbf{C}_c$  as input, renders the animatable Gaussians into image space, and is supervised solely on the multi-view video  $\mathbf{I}_{f,c}$ . During inference, ASH takes arbitrary skeletal poses  $\theta'$  and virtual cameras  $\mathbf{C}'$  as input and generates photorealistic rendering of the subjects at real-time frame rates.

**Gaussian Splatting.** We parameterize the actor representation as 3D Gaussians, which has been proven to be an efficient representation for modeling and rendering static 3D scenes [22]. Here, the static scene is depicted as a collection of 3D Gaussians

$$G(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x})^T \Sigma^{-1}(\mathbf{x})} \quad (1)$$

where  $\Sigma$  denotes the covariance matrix and the Gaussian is centered at  $\mu$ . In Kerbl et al. [22], the Gaussians are parameterized with the set  $\mathcal{G}_i = (\mu_i, \mathbf{q}_i, \mathbf{s}_i, \alpha_i, \eta_i)$ , each defined by its position  $\mu_i \in \mathbb{R}^3$ , rotation quaternion  $\mathbf{q}_i \in \mathbb{R}^4$ , scaling  $\mathbf{s}_i \in \mathbb{R}^3$ , opacity  $\alpha_i \in \mathbb{R}$ , and spherical harmonics coefficients  $\eta_i \in \mathbb{R}^{48}$ . To render the Gaussians into a particular camera view  $c$ , the Gaussian  $i$  has to be projected into image-space by updating the covariance as

$$\Sigma_{i,c} = \mathbf{J}_c \mathbf{C}_c \mathbf{R}_i \mathbf{S}_i \mathbf{S}_i^T \mathbf{R}_i^T \mathbf{C}_c^T \mathbf{J}_c^T \quad (2)$$

where  $\mathbf{R}_i$  and  $\mathbf{S}_i$  are rotation and scaling matrices obtained from the quaternion conversion  $\mathbf{q}_i$  and the scaling coefficients  $\mathbf{s}_i$ .  $\mathbf{J}_c$  is the Jacobian of the affine approximation of the projective transformation  $\mathbf{C}_c$ .

To render the color  $\mathbf{c}_p$  of a pixel  $\mathbf{p}$  in camera  $c$ , 3D Gaussian splatting [22] adopts a point-based splatting formulation, which blends the spherical harmonics  $\eta_i$  of the depth-ordered Gaussian splats overlapping with the pixel as

$$\mathbf{c}_p = \sum_{j \in \mathcal{N}_p} H(\eta_i, \mathbf{d}_p) \alpha'_j \prod_{k=1}^{j-1} (1 - \alpha'_k), \quad (3)$$

where  $\mathcal{N}_p$  denotes the set of Gaussian splats covering pixel  $\mathbf{p}$ .  $\alpha'_j$  refers to the opacity for the  $j$ th ordered Gaussian splat with respect to the current pixel, i.e.  $\alpha'_j = \alpha_j G_j(\mathbf{p})$ .  $H(\cdot)$  indicates the function that converts the spherical harmonics coefficients  $\eta_i$  and the view direction  $\mathbf{d}_p$  to an RGB color.

While 3D Gaussian splatting can produce high-quality renderings at very high frame rates (more than 100 fps), its

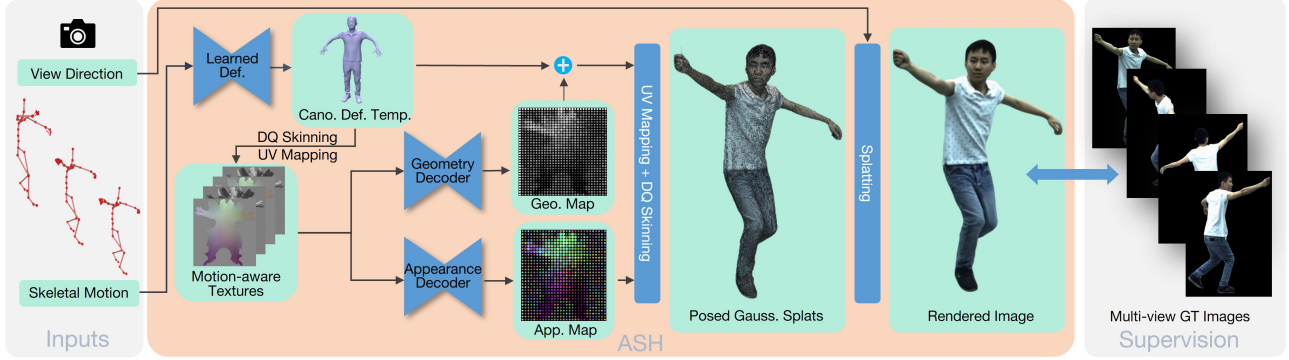


Figure 2. ASH generates high-fidelity rendering given a skeletal motion and a virtual camera view. A motion-dependent, canonicalized template mesh is generated with a learned deformation network. From the canonical template mesh, we can render the motion-aware textures, which are further adopted for predicting the Gaussian splat parameters with two 2D convolutional networks, i.e., the Geometry and Appearance Decoder, as the texels in the 2D texture space. Through UV mapping and DQ skinning, we warp the Gaussian splats from the canonical space to the posed space. Then, splating is adopted to render the posed Gaussian splats.

usage is primarily demonstrated for static scenes, and it is non-trivial to adopt this concept for controllable, detailed, and dynamic 3D human avatars. What is required here is *animatable* 3D Gaussians, i.e. we want to model the set of Gaussian parameters  $\{\mathcal{G}_i\}_{N_g}$  as a function of skeletal motion  $\bar{\theta}_f$  where  $N_g$  denotes the total number of Gaussians. Note that we consider motion rather than pose to account for potential surface dynamics.

