

TexVocab: Texture Vocabulary-conditioned Human Avatars

Yuxiao Liu¹ Zhe Li² Yebin Liu² Haoqian Wang¹

¹Shenzhen International Graduate School, Tsinghua University ²Tsinghua University



Figure 1. Overview of TexVocab. Given multi-view RGB videos of one character, we construct a texture vocabulary, and create realistic animatable human avatars.

Abstract

To adequately utilize the available image evidence in multi-view video-based avatar modeling, we propose *TexVocab*, a novel avatar representation that constructs a texture vocabulary and associates body poses with texture maps for animation. Given multi-view RGB videos, our method initially back-projects all the available images in the training videos to the posed SMPL surface, producing texture maps in the SMPL UV domain. Then we construct pairs of human poses and texture maps to establish a texture vocabulary for encoding dynamic human appearances under various poses. Unlike the commonly used joint-wise manner, we further design a body-part-wise encoding strategy to learn the structural effects of the kinematic chain. Given a driving pose, we query the pose feature hierarchically by decomposing the pose vector into several body parts and interpolating the texture features for synthesizing fine-grained human dynamics. Overall, our method is able to create animatable avatars with detailed and dynamic appearances from RGB videos, and the experiments show that our method outperforms state-of-the-art approaches. The project page can be found at <https://texvocab.github.io/>.

1. Introduction

Animatable human avatar modeling attracts a lot of attention due to its enormous potential in AR/VR applications

including games, movies and holoportation. Animatable human avatars usually take the skeletal body pose as the input signal and output the pose-conditioned dynamic human appearances. However, how to effectively learn the mapping between the driving signals and dynamic appearances is still full of challenges.

Previous works [24, 71] usually directly map the pose input, e.g., the pose vectors, to the human appearances using a conditional neural radiance field (NeRF) [37] represented by an MLP. However, the pose input does not involve any information about dynamic human appearances, so it remains difficult for NeRF MLPs to regress high-fidelity dynamic details solely from the pose input. Although some works [26, 44, 64] propose to auto-decode [40] latent embeddings to encode the dynamic appearances at the input end of NeRF, they still suffer from the limited representation ability of global codes [44] or feature lines [26], resulting in blurry synthesized avatars.

On the other hand, image-based reconstruction methods like pixelNeRF [69] and SparseFusion [73] have already proved that taking pixel-aligned features as the input of NeRF can significantly improve the quality and details for static scene rendering. Inspired by these image-conditioned representations, we propose *TexVocab*, a texture vocabulary that adequately utilizes explicit image evidence to guide the implicit conditional NeRF to learn the dynamics from expressive texture conditions. To associate the multi-view images with the dynamic human body, we back-project all the available images of corresponding training poses to the

posed SMPL surface and transform them to the SMPL UV domain like NeuralActor [29], obtaining a set of texture maps. Compared with other 3D representations like feature lines [26] and feature planes [4, 5], the 2D texture maps can compactly cover the whole 2D manifold of the human body while avoiding excessive memory cost. Besides, the texture maps provide pixel-aligned features that already involve detailed human appearances so that they can serve as effective conditions to decode fine-grained details.

Our objective is to not only reconstruct all the frames in the training dataset, but also synthesize high-quality dynamic human appearances under novel poses. Therefore, it is necessary to bridge the human poses and texture maps for animation. Drawing inspiration from PoseVocab [26], we sample key human poses and associate them with corresponding texture maps. Unlike PoseVocab which constructs a separate vocabulary for each joint of the SMPL model [31], we observe that the dynamic wrinkles of clothed humans are not individually influenced by the relative rotations of a single joint. Instead, they are typically influenced by the kinematic chain of a body part, such as the entire right arm. To this end, we propose to decompose all the SMPL skeletons into several body parts, then sample key body parts and assign corresponding texture maps to them. Such a body-part-wise decomposition explicitly models the structural motion of each body part, enabling better generalization to novel poses.

