

# DRaCoN – Differentiable Rasterization Conditioned Neural Radiance Fields for Articulated Avatars

Amit Raj<sup>1\*</sup> Umar Iqbal<sup>2</sup> Koki Nagano<sup>2</sup> Sameh Khamis<sup>2</sup> Pavlo Molchanov<sup>2</sup>  
James Hays<sup>1</sup> Jan Kautz<sup>2</sup>  
<sup>1</sup>Georgia Institute of Technology <sup>2</sup>NVIDIA Research



Figure 1. We present a Differentiable Rasterization conditioned Neural Radiance field (*DRaCoN*) as an approach for learning multiple dynamic full body avatars from a monocular or multi-view RGB videos. A training frame of an identity is shown in the first column. Our model provides articulation-based controls for the body model and allows novel view rendering of multiple avatars in arbitrary poses.

## Abstract

Acquisition and creation of digital human avatars is an important problem with applications to virtual telepresence, gaming, and human modeling. Most contemporary approaches for avatar generation can be viewed either as 3D-based methods, which use multi-view data to learn a 3D representation with appearance (such as a mesh, implicit surface, or volume), or 2D-based methods which learn photo-realistic renderings of avatars but lack accurate 3D representations. In this work, we present, *DRaCoN*, a framework for learning full-body volumetric avatars which exploits the advantages of both the 2D and 3D neural rendering techniques. It consists of a Differentiable Rasterization module, *DiffRas*, that synthesizes a low-resolution version of the target image along with additional latent features guided by a parametric body model. The output of *DiffRas* is then used as conditioning to our conditional neu-

ral 3D representation module (*c-NeRF*) which generates the final high-res image along with body geometry using volumetric rendering. While *DiffRas* helps in obtaining photo-realistic image quality, *c-NeRF*, which employs signed distance fields (*SDF*) for 3D representations, helps to obtain fine 3D geometric details. Experiments on the challenging *ZJU-MoCap* and *Human3.6M* datasets indicate that *DRaCoN* outperforms state-of-the-art methods both in terms of error metrics and visual quality.

## 1. Introduction

The rising interest in creating digital embodiments of real people for personalized gaming and immersive telepresence applications necessitates the creation of photorealistic full-body 3D digital avatars that are both scalable and low-cost. Creating controllable full-body avatars from monocular or multi-view videos of people in normal cloth-

\*Work done during an internship at NVIDIA.

ing is a challenging problem due to the wide range of plausible human motions, the lighting variations, and pose-dependent clothing deformations. Most recent approaches for relevant avatar generation can be categorized into: (i) 2D neural rendering methods, or (ii) 3D neural rendering methods<sup>1</sup>.

2D neural rendering methods use 2D or 2.5D guiding information to conditionally render images of avatars in novel views or poses [15, 22, 28, 30, 36, 44]. These methods rasterize some intermediate semantic labels or geometric representations such as UV-images, shapes, or poses and use 2D CNN-based generators (*i.e.*, image translation network) to refine or render the final image from these intermediate representations. The main advantage of these methods is that they produce high-resolution photorealistic rendering results thanks to the high-resolution screen-space intermediate representations and high-capacity 2D CNN-based generators. However, these 2D methods show limited multi-view and temporal consistency in their rendering of novel views or poses despite requiring a large amount of training data, and do not learn the 3D representation of the object.

3D neural rendering methods, on the other hand, aim to learn the 3D representation and appearance of the object or person which can be rendered in novel views and poses. The state-of-the-art methods in this direction use implicit functions [31, 33] to model human avatars and use volumetric rendering for image generation [14, 25, 26, 39]. These methods not only provide 3D-consistent neural rendering but can also model pose-dependent deformations. While these implicit representations allow modeling the output as continuous functions, the complexity of the signals they can represent can be potentially limited when trying to model complex human body articulations and geometric variations, leading to less photorealistic outputs compared to the 2D-based neural rendering methods [15, 28, 30, 36, 44]. Furthermore, these approaches which only rely on fully implicit architectures usually train a single network per identity and lack the strong generalization power seen in the 2D CNN-based generators.

In this work, we propose Differentiable Rasterization Conditioned Neural Radiance Field (*DRaCoN*), that leverages the advantages of both 2D and 3D neural rendering methods. Our hybrid approach not only provides view-consistent rendering but also achieves higher photorealism. *DRaCoN* learns a radiance field conditioned on pixel-aligned features that are obtained by differentially rasterizing a high-dimensional latent neural texture for the target body pose and view using a parametric human body model, SMPL [18]. We enhance the rasterized neural texture with a 2D convolutional neural network which not only models the inter-part dependencies, but also allows us to render a low-resolution version of the target view before the volumetric

rendering. We provide this low-resolution image as input to the above-mentioned conditional radiance field network with additional neural and geometric features, thereby, simplifying the task of learning 3D view synthesis. We also use the SMPL body model to warp the posed avatars into a canonical space. This allows us to generate avatars in arbitrary poses since the radiance fields are learned in the canonical space without having to allocate the capacity for modeling the body pose deformations. We demonstrate the effectiveness of our proposed approach on the challenging ZJU-MoCap and Human3.6M datasets.

To summarize, our contributions are as follows:

1. We present a novel hybrid framework for generating full-body, clothed, and dynamic avatars that combines screen-space and volumetric neural representations.
2. Our framework learns identity-specific appearance space, allowing fine-grained dynamic texture synthesis during pose retargeting or change without having input images during inference.
3. We demonstrate the state-of-the-art performance on novel view and novel pose synthesis compared to recent volumetric approaches.