### 3.2. Animatable Gaussian Splats

Intuitively, we want to learn a function  $\mathcal{F}(\bar{\theta}_f) = \{\mathcal{G}_i\}_{N_g}$  that maps the skeletal motion to animatable 3D Gaussian parameters. However, more than 20,000 Gaussian splats are typically required to achieve high-fidelity rendering of clothed humans. Thus, modeling and learning such a function can be challenging, especially when modeling it in 3D. Instead, our idea is to attach the Gaussian splats onto an animatable template mesh of the human, and parameterize the Gaussian splats in 2D texture space, i.e., each texel of the template mesh (covered by a face) stores the parameters of a 3D Gaussian. This enables ASH to efficiently learn the Gaussian parameters in 2D texture space, which we will now describe in more detail.

**Animatable Template.** To achieve this, we require an animatable human template denoted as  $M(\theta_f) = \mathbf{V}_f$ , which takes the skeletal motion and computes posed and deformed 3D vertices  $\mathbf{V}_f$  of a person-specific template mesh  $\mathbf{V}_m$ . In practice, we leverage the character model of Habermann et al. [13] and refer to the supplemental for further details. To generate the animatable template mesh  $\mathbf{V}_f$ , we first non-rigidly deform the template mesh vertices  $\mathbf{V}_m$  in the unposed-canonical space, denoted as  $\bar{\mathbf{V}}_f$ , with skeletal motion-dependent, i.e.  $\bar{\theta}_f$ , and learned embedded deformations [55] and per-vertex displacements. Given the skeletal pose  $\theta_f$ , the canonically deformed template mesh

vertices  $\bar{\mathbf{V}}_f$  can then be posed using Dual Quaternion skinning [21], denoted as  $\mathbf{V}_f$ .

**Animatable Gaussian Textures.** ASH depicts the character’s appearance with a *fixed* number of animatable Gaussian splats  $\{\mathcal{G}_i\}_N = (\bar{\boldsymbol{\mu}}_{uv,i}, \bar{\mathbf{d}}_{uv,i}, \mathbf{q}_{uv,i}, \mathbf{s}_{uv,i}, \alpha_{uv,i}, \boldsymbol{\eta}_{uv,i}) \in \mathbb{R}^{N \times 62}$  as the texels on the texture space of the animatable template mesh  $M(\theta_f)$ . Here,  $N$  denotes the number of texels that are covered by triangles in the UV map. Specifically,  $\bar{\boldsymbol{\mu}}_{uv,i}$  denotes the base position for Gaussian splats in the canonical space, which can be derived from the canonical animatable template mesh vertices  $\bar{\mathbf{V}}_f$  through texture mapping:

$$\bar{\boldsymbol{\mu}}_{uv,i} = w_{a,i} \bar{\mathbf{V}}_{f,j} + w_{b,i} \bar{\mathbf{V}}_{f,k} + w_{c,i} \bar{\mathbf{V}}_{f,l}, \quad (4)$$

where  $w_{(\cdot),i}$  denotes the barycentric weights for the texels and  $\bar{\mathbf{V}}_{f,(\cdot)}$  stands for the canonical vertex position for the triangle that covers the texel. Similar to the animatable template, we can pose the Gaussian splats  $\{\mathcal{G}_i\}$  stored in texels, from the canonical position  $\bar{\boldsymbol{\mu}}_i$  to the posed space, through Dual Quaternion skinning [21]:

$$\boldsymbol{\mu}_{uv,i} = \mathbf{T}_{uv,i}(\bar{\boldsymbol{\mu}}_{uv,i} + \bar{\mathbf{d}}_{uv,i}), \quad (5)$$

where  $\mathbf{T}_{uv,i}$  denotes the Dual Quaternion skinning transformation matrix for the  $i$ th texel.  $\bar{\mathbf{d}}_{uv,i}$  refers to a learned per-texel offset in the canonical space, which captures fine motion-dependent deformations of the Gaussian splats.

Parameterizing Gaussian splats as 2D texels enables us to predict them using efficient 2D convolutional architectures. Moreover, the shared canonical 2D space facilitates the learning of the motion-dependent Gaussian parameters.

**Gaussian Texture Decoder.** Due to the texel-based 2D parameterization of the 3D Gaussian splats, we can leverage the well-established, efficient 2D convolutional architectures. To formulate the mapping between the 3D skeletal

motion  $\bar{\theta}_f$  and the dynamic Gaussian splats  $\{\mathcal{G}_i\}_N$  on 2D texture space as a image-2-image translation problem [18], we adopt the motion-aware textures  $(\mathbf{T}_{n,f}, \mathbf{T}_{p,f})$  to depict the 3D skeletal motions  $\bar{\theta}_f$  in the 2D texture space. The normal textures  $\mathbf{T}_{n,f}$  and position textures  $\mathbf{T}_{p,f}$  can be computed from the posed deformed template mesh  $\mathbf{V}_f$  vertices of previous frames through inverse texture mapping. Consequently, we propose motion-aware 2D convolutional neural networks, i.e., the geometry network  $\mathcal{E}_{\text{geo}}$ , and the appearance network  $\mathcal{E}_{\text{app}}$ , predicting the geometry and appearance parameters of the Gaussian splats from the motion-aware textures  $(\mathbf{T}_{n,f}, \mathbf{T}_{p,f})$ . The geometry network  $\mathcal{E}_{\text{geo}}$  predicts the shape-related parameters, namely, the canonical offset  $\bar{\mathbf{d}}_{uv,i}$ , scale  $\mathbf{s}_{uv,i}$ , rotation quaternions  $\mathbf{q}_{uv,i}$ , and opacity  $\alpha_{uv,i}$ :

$$\mathcal{E}_{\text{geo}}(\mathbf{T}_{n,f}, \mathbf{T}_{p,f}) = (\bar{\mathbf{d}}_{uv,i}, \mathbf{s}_{uv,i}, \mathbf{q}_{uv,i}, \alpha_{uv,i}). \quad (6)$$

A separated motion-aware convolution decoder  $\mathcal{E}_{\text{app}}$  is adopted for learning the appearances characterized by the Spherical Harmonics  $\boldsymbol{\eta}_{uv,i}$ :

$$\mathcal{E}_{\text{app}}(\mathbf{T}_{n,f}, \mathbf{T}_{p,f}, \Phi_f) = \boldsymbol{\eta}_{uv,i}, \quad (7)$$

where  $\Phi_f$  indicates the global appearance features, which encodes the global root transition of the character with a shallow MLP, to account for the spatially varying lighting conditions within the capture space.