In summary, our technical contributions are below:

- TexVocab, a new avatar representation that constructs a texture vocabulary to leverage expressive texture conditions for high-quality avatar modeling.
- A body-part-wise encoding method that not only disentangles the effects of different joints on the dynamic appearance, but also retains the structural effects of the kinematic chain, enabling better pose generalization.
- Experiments demonstrate that our method can create higher-fidelity avatars with dynamic pose-dependent details compared to other SOTA approaches.

2. Related Works

Implicit Neural Representations. In the last few years, implicit scene representation is becoming increasingly popular as it can produce impressive results both in geometry [7, 36, 40] and appearance [30, 37, 54, 59] modeling. Recent works show that implicit representations like occupancy networks [36], signed distance fields (SDF) [40] and neural radiance fields (NeRF) [37] lead to higher-resolution and topology-free 3D scene modeling compared with explicit ones, e.g., voxels [35, 46, 55], points [47, 48] and meshes [2, 14, 62]. In particular, NeRF [37] shows impressive rendering quality and good differentiable properties, attracting much attention in static scene rendering [5, 12, 18, 28, 38, 68]. Many other works add the

time dimension to extend NeRF to dynamic scene modeling [10, 11, 20, 41, 53, 60]. However, the motivation of these works is to reconstruct the dynamic scene, object or character in each frame of the given video without considering generating dynamic appearances under novel conditions, e.g., body poses. While our goal is to not only reconstruct the 3D human under the training poses, but also to synthesize dynamic details under unseen poses.

Geometric Avatar Modeling. Geometric based methods aim at training a pose-conditioned human model based on observations from geometric data like scans [6, 25, 33, 50] or depth images [3, 9, 63]. Scan-based methods like SCANimate [50], SCALE [32], POP [33] and SNARF [6] adopt the SMPL [31] poses or position maps as the pose conditions to learn the pose-dependent geometric surface. FITE [27] and GeoTexAvatar [25] render positional maps of posed SMPL models for the spatial continuity of the learned pose-dependent warping field. MetaAvatar [63], PINA [9] and DSFN [3] learn the pose-dependent dynamics from depth sequences. However, these methods require 3D data for training the avatar, limiting their applications to more accessible RGB videos.

RGB Video-based Avatar Modeling. On the other end of the spectrum, lots of works focus on creating animatable textured avatars from RGB videos [21, 22, 24, 26, 29, 44, 64, 71, 72]. Many approaches directly map the SMPL-derived inputs like pose vectors to the human appearances using a conditional NeRF [37] to decode the dynamic character. Specifically, TAVA [24] learns the non-rigid warping fields and shading effects directly conditioned on SMPL pose vectors. SelfRecon [21] learns the canonical template mesh and pose-dependent deformation from a self-rotating video. SLRF [71] and AvatarRex [72] sample nodes attached to SMPL, define a set of local radiance fields and learn the mapping from SMPL pose vectors to node translations and dynamic appearances. However, the pose input does not contain any information about dynamic human appearances, so it remains difficult for MLPs to predict dynamics among various poses, thus limiting the quality of avatars. AniNeRF [44], ARAH [64], NeuralBody [45] TotalSelfScan [8] and PoseVocab [26] auto-decode [40] latent embeddings to encode the dynamic appearances with per-frame latent code or joint-structured feature lines. Although adding extra embeddings variables like global codes [44] and feature lines [26] offloads the network, the low capability and poor representation ability of these embeddings still limits the avatar quality. NeuralActor [29] predicts pose-dependent texture maps from SMPL normal maps through the vid2vid model [65] using ground-truth texture maps as the monitoring signals, extracts and feeds textural feature to an MLP for decoding the dynamic human appearances. TexDVA [49] and LookingGood [34] take 2 or 3 RGB images as the driving signals for synthesizing dynamic human