## 2. Related Work

The existing works in learning articulated human avatars from RGB images or videos can be mainly classified into two categories – 2D neural rendering-based methods and 3D neural rendering-based methods [41].

### 2.1. 2D Neural Rendering Methods

These methods aim to directly generate images of a person in novel views and poses using image-space rendering. Earlier methods treat the problem as an image-to-image translation task. Given an image of a target person with image representations of keypoints or dense meshes or semantic labels, these methods use an image translation model to directly render the target person in the style or pose of the source person possibly in different views [4, 5, 12, 19, 23, 27, 29, 44, 46, 49, 53]. The main limitation of these methods is that they lack multi-view and temporal consistency under significant viewpoint and pose changes since these 2D-based method generally do not learn any notion of 3D space. Instead of fully relying on the image-space translation network, Liquid Warping GAN [15, 16] uses UV-correspondences between the source and target meshes to explicitly warp the source image to the target pose and then uses an additional network to refine the warped image. Given a target pose in form of 2D keypoints, Huang *et al.* [8] predict the UV-coordinates of a dense mesh in image space and then generate the target image by sampling from a learned UV texture map. The main

<sup>1</sup>We follow the 2D and 3D neural rendering definitions by [41]

limitation of these methods is that they require some exemplar images to be available during inference, limiting their semantic editing capability, such as fine-grained control of the target appearance.

Textured Neural Avatar [36] removes the need of exemplar images but relies on multi-view data for training. Articulated Neural Renderer (ANR) [30] builds on a similar idea but instead uses reconstructed 3D meshes to rasterize a learned neural texture on image space from which the final rendered image is obtained using a convolutional network. While these approaches show impressive results and have the capability to generate photorealistic images, their output is limited to 2D images only and does not provide 3D representations of the person. Inspired by these methods, we propose a method that not only renders photorealistic images but also learns 3D representation of the person, allowing view-consistent rendering under large view or pose changes.

## 2.2. 3D Neural Rendering Methods.

These methods aim to create geometric representations of avatars as well as their 3D appearance that can be rendered in novel viewpoints and poses.

**Mesh-Based Approaches.** Given a monocular or multi-view video, earlier methods reconstruct the detailed geometry and textures using parametric mesh fitting [1, 2, 7, 52]. Alldieck *et al.* [2] fit the SMPL body model to all frames and optimize for per-vertex offsets using person silhouette information to capture the clothing and hair details. A detailed texture map is then generated by back-projecting the image colors to all visible vertices. The novel views and poses of the person can be easily generated by articulating the SMPL mesh and rendering it to the image with the obtained texture map. The methods [1, 52] adopt a similar strategy but replace the optimization-based framework used in [2] with learning-based components for faster processing. In all cases, they produce a parametric mesh with a texture map which can be articulated and rendered in a new view. However, the main limitation of these methods is that they rely on a fixed mesh topology which cannot accurately capture complex clothing and hair geometries. Also, the articulated meshes do not account for pose-conditioned geometric deformations and changes in texture. Hence, the rendered results lack realism and high-frequency details. In contrast, our volumetric representation directly learns these details that are not captured by the parametric models.

**Implicit Approaches.** More recently, PIFu and PIFuHD [32, 34] proposed to use implicit functions and pixel-aligned features for human reconstruction from single or multi-view RGB images. Since the implicit functions represent the body’s surface in a continuous 3D space, these methods can reconstruct highly-detailed human body shape. However, the reconstructed geometries are not rigged and can-

not be articulated in novel poses. To this end, ARCH [9] and S3 [47] also reconstruct meshes that can be articulated with novel poses. However, similar to other mesh-based models, rendering the reconstructed bodies in novel poses does not account for pose-dependent deformations. Scanimate [35] learns to fit a deformable implicit function of surface data to RGBD scans, allowing for additional articulation but requires 3D supervision to train.

The most recent methods take inspiration from NeRF [21] or its variants [3, 13, 24, 43, 50] and represent human avatars using pose-conditioned implicit 3D representations [11, 14, 25, 26, 42]. We refer the readers to the comprehensive report on advances in neural rendering [41] for additional insight about these approaches. These methods not only allow photorealistic novel view synthesis but also provide detailed human body geometry with pose dependent variations. NeuralBody [26] exploits the SMPL body model and learn latent codes corresponding to each vertex. These latent codes are then diffused using a sparse convolution module and sampled according to the target body pose and used to condition the NeRF model. A-NeRF [39] and AnimatableNeRF [25] use body pose information to canonicalize the sampled rays and learn neural radiance fields in the canonical space, which helps the learned avatar to generalize across different poses. Our hybrid framework leverage advantages of multiple representations: the parametric model provides articulation-based controls and promotes generalization through learnable neural textures enhanced by a 2D CNN-based generator, while our NeRF enables conditional 3D view synthesis and high-fidelity 3D geometry extraction.