### 3.3. Training Strategy

Unlike static scenes, dynamic clothed humans exhibit motion-dependent appearances and varying geometry throughout the frames, posing a significant challenge in training. To make it tractable, we propose a carefully designed training paradigm, which decomposes the learning of the motion-aware convolutions into two stages, namely, the warmup stage, and the final training.

**Warmup Stage.** As mentioned in Sec. 4, the DynaCap dataset [13] and our proposed dataset feature long training sequences with various motion-dependent detailed appearances. Therefore, naively training the proposed motion-aware decoders  $\mathcal{E}_{\text{geo}}$  and  $\mathcal{E}_{\text{app}}$ , from scratch without proper initialization will not converge during training. To tackle this problem, we propose a warmup stage, providing a better weight initialization for the motion-aware decoders.

We first sample  $t$  frames evenly across the training sequence and learn 3D Gaussian splat parameters  $\{\mathcal{G}_i''\}_{N_g}$  separately, which serves as a pseudo ground truth for the Gaussian splat parameters. In contrast to the original implementation for static 3D Gaussian splatting [22], we fix the position of the Gaussian splats  $\boldsymbol{\mu}_{uv,i}''$  throughout the training while only optimizing the remaining parameters. Specifically, the initial value for the Gaussian splat positions  $\boldsymbol{\mu}_{uv,i}''$  can be read out from the texture texels of the pose-deformed

template mesh  $\boldsymbol{\mu}_{uv,i}$ . Additionally, to preserve the correspondences across pseudo ground truth frames, we remove the splitting/merging of the Gaussian splats and keep the number of Gaussian splats fixed. The pretraining optimizes the L2 loss between the pseudo ground truth  $\{\mathcal{G}_i''\}$  and the Gaussian splat parameters produced by the motion-aware decoders  $\{\mathcal{G}_i'\}$ :

$$\mathcal{L}_{\text{pre}} = \mathcal{L}_2(\{\mathcal{G}_i'\}, \{\mathcal{G}_i''\}). \quad (8)$$

**Final Training.** After the warmup stage, we can further train the motion-aware decoder on the whole training sequence by minimizing the pixel-wise L1 and structural-similarity-index loss between the generated images  $\mathbf{I}'_{f,c}$  and the multi-view ground truth images  $\mathbf{I}_{f,c}$ :

$$\mathcal{L}_{\text{main}} = \lambda_{\text{pix}} \mathcal{L}_1(\mathbf{I}_{f,c}, \mathbf{I}'_{f,c}) + \lambda_{\text{str}} \mathcal{L}_{\text{ssim}}(\mathbf{I}_{f,c}, \mathbf{I}'_{f,c}), \quad (9)$$

where  $\mathcal{L}_{\text{ssim}}$  denotes the structural similarity index loss measuring the structural difference between two images.  $\lambda_{\text{pix}}$  and  $\lambda_{\text{ssim}}$  are set to 0.1 and 0.9, respectively.

## 4. Results

**Dataset** We adopted the DynaCap dataset [13] to quantitatively and qualitatively assess the effectiveness of our approach. We selected two representative subjects from the DynaCap dataset wearing loose and tight types of apparel for evaluating the accuracy of novel-view rendering and generalization ability to novel poses. Following the protocol proposed in DDC [13], we train our model using the training splits from the DynaCap dataset. Here, we hold out 4 camera views to assess the novel-view rendering accuracy. Moreover, we evaluate the model’s generalization ability to novel poses with motion sequences from the testing splits.

In addition to the DynaCap dataset, we recorded two novel sequences featuring distinct subjects to showcase the performance of our model qualitatively. The recorded subjects perform everyday motions such as dancing, jogging, and jumping. The sequences are recorded using a calibrated multi-camera system with 120 cameras at a frame rate of 25 fps. Separate training and testing sequences are recorded with a duration of 27,000 frames and 7,000 frames, respectively. All the captured frames are annotated with 3D skeletal poses [58] and foreground segmentations [23, 48].

### 4.1. Qualitative Results

We evaluate the performance of ASH on subjects from the DynaCap [13] dataset and our newly recorded sequences.

**Novel View Synthesis.** Fig. 3 presents the novel view synthesis results rendered from camera views unseen during training. ASH yields photorealistic rendering in real time, capturing sharp wrinkles details and view-dependent appearances. Remarkably, it can even generalize to loose




Figure 3. **Qualitative Results.** We present the results generated with ASH regarding novel view and pose synthesis. Note that our methods can produce high-quality rendering with delicate, motion-aware details for novel views and skeletal motions.

types of apparel and faithfully recovers the clothing dynamics, e.g., the swing of the skirts.

**Novel Pose Synthesis.** We further show the results generated on novel poses extracted from the testing sequences in Fig. 3. Given poses that significantly deviate from the training poses, our method still generates high-quality renderings with motion-aware appearances. For the dynamic results, we refer to the supplemental video.