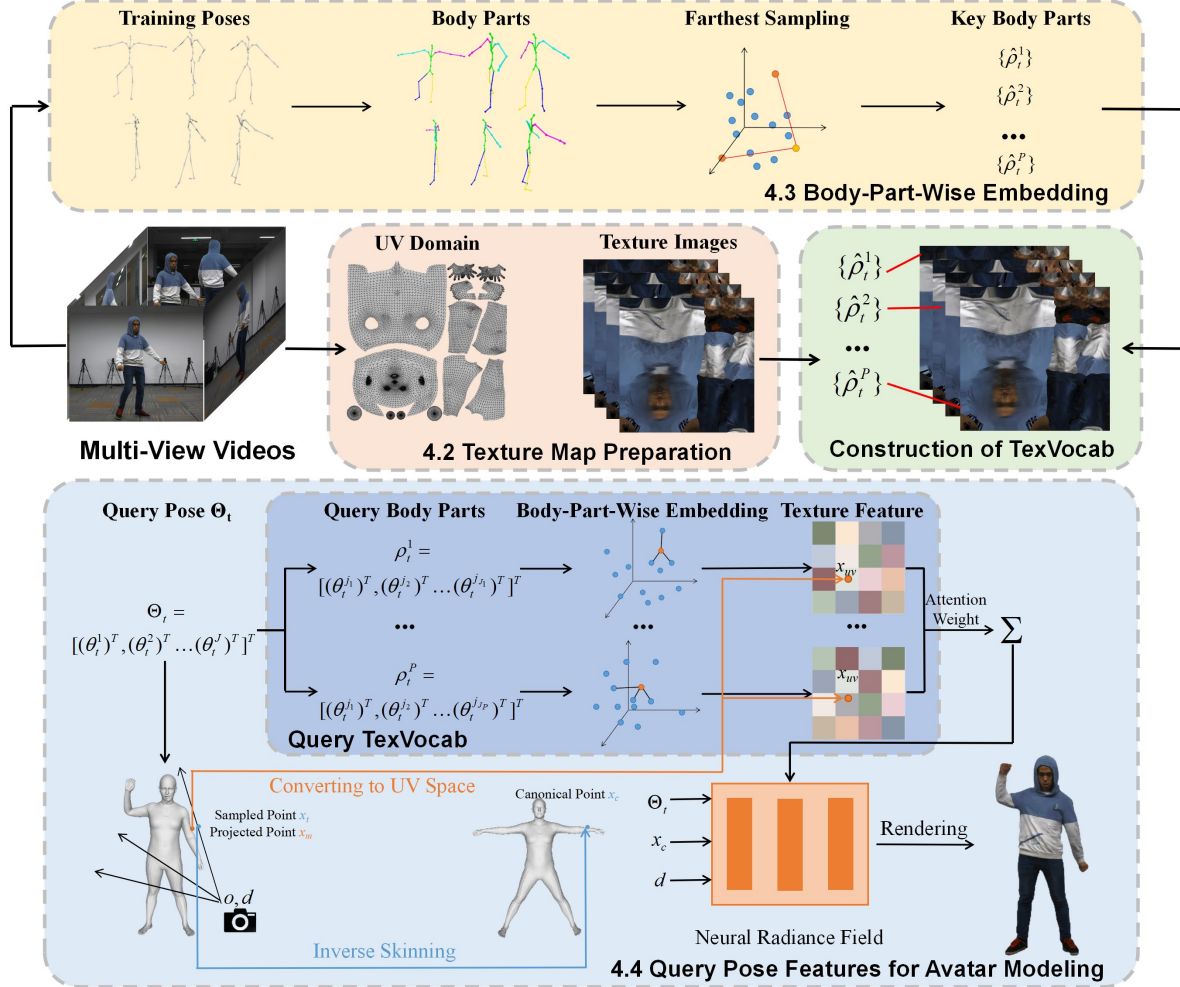


Figure 2. **Framework of TexVocab.** We first construct TexVocab by decomposing SMPL poses into body parts, sampling key body parts and gathering corresponding texture maps. Then given a query pose and a 3D coordinate, we decompose the pose into body parts, interpolate key body parts and sample texture maps as the pose conditioned feature. We finally utilize NeRF represented as an MLP to decode the dynamic character and render human appearance with detailed pose-dependent dynamics.

appearances. Stylepeople [13] uses simple neural rendering instead of NeRF which allows to render photo-realistic 2D images of individuals in baggy clothes in different poses.