## 3. Method

Our goal is to learn a controllable full-body avatar that can synthesize geometrically consistent images of the person in novel views using 3D neural rendering, while also leveraging the advantages of 2D neural rendering for photorealism. For training, we assume a set of calibrated multi-view or monocular videos is available. We also assume that the human bodies in the training video frames have been tracked using the SMPL [18] body model. An overview of our approach can be seen in Figure 2. Our framework consists of three modules: *differentiable rasterization (DiffRas)* (Section 3.2), *WarpField* (Section 3.3) module and *conditional-NeRF* module (Section 3.4). Given a target pose represented as the parameters of SMPL body model, DiffRas rasterizes high-dimensional identity-specific neural features in the target view. These pixel-aligned features are then used as a condition during 3D neural rendering. To simplify the learning of otherwise highly dynamic human bodies, the WarpField module maps the sampled 3D points in 3D world space to a canonical space. Finally, the canonicalized 3D points and pixel-aligned features are fed

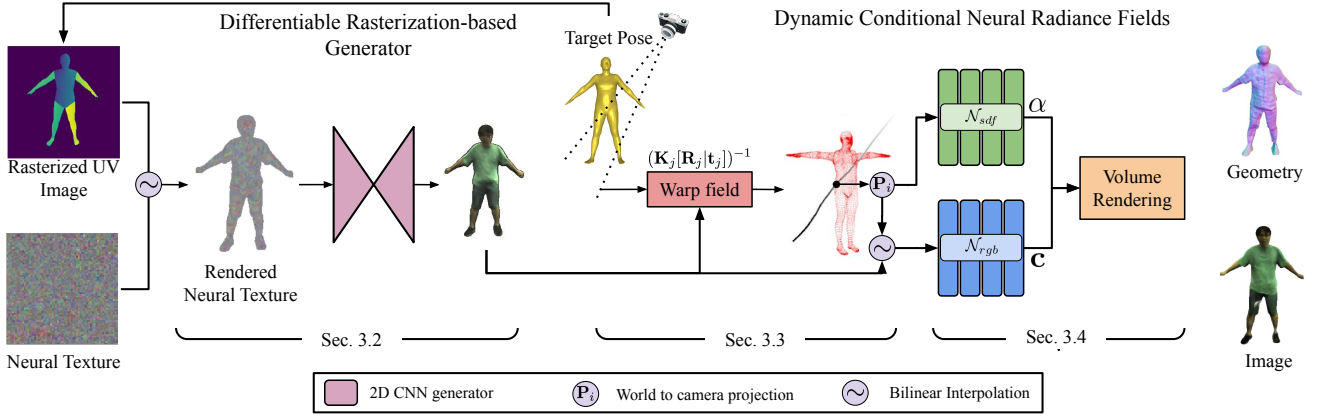


Figure 2. Overview of the proposed approach. Given a target camera and a parametric body model of an arbitrary pose, DRaCoN *i*) employs a differentiable rasterizer with identity-specific neural textures to rasterize pixel-aligned neural features for the target view; *ii*) warps the 3D rendering points from the posed space to canonical volume using a learned warping field and samples neural features; *iii*) uses the rendered pixel-aligned features as input to the conditional neural 3D representation module (c-NeRF) which produces the final image using volumetric rendering. We can also extract fine surface geometry as the zero-level set of the predicted signed distance values.

as input to the conditional neural 3D representation module (c-NeRF) which generates the final high-res image and detailed 3D surface geometry using volumetric rendering. In the following, we explain each of these modules in greater details.

### 3.1. Preliminaries

Given a set of posed images  $\{\mathbf{I}_i \in \mathbb{R}^{H \times W \times 3}\}_{i=1}^N$  with their associated intrinsic  $\{\mathbf{K}_i\}_{i=1}^N$  and extrinsic camera parameters (rotation  $\{\mathbf{R}_i\}_{i=1}^N$  and translation  $\{\mathbf{T}_i\}_{i=1}^N$ ), we define the world to camera projection as  $\mathbf{P}_i = \mathbf{K}_i [\mathbf{R}_i | \mathbf{T}_i]$ . We also assume that the pose and shape parameters ( $\theta_i \in \mathbb{R}^{24 \times 3}, \beta_i \in \mathbb{R}^{10}$ ) of the tracked SMPL meshes for each frame  $i$  are available from which the 3D mesh vertices  $\mathbf{V}$  can be obtained using a linear function  $\mathbf{V}_i = \mathcal{M}(\theta_i, \beta_i)$  as defined in [18].

NeRF [21] learns a function  $\mathcal{N}(\mathbf{x}, \mathbf{d}) : \mathbb{R}^{N_x \times N_d} \rightarrow \mathbb{R}^4$  which takes the positional encodings of a world space position  $\mathbf{x}$  and viewing direction  $\mathbf{d}$  as input and returns the radiance  $\mathbf{c} \in \mathbb{R}^3$  and associated differential opacity  $\alpha \in \mathbb{R}$  at that location:

$$\mathbf{c}, \alpha = \mathcal{N}(\gamma_L(\mathbf{x}), \gamma_M(\mathbf{d})),$$

where  $\gamma_L(\mathbf{x})$  and  $\gamma_M(\mathbf{d})$  are the positional encodings of  $\mathbf{x}$  and  $\mathbf{d}$  with  $L$  and  $M$  octaves, respectively, as described in [20]. For ease of readability we drop the dependence on  $\mathbf{d}$  in all future discussion.

The NeRF network is trained by tracing rays into the scene from a camera origin through each pixel of the image. For image  $i$  the  $j$ th pixel location is given by  $u_i^j \in \mathbb{R}^2$  and the camera origin  $\mathbf{o} \in \mathbb{R}^3$  and ray direction  $\mathbf{d} \in \mathbb{R}^3$  are calculated as follows:

$$\mathbf{o}_i = -\mathbf{R}_i^{-1} \mathbf{T}_i, \quad (1)$$

$$\mathbf{d}_i^j = \frac{\mathbf{P}_i^{-1} u_i^j - \mathbf{o}_i}{\|\mathbf{P}_i^{-1} u_i^j - \mathbf{o}_i\|}. \quad (2)$$

We then sample points uniformly along the ray given by

$$\mathbf{x}_i^j(t) = \mathbf{o}_i + t \mathbf{d}_i^j.$$

We drop the parameterization variable  $t$  in all subsequent equations for notational simplicity.