## 4.2. Comparisons

**Competing Methods.** We compare our model with the state of the arts on animatable neural human rendering: 1) DDC [13] features a mesh-based approach where the geometry is represented with a learned embedded graph, and the appearance is encoded using learned dynamic textures. Specifically, DDC is the only real-time approach among competing methods, while other hybrid methods typically take seconds to volume-render an image. 2) TAVA [26] is a hybrid approach depicting the shape, appearance, and skinning weights as implicit fields in canonical space. The samples in the posed space are canonicalized w.r.t. the skeleton through iterative root finding. 3) NA [29] conditions

the canonical color and density field of dynamic characters on the learned feature texture of the parametric human body models. The canonicalization of spatial samples is achieved by inverse kinematics. 4) HDHumans [14] models the appearance of dynamic humans as the appearance and density fields conditioned on the feature texture map of the motion-aware deformable template mesh. Notably, the template mesh will deform w.r.t. the implicit density field, improving the alignment between the observation and canonical space. **Metrics.** We adopt the Peak Signal-to-Noise Ratio (PSNR) metric to measure the quality of the rendered image. Moreover, we adopt the learned perceptual image patch similarity (LPIPS) [75] that better mirrors human perception. Note that the metrics are assessed at a 1K resolution, averaged across every 10th frame throughout the sequence. Here, we denote the subject with tight outfits as *Tight Outfits*, and the other wearing loose clothing as *Loose Outfits*.

**Quantitative Comparison.** Tab. 1 and Tab. 2 illustrate the quantitative comparison against the competing methods on novel view and pose synthesis. Compared with the real-time capable methods, our method significantly outperforms DDC [13] in PSNR and LPIPS regarding novel-view

Methods	RT	Tight Outfits		Loose Outfits	
		PSNR	LPIPS	PSNR	LPIPS
TAVA	✗	24.61	62.26	27.31	37.55
NA	✗	30.33	23.71	25.30	50.01
HDHumans	✗	30.98	15.09	29.24	15.79
DDC	✓	31.21	22.56	28.10	31.68
<b>Ours</b>	✓	35.84	11.92	35.47	8.30

Table 1. **Quantitative Comparison on Novel View Synthesis.** We quantitatively compare ASH with other methods on seen skeletal motions but unseen views. We highlight the **best** and **second-best** scores. We outperform previous real-time and even non-real-time methods in all matrices by a large margin.

Methods	RT	Tight Outfits		Loose Outfits	
		PSNR	LPIPS	PSNR	LPIPS
TAVA	✗	28.30	37.47	26.31	50.11
NA	✗	28.78	25.78	25.03	44.20
HDHumans	✗	28.17	20.69	26.71	22.75
DDC	✓	27.77	30.16	26.43	32.22
<b>Ours</b>	✓	28.90	22.83	27.12	20.22

Table 2. **Quantitative Comparison on Novel Pose Synthesis.** We quantitatively compare ASH with other methods on unseen skeletal motions and unseen views. ASH achieves the highest PSNR and the second-best LPIPS on the subject with tight outfits, and outperforms other methods for the subject with loose clothing.

synthesis, highlighting our method’s superiority in capturing the motion-aware appearances from the training data. In novel pose synthesis, compared to DDC [13], our method demonstrates significantly improved performance. This underscores our method’s generalization ability to novel motions. As for the comparison against the non-real-time approaches, our method consistently surpasses previous works regarding PSNR and LPIPS. Notably, our method is capable of real-time rendering and achieves much better quantitative performances than HDHumans in novel-view synthesis, and comparable performances in novel-pose synthesis.

**Qualitative Comparison.** Fig. 4 comprises the qualitative comparison on the novel-view and novel-pose rendering: TAVA [26] struggles to handle various motions in the DynaCap dataset [13], resulting in blurry rendering. While NA [30] effectively captures details in subjects wearing tight apparel, it struggles with significant artifacts for subjects in loose outfits. This issue arises from the inherent challenge of representing loose clothing as residual displacements on the parametric human body model. HDHumans [14] stands out among non-real-time competing methods, producing high-fidelity renderings with sharp details. Notably, it preserves the wrinkles and delicate textures on the skirt. However, due to the extensive sampling needed for volume rendering, it takes seconds for HDHumans to render a single frame. In contrast, ASH excels by deliver-

Methods	Training Pose		Testing Pose	
	PSNR	LPIPS	PSNR	LPIPS
w/o mot.	27.19	33.16	26.86	32.35
w/o disp.	33.21	17.34	27.33	25.32
w/ 128.res.	35.15	11.26	27.13	22.00
w/ 512.res.	35.28	8.60	27.00	21.13
<b>Ours</b>	35.47	8.30	27.13	20.22

Table 3. **Ablation Study.** We assess our design choices on the image synthesis tasks on the subject with loose outfits. Our method achieves better performance against the design alternatives.

ing rendering quality that matches or exceeds HDHumans in real time.

DDC [13] is the only competing method with real-time capability. Although it captures coarse motion-aware appearances, its output tends to be blurry and lacks detail. ASH matches the real-time capability as DDC, while generating renderings with much finer details.

### 4.3. Ablations

To assess the effectiveness of the major components of our model, we conduct the following ablative experiments on the novel view and motion synthesis tasks.

**Motion Conditions.** Our method depicts the appearance of the clothed human through motion-aware, deformable Gaussian splats in the canonical space. To assess the efficacy of the motion conditions, we remove the motion-aware decoder and learn the appearance parameters of Gaussian splats from a truncated training sequence of 1,000 frames, termed as **w/o mot.**. As seen in Tab. 3 and Fig. 5, without motion conditioning, the synthesized results fail to recover the clothing dynamics and suffer from severe artifacts.