Avatar Modeling with Simulation. Unlike the data-driven methods mentioned above, another line of methods use physical simulation for dynamic garment reconstruction to model the clothed human avatars. Pioneer learning-based garment simulation methods [39, 42, 51, 61] use predefined simulators and pre-generated data, while the deep learning framework does not contain any physical models. PBNS [1] and SNUG [52] incorporate physical constraints into dynamic cloth simulation and generate realistic simulation results using unsupervised training. Caphy [56] proposes to optimize the parameters of the garment physics priors to obtain better physical properties from the 3D scans instead of using the given fixed physical parameters of the

fabric. Compared with these works, our method is data-driven without the requirement of complicated simulation.

3. Preliminary

Given a set of multi-view RGB videos of a performer with T frames captured by N synchronized cameras, we aim at training a model that can output high-fidelity dynamic appearances animated by skeletal poses. We denote the RGB sequences as $\{I_t^n | 1 \leq n \leq N, 1 \leq t \leq T\}$. We assume the access to the body poses of all the frames, denoted as $\Theta_t \in \mathbb{R}^{J \times 3}, 1 \leq t \leq T$, where J denotes the number of joints of the human body. We assume that the former T_1 poses are used for training, while the rest poses are for testing. Similar to other avatar representations [44, 58], we utilize linear blend skinning (LBS) to transform sampled

points from the observation space of pose Θ_t to the canonical space. The sampled point and the canonical point can be denoted as x_t and x_c respectively. Then we represent the canonical 3D character as a pose-conditioned neural radiance field (NeRF) [37] that takes a canonical 3D coordinate x_c , a viewing direction d , and the positional conditioned pose feature $f(\Theta_t, x_c)$ as the input:

$$g(\gamma_x(x_c), \gamma_d(d), f(\Theta_t, x_c)) = (\sigma_t(x_c), c_t(x_c)) \quad (1)$$

where γ_x and γ_d is the positional encoding [57] for spatial location and viewing direction. σ_t and c_t denote the output density and color conditioned on the pose Θ_t . Following NeRF [37], we render the images using output density values and color values.

4. Method

4.1. Overview

$f(\Theta_t, x_c)$ in Eq. (1) plays an important role as it provides pose conditions for the MLP to decode the dynamic 3D character. Pioneer works like AniNeRF [44], ARAH [64] and PoseVocab [26] assign global latent codes or joint-structured feature lines as the pose conditions. However, these representations produce limited animation results because of the low representation ability of global codes and feature lines. To provide the pose-conditioned features with higher quality, we propose TexVocab, a novel method that adequately exploits explicit image evidence to guide the neural networks to learn the pose-dependent dynamic details. The framework of our approach is shown in Fig. 2. First, we prepare texture maps, decompose all the training poses and sample key body parts to construct TexVocab. Then given a SMPL pose and a 3D position, we query key body parts using K nearest neighbors (KNN), interpolate, and sample texture maps as the pose-conditioned feature according to the UV coordinate of the 3D position. Finally, we decode the dynamic 3D character using NeRF which is represented as an MLP.

4.2. Texture Map Preparation

To utilize all the image evidence more efficiently, we propose to gather all the available training views to a particular UV domain and acquire texture maps.