### 3.2. Differentiable Rasterization (DiffRas)

Our goal is to generate pixel-aligned features  $\mathbf{F} \in \mathbb{R}^{H \times W \times D}$  capturing the geometric and appearance properties of the person in the target viewpoint and body pose. In order to make sure that these features are consistent across different body poses and views, we define them in the canonical UV-space of the SMPL body model and rasterize them with the 3D mesh. Specifically, we learn a set of  $D$ -dimensional ( $D=128$ ) neural features  $\mathbf{Z}_i \in \mathbb{R}^{M \times M \times D}$  where the spatial coordinates correspond to the UV-coordinates of SMPL. Given the SMPL vertices  $\mathbf{V}_i$ , DiffRas differentially rasterizes these neural features to the image space:  $\text{DiffRas}(\mathbf{Z}, \mathbf{V}_i) \rightarrow \mathbf{F}_i$ . The learnable neural features  $\mathbf{Z}$  are optimized end-to-end using the image reconstruction losses over all training frames.

Unlike previous works which rely on pixel-aligned features from exemplar images, [31, 51], the addition of the differentiable rasterizer eschews the need for input images during inference as the pixel-aligned features can now be rasterized on the fly for any arbitrary viewpoint for any identity. In practice, we enforce the first three channels of  $\mathbf{F}_i$  to be a low resolution version the target image. This allows to exploit the latest advances in 2D neural rendering, for example, by using photometric and perceptual losses on the generated images to improve photo-realism. Also, since

the features are already rasterized in the target pose, it simplifies the problem for subsequent components of our approach. In contrast to existing methods [25, 26, 39] which encode all appearance and geometric details in the NeRF module using a highly non-linear mapping, our approach learns these details in a well-defined UV-space thereby simplifies neural avatar learning.

### 3.3. Warping Fields

The NeRF model as originally proposed in [21] is trained for static scenes only. To simplify learning of the complex articulated geometries of human bodies, we map sampled points  $\mathbf{x}$  in posed space to canonical space before passing them to c-NeRF using a warping field  $\mathcal{W}$ . Particularly, each point  $\mathbf{x}_i^j$  in posed space is mapped to point  $\hat{\mathbf{x}} = \mathcal{W}(\mathbf{x})$  in the canonical space.

However, as shown in [17], learning the warping field and  $\mathcal{N}(\mathbf{x})$  simultaneously is an under-constrained problem. To this end, we exploit the tracked SMPL meshes to regularize the learning of the warping field. Specifically, let  $G_k(\theta)$  represent the local transformation of joint  $k$  with respect to its parent in the SMPL body model, we define the transform to warp a canonical vertex  $\hat{\mathbf{v}}_l \in \mathbb{R}^3$  to posed space as follows:

$$\phi(\mathbf{v}_l) = \sum_{m=1}^J w_m(\mathbf{v}_l) \left( \prod_{k \in \mathcal{K}_m} G_k(\theta) \right), \quad (3)$$

where  $\mathcal{K}_m$  is the kinematic chain for joint  $m$  and  $w_m(\mathbf{v})$  are the skinning weight associated with vertex  $\mathbf{v}_l$ . The location of the posed vertex is obtained as:  $\mathbf{v}_l = \phi(\hat{\mathbf{v}}_l)\hat{\mathbf{v}}_l$ .

We use this skinning field to initialize the warping field  $\mathcal{W}$ . Specifically, for any point  $\mathbf{x}_i^j$  in posed space, we find the vertex  $\mathbf{v} \in \mathbf{V}_i$  that is closest to this point, and use the associated skinning weights to canonicalize  $\mathbf{x}_i^j$ :

$$p(\mathbf{x}) = \min_{\mathbf{v} \in \mathbf{V}} \|\mathbf{x} - \mathbf{v}\|, \quad (4)$$

$$\hat{\mathbf{x}} = \mathcal{W}(\mathbf{x}) = \phi^{-1}(p(\mathbf{x}))\mathbf{x}. \quad (5)$$

**Learning Residuals for Warp Field.** The skinning weights associated with the naked body model do not account for the geometry outside the parametric mesh or pose-based geometric deformation of clothing. To account for these variations, we learn a residual field  $\Delta\mathcal{W}(\mathbf{x}, \mathbf{f}_i^{\mathbf{x}}) : \mathcal{R}^{3+d} \rightarrow \mathcal{R}^3$  and obtain the final canonicalized point  $\hat{\mathbf{x}}^c$  as:

$$\hat{\mathbf{x}}^c = \hat{\mathbf{x}} + \Delta\mathcal{W}(\mathbf{x}, \mathbf{f}_i^{\mathbf{x}}).$$

Here  $\mathbf{f}_i^{\mathbf{x}} \in \mathbf{F}_i$  are the pixel-aligned feature associated with point  $\mathbf{x}$  obtained from the differentiable rasterizer (Section 3.2).