**Motion-aware Offset.** The motion-aware offset is adopted to account for the non-rigid motion-dependent deformation of the Gaussian splats. We remove the motion-aware offset applied to the canonical Gaussian splats, only allowing the appearance to be motion-dependent, termed as **w/o disp.**. As shown in Tab. 3, excluding the learned motion-aware offset leads to worse quantitative performance and noticeable blurry artifacts on the rendered images.

**Texture Resolution.** The animatable Gaussian splats are parameterized as texels in the texture space of the deformable template mesh, where the resolution is set to 256. To study the impact on the resolution of the texture space, we conducted ablative studies with different resolutions, i.e., halved resolution termed as **w/ 128.res.**, and doubled resolution termed as **w/ 512.res.**. As illustrated in Tab. 3 and Fig. 5, doubling the resolution results in comparable results, while it significantly increases computational complexity in both the U-Net [46] evaluation and tile-based rasterization, preventing the model from being real-time compatible. On the other hand, reducing the resolution to 128 leads to a sig-



Figure 4. **Qualitative Comparison.** We compared our methods with the state of the arts, i.e., TAVA [26], NA [30], HDHumans [14], DDC [13], in novel view and novel motion synthesis. Note that our results significantly outperform the real-time methods in quality while showing comparable or even better results than the offline methods.

nificant decline in perceptual metrics and blurry rendering.

As seen in Tab. 3 and Fig. 5, our method outperforms the design alternatives quantitatively and qualitatively.

## 5. Conclusion

In this paper, we introduce ASH, a real-time method for high-quality rendering of animated humans, learned solely from multi-view videos. ASH attaches the 3D Gaussians splats, initially designed for modeling static scenes, onto a

deformable mesh template. Bridged by the mesh’s UV parameterization, we can efficiently learn the 3D Gaussians in 2D texture space as a texture translation task. ASH demonstrates significantly better performances quantitatively and qualitatively than state-of-the-art, real-time capable methods on animatable human rendering, and even better performance than the state-of-the-art offline methods. Currently, ASH does not update the underlying deformable template mesh. In the future, we will explore whether the Gaussian splatting can directly improve the 3D mesh geometry.






Figure 5. **Qualitative Ablation.** We compare our design choices on the image synthesis task. Our method excels in rendering quality and detail recovery. Our method shows compatible rendering quality as **w/ 512.res.** with doubled texel resolution and much sharper rendering than **w/ 128.res.** with halved texel resolution.

## References

- [1] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *International Conference on 3D Vision*, pages 98–109, 2018. 2
- [2] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1175–1186, 2019. 2
- [3] 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 Transactions on Graphics (TOG)*, 40(4):1–17, 2021. 2
- [4] Alexander Bergman, Petr Kellnhofer, Wang Yifan, Eric Chan, David Lindell, and Gordon Wetzstein. Generative neural articulated radiance fields. *Adv. Neural Inform. Process. Syst.*, 35:19900–19916, 2022. 3
- [5] Dan Casas, Marco Volino, John Collomosse, and Adrian Hilton. 4d video textures for interactive character appearance. *Comput. Graph. Forum*, 33(2):371–380, 2014. 2
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *Eur. Conf. Comput. Vis.*, pages 333–350. Springer, 2022. 2
- [7] 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
- [8] Paolo Cignoni, Guido Ranzuglia, M Callieri, M Corsini, F Ganovelli, N Pietroni, M Tarini, et al. Meshlab. 2011. 12
- [9] Yao Feng, Jinlong Yang, Marc Pollefeys, Michael J. Black, and Timo Bolkart. Capturing and animation of body and clothing from monocular video. In *SIGGRAPH Asia 2022 Conference Papers*, 2022. 3
- [10] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5501–5510, 2022. 2
- [11] Qingzhe Gao, Yiming Wang, Libin Liu, Lingjie Liu, Christian Theobalt, and Baoquan Chen. Neural novel actor: Learning a generalized animatable neural representation for human actors. *IEEE Trans. Vis. Comput. Graph.*, 2023. 3
- [12] Michael Garland and Paul S Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 209–216, 1997. 12
- [13] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM Trans. Graph.*, 40(4), 2021. 1, 2, 3, 4, 5, 6, 7, 8, 12
- [14] Marc Habermann, Lingjie Liu, Weipeng Xu, Gerard Pons-Moll, Michael Zollhoefer, and Christian Theobalt. Hdhumans: A hybrid approach for high-fidelity digital humans. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(3):1–23, 2023. 1, 3, 6, 7, 8
- [15] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Escaping plato’s cave: 3d shape from adversarial rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9984–9993, 2019. 2
- [16] Shoukang Hu, Fangzhou Hong, Liang Pan, Haiyi Mei, Lei Yang, and Ziwei Liu. Sherf: Generalizable human nerf from a single image. *arXiv preprint*, 2023. 3
- [17] Mustafa Işık, Martin Runz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Niessner. Humanrf: High-fidelity neural radiance fields for humans in motion. *ACM Transactions on Graphics (TOG)*, 42(4):1–12, 2023. 2
- [18] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. 5
- [19] T. Jiang, X. Chen, J. Song, and O. Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16922–16932, 2023. 3
- [20] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8320–8329, 2018. 3
- [21] Ladislav Kavan, Steven Collins, Jiří Žára, and Carol O’Sullivan. Skinning with dual quaternions. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games*, pages 39–46, 2007. 4, 12
- [22] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):1–14, 2023. 2, 3, 5, 13
- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer White-