For each training pose Θ_t , $1 \leq t \leq T_1$, as shown in Fig. 3, we gather the available views and back-project every pixel x_p to a posed SMPL mesh $(\mathcal{V}_t, \mathcal{F})$:

$$\text{find}(u^*, v^*, f^*) \text{ s.t. } \|x_p - \mathcal{P}(\mathcal{B}_{u,v}(\mathcal{V}_{t[\mathcal{F}(f)]}))\|_2 = 0 \quad (2)$$

where $\mathcal{V}_t \in \mathbb{R}^{N_V \times 3}$ and $\mathcal{F}_t \in \mathbb{R}^{N_F \times 3}$ denotes the vertices and faces, N_V and N_F are the number of vertices and faces of the SMPL mesh respectively. $1 \leq f \leq N_F$ is the triangle index, $\mathcal{V}_{t[\mathcal{F}(f)]}$ is the three vertices of the triangle $\mathcal{F}(f)$,

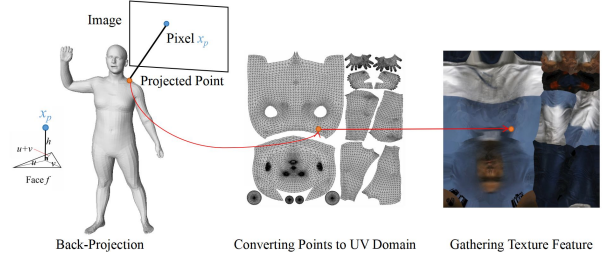


Figure 3. Overview of texture map preparation. First, we back-project all the available pixels to the posed SMPL mesh. Then we convert the projected points on the SMPL mesh to a particular UV domain. Finally, we gather and average all the available pixels, and obtain texture maps based on multi-view images.

$(u, v) : u, v, u + v \in [0, 1]$ represent the barycentric coordinates on the face, $\mathcal{B}_{u,v}(\cdot)$ is the barycentric interpolation function, and $\mathcal{P}(\cdot)$ stands for perspective projection.

Then we transfer the projected point (u^*, v^*, f^*) to a particular UV domain based on the SMPL-defined UV parameterization $\mathcal{A} \in \mathbb{R}^{N_F \times 3 \times 2}$ which maps 3D mesh surface points to a 2D UV plane. Once all the image pixels have been mapped, we average it across all the available views to gather the texture image U_t . Finally, we extract the feature map F_t of the texture image U_t using a convolutional neural network (CNN) similar to PixelNeRF [69].

4.3. Body-Part-Wise Embedding

We observe that the dynamic wrinkles of clothed humans are not influenced by the rotation of a single joint individually, but governed by a kinematic chain of a body part. To this end, we propose to divide poses into several body parts, which can not only disentangle the effects of different joints on the dynamic appearances but also retain topology information among the kinematic chains. Specifically, as shown in Table 1 and Fig. 4, we divide the SMPL pose with $J = 24$ joints into $P = 5$ body parts, and each body part contains a particular set of joints. We denote the p -th body part of the t -th frame as

$$\rho_t^p = [(\theta_t^{j_1})^T, (\theta_t^{j_2})^T, \dots, (\theta_t^{j_{J_p}})^T]^T \quad (3)$$

here $\theta_t^j \in so(3)$ denotes the j -th joint rotation of pose Θ_t , $1 \leq p \leq P$ is the rank of the body part, and J_p denote the number of joints contained in the p -th body part.

After the division, we sample key body parts and assign embeddings to them. Given the p -th body parts $\{\rho_t^p | 1 \leq t \leq T_1\}$ of the training poses, we first sample M key body parts via farthest point sampling. The distance metric between two body parts ρ_1^p, ρ_2^p is calculated as the sum of the distance [19] of each joint rotation in the p -th body part:

$$d(\rho_1^p, \rho_2^p) = \sum_{k=1}^{J_p} (1 - |q(\theta_1^{j_k})^T q(\theta_2^{j_k})|) \quad (4)$$

body part rank	Joints contained in the body part
1(main body)	0,1,2,3,6,9,12,15
2(left leg)	4,7,10
3(right leg)	5,8,11
4(left arm)	13,16,18,20,22
5(right arm)	14,17,19,21,23

Table 1. The division of the body parts. To maintain the information of the kinematic chains, we decompose SMPL skeletons into body parts instead of joints.