### 3.4. Conditional 3D Representation

The DiffRas module produces a low resolution image with no associated geometric information. To this end, we learn a conditional neural 3D representation (c-NeRF) model that takes as input  $\hat{\mathbf{x}}^c$  and pose-specific features  $\mathbf{f}_i^{\mathbf{x}}$ , and produces a high-fidelity 3D surface geometry as well as generates a higher resolution version of the image. We follow [43, 48] and represent the avatar using two functions  $\mathcal{N}_{sdf}$  and  $\mathcal{N}_{color}$  responsible for modeling surface geometry and texture, respectively. The neural signed distance function (SDF)  $\mathcal{N}_{sdf} : \mathbb{R}^{L+D} \rightarrow \mathbb{R}$  takes as input the positional encoding of the canonicalized point  $\hat{\mathbf{x}}^c$  and pose-specific features  $\mathbf{f}_i^{\mathbf{x}}$  and produces the signed distance of the point w.r.t. body surface geometry. Following [48], we generate the differential opacity at the point from the SDF using the following function:

$$\alpha = \Psi(\mathcal{N}_{sdf}(\hat{\mathbf{x}}^c, \mathbf{f}_i^{\mathbf{x}}), \beta),$$

where  $\Psi$  is the cumulative distribution function (CDF) of a laplacian distribution with scale  $\beta$  which is a learnable parameter. The function  $\mathcal{N}_{rgb}$  also takes  $\hat{\mathbf{x}}^c$  and pose-specific features  $\mathbf{f}_i^{\mathbf{x}}$  and produces the color,  $\mathbf{c} = \mathcal{N}_{rgb}(\hat{\mathbf{x}}^c, \mathbf{f}_i^{\mathbf{x}})$ , associated with the point. To better inform the input, we also concatenate the UV-coordinates of the sampled features as an additional information. The final RGB image is generated using volume rendering and the 3D geometry is obtained as the the zero-level set of the SDF.

Note that using SDF to represent body surface has the advantage that it is interpretable and the geometry can be extracted easily without the need of a careful selection of threshold for density values as used in existing NeRF-based neural avatar methods [25, 26, 39].

### 3.5. Training and Losses

**Texture Branch:** The c-NeRF network is trained with a photometric reconstruction loss for the color  $C(u_i^j)$  accumulated along each ray w.r.t ground truth color  $\hat{C}(u_i^j)$ :

$$\mathcal{L}_{\text{recon}} = \|C(u_i^j) - \hat{C}(u_i^j)\|_2.$$

We also find that the warping module can “cheat” by generating large  $\Delta\mathbf{x}$  values that place the sampled points at arbitrary locations to help optimize the learning. This causes generalization to suffer as the mapping for novel poses become uncertain. We constrain the norm of the deformation to limit the maximal displacement that the warping module can produce:

$$\mathcal{L}_{\text{delta}} = \|\Delta\mathcal{W}(\mathbf{x}, \mathbf{f}_i^{\mathbf{x}})\|_2$$

The c-NeRF network is then optimized as follows:

$$\mathcal{L}_{\text{nerf}} = \lambda_{\text{recon}}\mathcal{L}_{\text{recon}} + \lambda_{\text{delta}}\mathcal{L}_{\text{delta}} \quad (6)$$

**DiffRas Branch:** While there is no exact supervision for the output of DiffRas module  $\mathbf{F}$ , we constrain the first three channels of  $\mathbf{F}$  as in ANR [30] as follows:

$$\begin{aligned}\mathcal{L}_{\text{pixel}} &= \|\mathbf{F}_i[:3] - \mathbf{I}_i\|. \\ \mathcal{L}_{\text{vgg}} &= \sum_l w_l \|\phi_l(\mathbf{F}_i[:3]) - \phi_l(\mathbf{I}_i)\|.\end{aligned}$$

Where,  $\phi_l$  return features at different layers of a pretrained VGG-network [37]. Additionally, we constrain the fourth channel to learn the mask of the generated identity.

$$\mathcal{L}_{\text{mask}} = BCE(\mathbf{F}_i[4], \mathbf{M}_i).$$

The total loss for the DiffRas branch is then computed as:

$$\mathcal{L}_{\text{DiffRas}} = \lambda_{\text{pixel}} \mathcal{L}_{\text{pixel}} + \lambda_{\text{vgg}} \mathcal{L}_{\text{vgg}} + \lambda_{\text{mask}} \mathcal{L}_{\text{mask}}. \quad (7)$$

**Geometry Branch:** The geometry network,  $\mathcal{N}_{\text{sdf}}$ , is encouraged to learn SDF by using the eikonal regularization [6]:

$$\mathcal{L}_{\text{eik}} = \|\Delta \mathcal{N}_{\text{sdf}}(\mathbf{x}, \mathbf{f}^{\mathbf{x}}) - 1\|. \quad (8)$$

Additionally, to encourage smoothness of the learnt geometry we apply a Total Variation (TV) loss over the estimated density as follows:

$$\mathcal{L}_{\text{TV}} = \|\Psi(\mathbf{x}) - \Psi(\mathbf{x} + \delta\mathbf{x})\|. \quad (9)$$