- head, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 5
- [24] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *Adv. Neural Inform. Process. Syst.*, 2021. 2
- [25] Youngjoong Kwon, Lingjie Liu, Henry Fuchs, Marc Habermann, and Christian Theobalt. Deliffas: Deformable light fields for fast avatar synthesis. *Adv. Neural Inform. Process. Syst.*, 2023. 3
- [26] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. 2022. 3, 6, 7, 8
- [27] Zhengqi Li, Simon Niklaus, Noah Snaveley, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021. 2
- [28] Siyou Lin, Hongwen Zhang, Zerong Zheng, Ruizhi Shao, and Yebin Liu. Learning implicit templates for point-based clothed human modeling. In *ECCV (3)*, pages 210–228, 2022. 2
- [29] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Adv. Neural Inform. Process. Syst.*, 33:15651–15663, 2020. 1, 2, 6
- [30] 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.(ACM SIGGRAPH Asia)*, 2021. 3, 7, 8
- [31] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhofer, Yaser Sheikh, and Jason M. Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Trans. Graph.*, 40(4):59:1–59:13, 2021. 2
- [32] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. 3
- [33] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J. Black. The power of points for modeling humans in clothing. In *Int. Conf. Comput. Vis.*, pages 10974–10984, 2021. 2
- [34] Marko Mihajlovic, Yan Zhang, Michael J. Black, and Siyu Tang. Leap: Learning articulated occupancy of people. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10461–10471, 2021. 2
- [35] 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 *Eur. Conf. Comput. Vis.*, 2020. 1, 2
- [36] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):1–15, 2022. 2
- [37] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *Int. Conf. Comput. Vis.*, 2021. 3
- [38] Ahmed A. A. Osman, Timo Bolkart, and Michael J. Black. Star: Sparse trained articulated human body regressor. In *Eur. Conf. Comput. Vis.*, pages 598–613, 2020. 3
- [39] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *Int. Conf. Comput. Vis.*, 2021. 2
- [40] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *ACM Trans. Graph.*, 40(6), 2021. 2
- [41] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10975–10985, 2019. 3
- [42] 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 *Int. Conf. Comput. Vis.*, pages 14314–14323, 2021. 1, 3
- [43] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9054–9063, 2021. 2
- [44] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2
- [45] Edoardo Remelli, Timur M. Bagautdinov, Shunsuke Saito, Chenglei Wu, Tomas Simon, Shih-En Wei, Kaiwen Guo, Zhe Cao, Fabian Prada, Jason M. Saragih, and Yaser Sheikh. Drivable volumetric avatars using texel-aligned features. In *SIGGRAPH (Conference Paper Track)*, pages 56:1–56:9, 2022. 2
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, Oct 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 7, 12
- [47] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2886–2897, 2021. 2
- [48] Soumyadip Sengupta, Vivek Jayaram, Brian Curless, Steve Seitz, and Ira Kemelmacher-Shlizerman. Background matting: The world is your green screen. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 5
- [49] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliiev, Renat Bashirov, Egor Burkov, Karim Isakov, Aleksei Ivakhnenko, Yury Malkov, Igor Pasechnik, Dmitry Ulyanov, et al. Textured neural avatars. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2387–2397, 2019. 1, 2
- [50] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deep-

- voxels: Learning persistent 3d feature embeddings. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 2
- [51] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Adv. Neural Inform. Process. Syst.*, 2019. 2
- [52] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*. Eurographics Association, 2007. 12
- [53] Shih-Yang Su, Frank Yu, Michael Zollhofer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. *Adv. Neural Inform. Process. Syst.*, 34:12278–12291, 2021. 3
- [54] Shih-Yang Su, Timur Bagautdinov, and Helge Rhodin. Danbo: Disentangled articulated neural body representations via graph neural networks. In *Eur. Conf. Comput. Vis.*, 2022. 3
- [55] Robert W. Sumner, Johannes Schmid, and Mark Pauly. Embedded deformation for shape manipulation. *ACM Trans. Graph.*, 26(3):80–es, 2007. 4, 12
- [56] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5449–5459, 2022. 2
- [57] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Wang Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. In *Comput. Graph. Forum*, pages 703–735. Wiley Online Library, 2022. 2
- [58] TheCapture. The Capture. <http://www.thecapture.com/>, 2020. 3, 5
- [59] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-gif: Neural generalized implicit functions for animating people in clothing. In *Int. Conf. Comput. Vis.*, pages 11688–11698, 2021. 2
- [60] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhofer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Int. Conf. Comput. Vis. IEEE*, 2021. 2
- [61] Marco Volino, Dan Casas, John Collomosse, and Adrian Hilton. Optimal representation of multiple view video. In *Brit. Mach. Vis. Conf.* BMVA Press, 2014. 2
- [62] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 2
- [63] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdfs. In *Eur. Conf. Comput. Vis.*, 2022. 3
- [64] Ziyang Wang, Timur Bagautdinov, Stephen Lombardi, Tomas Simon, Jason Saragih, Jessica Hodgins, and Michael Zollhofer. Learning compositional radiance fields of dynamic human heads, 2020. 2
- [65] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Vid2actor: Free-viewpoint animatable person synthesis from video in the wild. *arXiv preprint arXiv:2012.12884*, 2020. 3
- [66] 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 Conf. Comput. Vis. Pattern Recog.*, pages 16210–16220, 2022. 2
- [67] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9421–9431, 2021. 2
- [68] Donglai Xiang, Fabian Prada, Timur Bagautdinov, Weipeng Xu, Yuan Dong, He Wen, Jessica Hodgins, and Chenglei Wu. Modeling clothing as a separate layer for an animatable human avatar. *ACM Trans. Graph.*, 40(6):1–15, 2021. 2
- [69] Donglai Xiang, Timur Bagautdinov, Tuur Stuyck, Fabian Prada, Javier Romero, Weipeng Xu, Shunsuke Saito, Jingfan Guo, Breannan Smith, Takaaki Shiratori, et al. Dressing avatars: Deep photorealistic appearance for physically simulated clothing. *ACM Trans. Graph.*, 41(6):1–15, 2022. 2
- [70] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. In *Comput. Graph. Forum*, pages 641–676. Wiley Online Library, 2022. 2
- [71] 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. In *ACM SIGGRAPH 2011 papers*, pages 1–10. 2011. 2
- [72] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *Adv. Neural Inform. Process. Syst.*, 34:14955–14966, 2021. 3
- [73] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Int. Conf. Comput. Vis.*, pages 5752–5761, 2021. 2
- [74] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 2
- [75] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 6
- [76] Zerong Zheng, Xiaochen Zhao, Hongwen Zhang, Boning Liu, and Yebin Liu. Avatarrex: Real-time expressive full-body avatars. *ACM Trans. Graph.*, 42(4), 2023. 3