Furthermore, following Sitzmann et al. [38], we apply exponential regularization for encouraging non-zero SDF at almost all locations except the surface:

$$\mathcal{L}_{\text{exp}} = \mathbb{E}_{\mathbf{x}}[\exp(-\alpha * \mathcal{N}_{\text{sdf}}(\mathbf{x}))] \quad (10)$$

where  $\alpha \gg 1$ . Finally, we also use an SMPL guided SDF regularization,  $\mathcal{L}_{\text{face}}$ , for the face region that encourages SDF of the samples around face regions to be close to SDF w.r.t to SMPL mesh. The regularization for the geometry is then given by:

$$\mathcal{L}_{\text{geom}} = \lambda_{\text{eik}} \mathcal{L}_{\text{eik}} + \lambda_{\text{TV}} \mathcal{L}_{\text{TV}} + \lambda_{\text{exp}} \mathcal{L}_{\text{exp}} + \lambda_{\text{face}} \mathcal{L}_{\text{face}} \quad (11)$$

The training objective for our framework is thus given by

$$\mathcal{L}_{\text{DRaCoN}} = \underbrace{w_{\text{nerf}} \mathcal{L}_{\text{nerf}}}_{\text{NeRF objective}} + \underbrace{w_{\text{DiffRas}} \mathcal{L}_{\text{DiffRas}}}_{\text{DiffRas objective}} + \underbrace{w_{\text{geom}} \mathcal{L}_{\text{geom}}}_{\text{Geometry Regularization}} \quad (12)$$

### 3.6. Progressive Growing

In practice, we first train both the DiffRas and c-NeRF at a low resolution and then train c-NeRF at successively higher resolutions while keeping the resolution of DiffRas constant. This progressive training accelerates the convergence since the network mostly needs to focus on the coarse structures early in the training. We also calculate a per-frame affine transform that crops the image in the scene to the tightest  $256 \times 256$  image. This ensures that the DiffRas always sees similarly sized images regardless of where the actor is in the world space.

## 4. Experiments

We evaluate our results with datasets commonly used for learning full body dynamic avatars from videos and compare results with state-of-the-art baselines.

### 4.1. Dataset

**ZJU Mocap [26]:** The ZJU mocap dataset consists of actors performing various activities captured in 23 videos. In all our settings, we train our model on 4 views (as in [26]) on 300 frames. We train the multi-identity model on 4 simultaneous identities. We evaluate novel view information on the remaining cameras of training poses and novel pose information on test cameras of unseen poses.

**Human3.6M [10]:** We follow the protocol defined in [25] and train our model on 3 cameras and evaluate the performance of our method on the remaining cameras for novel view and poses on a selected subset of of Human3.6M dataset.

### 4.2. Baselines

We compare our approach with the following baselines: **Animatable NeRF [25]**. Canonicalizes sampled points based on the skinning weights of a parametric mesh and learns a single body model for each data. The main difference of [25] and our approach involves the residual calculation. Particularly, animatable-nerf learns residual blend weights rather than a residual in canonical space as in our framework.

**NeuralBody [26]**. NeuraBody learns to associate a latent code with each vertex of a deformable body model and then performs ray tracing on the diffused version of these learned latent codes.

**A-NeRF [39]**. Uses only pose information and defines pose relative encodings to learn an avatar that can be rendered in novel views and poses.

**NHR [45]**. Neural Human Rendering uses sparse point clouds to render multi-view images without explicit pose based reasoning.

### 4.3. Quantitative Analysis

We evaluate the approaches using MSE, PSNR, SSIM and LPIPS metrics for all testing frames in all settings. Table 1 and Table 2 compare our approach with the state-of-the-art methods on ZJU and Human3.6M respectively. We use S9 and S11 for evaluation for Human3.6M, which are the commonly used testing subjects. The generated images used for evaluation of the state-of-the-art methods are provided by [25].

### 4.4. Qualitative Analysis

Fig. 3 and Fig. 5 shows the performance of our method compared to above methods on the ZJU-Mocap and Hu-




Figure 3. Comparison of novel view and pose synthesis (a) A-NeRF [39] (b) Animatable NeRF [25] (c) NeuralBody [26] (d) Ours. Our model is able to retain textures better owing to conditional features from DiffRas while also produces better person shapes.

	ZJU-315				ZJU-377				User Study	
	SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$		SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$		Num Votes	Mean preference
			Alexnet	VGG			Alexnet	VGG		
A-NeRF	0.930	17.711	0.094	0.073	0.927	16.528	0.101	0.077	53	$7.5 \pm 1.5\%$
Anim-NeRF	0.917	16.774	0.101	0.083	0.940	17.971	0.075	0.066	15	$2.5 \pm 0.64\%$
NB	<b>0.947</b>	18.608	0.092	0.090	<b>0.950</b>	20.072	0.071	0.060	168	$24 \pm 3.25\%$
Ours	0.922	<b>21.315</b>	<b>0.053</b>	<b>0.046</b>	0.946	<b>21.084</b>	<b>0.048</b>	<b>0.044</b>	464	<b><math>66 \pm 4\%</math></b>

Table 1. Novel pose synthesis results on ZJU-Mocap dataset. Our approach outperforms recent baselines on most supervised metrics




Figure 4. Novel view (above) and novel pose (below) results of each stage of our model. We see that the c-NeRF module adds necessary details using intermediate output from the DiffRas stage. The geometry is obtained by sampling  $\mathcal{N}_{sdf}$  using a uniform grid and running marching cubes over it.

man3.6 dataset. We observe that Animatable-NeRF performs poorly on novel poses without additional refinement. Inference time refinement is inconvenient as it involves an optimization step which adds some cost to the inference

speeds. Additionally, as seen in the last row of Fig. 5, Animatable-NeRF fails to maintain the shape of the actor under novel poses. NeuralBody generates high quality results for novel views, however struggles with novel poses. It collapses in a multi-ID setting. We demonstrate performance comparison of our method against a multi-ID NeuralBody setting in the supplementary. Additionally, We see that A-NeRF overfits to training views due to its high dimensional pose relative encoding which require large number of poses to train. In contrast our c-NeRF renderer is able to generalize to novel poses, due to pixel level conditioning information from intermediate neural features. Fig 4 shows the intermediate output of the DiffRas stage and the reconstructed geometry for the novel pose. We see that the c-NeRF adds additional texture information while also providing access to deformable geometry of the actor.