## A. Overview

In this appendix, we provide more details regarding the following aspects: More implementation details (Sec. B); more ablative studies (Sec. C); runtime analysis for the major components (Sec. D); realtime applications built upon ASH (Sec. E); more detailed discussion on limitations and future directions (Sec. F).

## B. Implementation Details

In the main paper, we mentioned that ASH learns the Gaussian splat parameters in the 2D texture space of an animatable human template  $M(\theta_f) = \mathbf{V}_f$ . Here, we provide more details regarding the **deformable template mesh** and the **motion-aware decoders**.

**Deformable Template Mesh.** We adopt the formulation mentioned in Habermann et al. [13] for modeling the deformable template mesh, which deforms the template mesh vertices  $\mathbf{V}_m$  in the canonical space with a learned embedded deformation [52, 55]:

$$\bar{\mathbf{V}}_{f,i} = \mathbf{D}_i + \sum_{j \in \mathcal{N}_{nv,i}} w_{i,j} (R(\mathbf{A}_j)(\mathbf{V}_{m,i} - \mathbf{V}_{G,j}) + \mathbf{V}_{G,j} + \mathbf{T}_j) \quad (10)$$

where  $\bar{\mathbf{V}}_{f,i} \in \mathbb{R}^3$  denotes the deformed template vertices in the rest pose.  $\mathcal{N}_{nv,i} \in \mathbb{N}$  indicates the indexes for the embedded graph node that are connected to the  $i$ th vertex on the template mesh.  $\mathbf{V}_{G,j} \in \mathbb{R}^3$ ,  $\mathbf{A}_j \in \mathbb{R}^3$ , and  $\mathbf{T}_j \in \mathbb{R}^3$  stands for the rest positions, rotation Euler angles, and translations of the embedded graph nodes. Notably, the connectivity of the embedded graph  $\mathbf{V}_{G,j}$  can be adopted by simplifying the template mesh  $M$  using quadric edge collapse decimation [8, 12]. Moreover, the connection, as well as the connection weights  $w_{i,j}$ , between the template mesh  $\mathbf{V}_m$  and the embedded graph are generated following the configuration in Sumner et al. [55].  $R(\cdot) \in \mathbb{R}^{3 \times 3}$  denotes the function that converts the Euler angle to a rotation matrix.  $\mathbf{D}_i \in \mathbb{R}^3$  indicates the per-vertex displacement to model an even finer level of geometry details. Specifically, embedded graph parameters  $\mathbf{V}_{G,j}$ ,  $\mathbf{A}_j$ , and per-vertex displacements  $\mathbf{D}_i$  are derived from skeletal motion  $\bar{\theta}_f$  with structure-aware graph convolution neural networks. We refer to Habermann et al. [13] for more details.

**Motion-aware Decoders.** ASH adopts motion-aware 2D convolutional neural networks, i.e., the geometry network  $\mathcal{E}_{geo}$ , and the appearance network  $\mathcal{E}_{app}$ , predicting the geometry and appearance parameters of the Gaussian splats from the motion-aware textures  $(\mathbf{T}_{n,f}, \mathbf{T}_{p,f})$ . Both the geometry network  $\mathcal{E}_{geo}$  and the appearance network  $\mathcal{E}_{app}$  are U-Nets implemented following the configuration mentioned in Olaf et al. [46]. Specifically, we channel-wise concatenate the global appearance features  $\Phi_f \in \mathbb{R}^{16}$  to the bottleneck features of the appearance network  $\mathcal{E}_{app}$  to account for the lighting variations in the studio. The global appearance

Methods	Training Pose		Testing Pose	
	PSNR	LPIPS	PSNR	LPIPS
w/ MLP	28.05	34.93	26.79	35.00
w/ 12.cam	33.50	10.76	26.96	21.66
w/ 30.cam	35.37	8.52	27.08	20.76
w/ 60.cam	35.49	8.32	27.14	20.42
<b>Ours</b>	35.47	8.30	27.13	20.22

Table 1. **Ablation Study.** We further assess our design choices on the image synthesis tasks with the subject wearing loose outfits in the DynaCap [13] dataset. We highlight the **best** and the **second-best** scores.

feature  $\Phi_f$  is derived from positional encoded skeleton root translation with a 3-layer shallow MLP, of which the width is set to 32.

## C. Ablations

In this section, we provide more ablative studies to demonstrate the effectiveness of ASH.

**Number of Camera Views.** To assess the robustness of ASH against sparser camera view supervision, we conducted ablative experiments that take multi-view videos from 12, 30, and 60 cameras as supervision, termed as **w/ 12.cam**, **w/ 30.cam**, and **w/ 60.cam**. Note that the selected camera views are evenly distributed in the studio. As illustrated in Fig. 1 and Tab. 1, ASH can still accurately synthesize the imagery of the animatable characters when training with sparser input views.