Figure 5. Comparison of novel view (odd rows) and novel pose (even rows) synthesis w.r.t NHR, Neuralbody, and Animatable-NeRF. Our model preserves the shape of the actor under novel poses.




Figure 6. We demonstrate the importance of each module in our framework by comparing against models trained under the following settings: (a) without  $\mathcal{L}_{DiffRas}$  (b) without  $\mathcal{L}_{face}$  (c) without  $\mathcal{L}_{eik}$  and (d) without  $\mathcal{L}_{delta}$

#### 4.5. User study

Since most of the metrics used serve mostly as a proxy to perceptual quality, we perform a user study to compare avatars generated using different methods. Particularly, each user is presented with 20 randomly sampled views in novel poses across 4 different identities for 8 seconds. The user is asked to rate which of the 4 avatars looks most realistic. We collected this data over 35 random subjects. The demographics include both technical experts who work on digital avatars and general audience. The results (Table. 1) of the study indicate that our method is preferred 66 % of the time as compared to the the next highest rated method (NeuralBody [26]) at 24%.

#### 4.6. Ablation Study

We study the following important questions with respect to our framework.

	MSE ↓	SSIM ↑	PSNR ↑	LPIPS ↓	
				Alexnet	VGG
D-NeRF	2.34	0.946	26.47	0.082	0.058
NHR	0.82	0.976	27.91	0.026	0.024
NB	0.66	0.977	27.93	0.039	0.029
Anim-NeRF	0.50	0.979	28.11	0.035	0.028
Ours	<b>0.32</b>	<b>0.985</b>	<b>29.42</b>	<b>0.023</b>	<b>0.022</b>

Table 2. Novel pose synthesis results on Human3.6M dataset. Our approach outperforms recent baselines on most supervised metrics

**Need for DiffRas Objective** The DiffRas objective constrains the first four channels of the intermediate features, we see in Fig. 6 and Table 3 that without this constraint, the c-NeRF module cannot effectively learn higher frequency information and the results are significantly noisier. This constraint helps ease the learning problem by providing supervision for lower frequency information (since it is trained at a lower resolution).

**Effect of Face Regularization** The advantage of the SDF representation is seen by the use of this regularization. Since tracking errors can lead to significant facial blurring, regularizing the face shape w.r.t base SMPL mesh alleviates the texture averaging in facial region to some extent.

	MSE ↓	SSIM ↑	PSNR ↑	LPIPS ↓
Ours( $-\mathcal{L}_{DiffRas}$ )	3.511	0.957	24.546	0.053
Ours( $-\mathcal{L}_{eik}$ )	1.767	0.975	27.526	0.033
Ours( $-\mathcal{L}_{face}$ )	1.785	0.975	27.485	0.031
Ours( $-\mathcal{L}_{delta}$ )	<b>1.402</b>	0.977	27.531	0.031
Ours	1.568	<b>0.977</b>	<b>28.045</b>	<b>0.026</b>

Table 3. Ablation results on models trained on ID-337 of ZJU-mocap dataset

## 5. Discussion

**Limitations and Future Work.** Although our approach is robust to large pose or view changes, tracking errors during parametric body fitting can create inaccurate canonicalization or misalignments in neural textures. It may be fruitful to explore directions which do not require explicit parametric model fitting [35] or learn 3D surfaces directly from videos [24, 43, 48]. Also our end-to-end inference time is potentially slow for real-time applications and interesting future directions include scaling the approach to both high-resolution and real-time rendering.

**Ethical Considerations.** A technique to synthesize photo-realistic images of people from a monocular video could be misused to create manipulated imagery of real people. Such misuse of media synthesis techniques pose a societal threat. Viable countermeasure solutions include watermarking the model data or output [40, 51].

**Conclusion.** By combining parametric body models with high-dimensional neural textures and dynamic conditional



radiance field network, our approach demonstrates important steps toward the generation of controllable and photorealistic 3D avatars of multiple people. We believe the avatar creation from a casual video input may lead to the democratization of the high-quality personalized digital avatars.

## References

- [1] 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 *CVPR*, June 2019. 3
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, June 2018. 3
- [3] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. *ICCV*, 2021. 3
- [4] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *ICCV*, pages 5933–5942, 2019. 2
- [5] Patrick Esser, Ekaterina Sutter, , and Bjorn Ommer. A variational U-Net for conditional appearance and shape generation. In *CVPR*, 2018. 2
- [6] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *ICML*, pages 3569–3579. 2020. 6
- [7] Marc Habermann, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular human performance capture using weak supervision. In *CVPR*, 2020. 3
- [8] Zhichao Huang, Xintong Han, Jia Xu, and Tong Zhang. Few-shot human motion transfer by personalized geometry and texture modeling. In *CVPR*, 2021. 2
- [9] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. ARCH: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 3
- [10] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. *TPAMI*, 36(7):1325–1339, 2014. 6
- [11] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. In *NeurIPS*, 2021. 3
- [12] Christoph Lassner, Gerard Pons-Moll, and Peter V Gehler. A generative model of people in clothing. In *ICCV*, pages 853–862, 2017. 2
- [13] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. BARF: Bundle-adjusting neural radiance fields. *ICCV*, 2021. 3
- [14] 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. *SIGGRAPH ASIA*, 2021. 2, 3
- [15] Wen Liu, Zhixin Piao, Min Jie, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *ICCV*, 2019. 2
- [16] Wen Liu, Zhixin Piao, Zhi Tu, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan with attention: A unified framework for human image synthesis. *TPAMI*, 2021. 2
- [17] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics (TOG)*, 38(4):65, 2019. 5
- [18] 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)*, 2015. 2, 3, 4
- [19] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *NeurIPS*, pages 406–416, 2017. 2
- [20] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019. 4
- [21] 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, 4, 5
- [22] Koki Nagano, Jaewoo Seo, Jun Xing, Lingyu Wei, Zimo Li, Shunsuke Saito, Aviral Agarwal, Jens Fursund, and Hao Li. pagan: real-time avatars using dynamic textures. *ACM Transactions on Graphics (TOG)*, 37(6):1–12, 2018. 2
- [23] Natalia Neverova, Riza Alp Guler, and Iasonas Kokkinos. Dense pose transfer. In *ECCV*, pages 123–138, 2018. 2
- [24] Michael Oechsle, Songyou Peng, and Andreas Geiger. UNISURF: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *ICCV*, 2021. 3, 8
- [25] 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 *CVPR*, 2021. 2, 3, 5, 6, 7, 8
- [26] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 2, 3, 5, 6, 7, 8
- [27] Unsupervised person image synthesis in arbitrary poses. Albert pumarola and antonio agudo and alberto sanfelix and and francesc moreno-noguer. In *CVPR*, 2018. 2
- [28] Sergey Prokudin, Michael J Black, and Javier Romero. Smpix: Neural avatars from 3d human models. In *WACV*, 2021. 2
- [29] Amit Raj, Patsorn Sangkloy, Huiwen Chang, James Hays, Duygu Ceylan, and Jingwan Lu. Swapnet: Image based garment transfer. In *ECCV*, pages 679–695. Springer, 2018. 2
- [30] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. ANR-articulated neural rendering for virtual avatars. In *CVPR*, 2021. 2, 3, 6

- [31] Amit Raj, Michael Zollhoefer, Tomas Simon, Jason Saragih, Shunsuke Saito, James Hays, and Stephen Lombardi. Pva: Pixel-aligned volumetric avatars. In *CVPR*, 2020. 2, 4
- [32] 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
- [33] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, June 2020. 2
- [34] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, June 2020. 3
- [35] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, pages 2886–2897, 2021. 3, 8
- [36] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliev, Renat Bashirov, Egor Burkov, Karim Iskakov, Aleksei Ivakhnenko, Yury Malkov, Igor Pasechnik, Dmitry Ulyanov, et al. Textured neural avatars. In *CVPR*, pages 2387–2397, 2019. 2, 3
- [37] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [38] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *NeurIPS*, 33:7462–7473, 2020. 6
- [39] 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. 2, 3, 5, 6, 7
- [40] Matthew Tancik, Ben Mildenhall, and Ren Ng. Stegastamp: Invisible hyperlinks in physical photographs. In *CVPR*, 2020. 8
- [41] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Yifan Wang, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. *arXiv preprint arXiv:2111.05849*, 2021. 2, 3
- [42] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-GIF: Neural generalized implicit functions for animating people in clothing. In *CVPR*, 2021. 3
- [43] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NeurIPS*, 2021. 3, 5, 8
- [44] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *NeurIPS*, 2018. 2
- [45] Minye Wu, Yuehao Wang, Qiang Hu, and Jingyi Yu. Multi-view neural human rendering. In *CVPR*, 2020. 6, 8
- [46] Lingbo Yang, Pan Wang, Chang Liu, Zhanning Gao, Peiran Ren, Xinfeng Zhang, Shanshe Wang, Siwei Ma, Xiansheng Hua, and Wen Gao. Towards fine-grained human pose transfer with detail replenishing network. In *TIP*, 2020. 2
- [47] Ze Yang, Shenlong Wang, Sivabalan Manivasagam, Zeng Huang, Wei-Chiu Ma, Xinchun Yan, Ersin Yumer, and Raquel Urtasun. S3: Neural shape, skeleton, and skinning fields for 3d human modeling. In *CVPR*, 2021. 3
- [48] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *NeurIPS*, 34, 2021. 5, 8
- [49] Jae Shin Yoon, Lingjie Liu, Vladislav Golyanik, Kripasindhu Sarkar, Hyun Soo Park, and Christian Theobalt. Pose-guided human animation from a single image in the wild. In *CVPR*, 2021. 2
- [50] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. *ICCV*, 2021. 3
- [51] Ning Yu, Vladislav Skripniuk, Sahar Abdelnabi, and Mario Fritz. Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In *ICCV*, 2021. 4, 8
- [52] Tiancheng Zhi, Christoph Lassner, Tony Tung, Carsten Stoll, Srinivasa G Narasimhan, and Minh Vo. Texmesh: Reconstructing detailed human texture and geometry from rgb-d video. In *ECCV*, 2020. 3
- [53] Zhen Zhu, Tengting Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive pose attention transfer for person image generation. In *CVPR*, 2019. 2