**The Impact of 2D Learning.** To validate the efficacy of the 2D texel paradigm for 3D Gaussian splats, we conducted an ablative experiment that predicts the 3D Gaussian parameters directly from 3D, termed as **w/ MLP**. We adopted an 8-layer MLP that consumes the skeletal motion  $\bar{\theta}_f$  and the positional-encoded canonical Gaussian position  $\bar{\mu}_i$ , predicting the Gaussian splat parameters  $\{\mathcal{G}_i\}$  in the canonical space. Specifically, the width for each hidden layer of the MLP is set to 256. Similar to ASH, the canonical Gaussian splats are wrapped to observation space through DQ skinning [21]. As illustrated in Fig. 1 and Tab. 1, directly learning the Gaussian parameters in 3D will lead to blurry rendering and cannot preserve the motion-dependent wrinkle details. In contrast, ASH, which formulates the learning of 3D Gaussian splats as image translation in 2D texel space, delivers high-quality rendering with delicate details.

## D. Runtime Analysis

In this section, we conduct a detailed runtime analysis for each major component in ASH. Specifically, we record the runtime for each major component when rendering a 1K (1285 × 940) image on a single Nvidia Tesla A100 graphics device. Additionally, the runtime analysis is benchmarked on models with different texture space resolution, specifi-

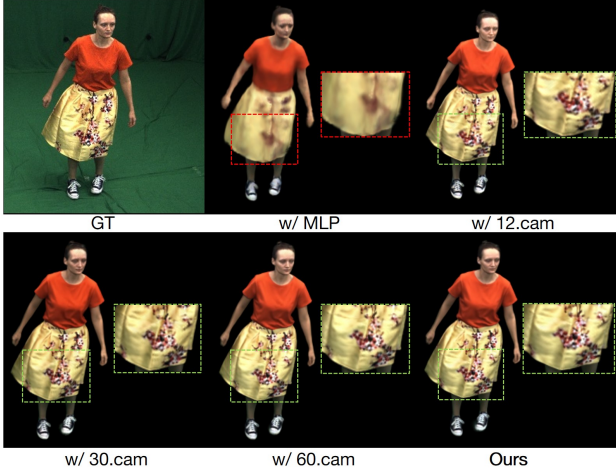


Figure 1. **Qualitative Ablation.** We compare ASH with the models that take alternative design choices. ASH excels in rendering quality than the model directly learns the Gaussian parameters from 3D canonical space (**w/ MLP**). Moreover, ASH exhibits robustness against less training views (**w/ 12.cam**, **w/ 30.cam**, **w/ 60.cam**).

Method	Stg.1	Stg.2	Stg.3	Stg.4	Time	FPS
w/ 128.res.	20.28	1.93	5.57	1.91	29.69	33.68
w/ 512.res.	24.25	16.34	18.45	4.31	63.35	15.79
<b>Ours</b>	<b>21.03</b>	<b>3.60</b>	<b>7.00</b>	<b>2.11</b>	<b>33.74</b>	<b>29.64</b>

Table 2. **Runtime Analysis.** We present detailed runtime for each major component in ASH measured in milliseconds. We also report the runtime of the models that take halved and doubled texel resolution, termed as **w/ 128.res.** and **w/ 512.res.**, respectively. Note that ASH can render high-quality animatable humans in a real-time frame rate.

cally at 128, 512, and 256, referred to as **w/ 128.res.**, **w/ 512.res.**, and **Ours**, respectively. Here, we divide the rendering pipeline of ASH into four steps as follows:

- Creating the deformable template meshes  $M(\theta_f)$  from skeletal motions  $\theta_f$  with structure-aware graph convolution networks, termed as **Stg.1**.
- Deriving motion-aware texture maps  $(\mathbf{T}_{n,f}, \mathbf{T}_{p,f})$  from deformable template meshes  $M(\theta_f)$ , termed as **Stg.2**.
- Predicting the canonical Gaussian splats  $\{\mathcal{G}_i\}$  with motion-aware geometry decoder  $\mathcal{E}_{\text{geo}}$  and appearance decoder  $\mathcal{E}_{\text{app}}$ , termed as **Stg.3**.
- Performing tile-based rasterization with the predicted Gaussian splats  $\{\mathcal{G}_i\}$ , termed as **Stg.4**.

Tab. 2 illustrates the runtime for each component in ASH for models with different 2D texel resolutions. While halving the texel resolution (**w/ 128.res.**) speeds up the image synthesis of the animatable humans, it may produce blurry details in the rendered images. Doubling the texel resolution (**w/ 512.res.**) results in comparable rendering quality.




Figure 2. **System Overview.** ASH Player is an interface runs in the browser, visualizing the imagery and skeletal poses of animatable characters. The renderings of the animatable humans are computed in real time from the GPU cluster server and streamed to the ASH Player front-end interface on a personal computer.

Nevertheless, it significantly increases computational complexity, preventing the model from being real-time compatible. In contrast, ASH can generate high-fidelity renderings of animatable characters in a real-time frame rate.

## E. Application

In this section, we introduce ASH Player, a real-time application built upon ASH.

Fig. 2 presents a screenshot of ASH Player, which runs in the web browser on a personal computer. The backend model of ASH Player, i.e., ASH, is deployed on the GPU cluster server. Once users specify the skeletal poses and virtual camera views, ASH Player will present the photoreal rendering of animatable characters, which is real-time computed and streamed from the GPU cluster server. Moreover, ASH Player allows users to inspect the animatable characters with spiral camera views. Please refer to the supplementary video for a more comprehensive visualization of ASH Player.

## F. Limitations

Although ASH enables the high-fidelity, real-time rendering of animatable human characters, it has certain limitations that we hope to address in the future.

Firstly, ASH does not extract detailed explicit geometry from the Gaussian splats. We will explore refining the explicit template meshes by backpropagating the gradient from image space into the template meshes using splatting. Additionally, ASH does not model topological changes like opening a jacket. Future research might focus on modeling the topological changes with the adaptive adding and removal of Gaussian splats introduced in the original 3D Gaussian splatting paper [22].