Figure 4. We decompose the SMPL skeletons into several body parts. Joints with the same color belong to the same body part.

where $q(\cdot)$ is a function that maps an axis-angle vector to a unit quaternion. We denote the sampled M key body parts as $\{\hat{\rho}_m^p | 1 \leq m \leq M\}$. For each sampled key body part $\hat{\rho}_t^p$, we assign the corresponding texture image feature \hat{F}_t to it. Notice that we assign the same texture map to different key body parts which belong to the same pose.

So far, we have constructed M pairs of keys and values for each body part based on the training poses. For each key body part, we can find the corresponding textural feature. These body-part-wise pose embeddings construct TexVocab and serve as discrete samples in the continuous pose feature space in the following query.

4.4. Query Pose Features for Avatar Modeling

Given a pose vector $\Theta_t = [(\theta_t^1)^T, (\theta_t^2)^T, \dots, (\theta_t^J)^T]^T$ and a 3D position x_t sampled in the observation space of pose Θ_t , we query the pose feature based on the constructed TexVocab in Sec. 4.2 and Sec. 4.3.

Query Body Parts. Given a pose Θ_t , we first decompose it to several body parts denoted as $\{\rho_t^p | 1 \leq p \leq P\}$. For the p -th query body part ρ_t^p , we search for K nearest body parts $\{\hat{\rho}_k^p | 1 \leq k \leq K\}$ according to Eq. (4). Then we interpolate the corresponding texture maps $\{\hat{F}_k | 1 \leq k \leq K\}$ according to the distance between body parts and make normalization:

$$F_{p,t} = \sum_{k=1}^K \hat{F}_k \cdot \frac{w(\rho_t^p, \hat{\rho}_k^p)}{\sum_{k=1}^K w(\rho_t^p, \hat{\rho}_k^p)} \quad (5)$$

$$w(\rho_t^p, \hat{\rho}_k^p) = J_p - d(\rho_t^p, \hat{\rho}_k^p) \quad (6)$$

where $w(\rho_t^p, \hat{\rho}_k^p)$ is the weight that varies inversely to the distance between body parts.

Texture Map Sampling. Given a 3D position x_t in the observation space of Θ_t , we project it to SMPL surface:

$$(u^*, v^*, f^*) = \arg \min_{u,v,f} \|x_p - \mathcal{B}_{u,v}(\mathcal{V}_{t[\mathcal{F}(f)]})\|_2^2 \quad (7)$$

and then we convert the projected point $x_m = (u, v, f)$ to the 2D coordinate x_{uv} in UV domain according to the fixed UV parameterization matrix \mathcal{A} defined in Sec. 4.2. For each weighted feature map $F_{p,t}$, we sample the feature $h(x_{uv}, F_{p,t})$ using bilinear interpolation. Moreover, following SCANimate [50], we also apply a skinning-weight-wise attention scheme on the feature of the p -th body part to reduce spurious correlations, which can be denoted as

$$\Omega(x_c, p) = \frac{\sum_{i=1}^{J_p} \omega(x_c, j_i)}{\sum_{j=1}^J \omega(x_c, j)} \quad (8)$$

where $\omega(x_c, j)$ is the pre-defined influence weight of the j -th joint on the corresponding canonical point x_c . Finally, we gather the sampled feature as

$$f(\Theta_t, x_c) = \sum_{p=1}^P \Omega(x_c, p) \cdot h(x_{uv}, F_{p,t}) \quad (9)$$

With the sampled pose feature $f(\Theta_t, x_c)$, we can feed it along with the pose vector Θ_t , the viewing direction d and the 3D coordinate x_c into the network described as Eq. (1) to decode NeRF that represents the dynamic 3D character.

Discussion. For the pixel x_p and the query 3D coordinate x_t , we both project them to the posed SMPL mesh and convert them to a particular UV domain. These guarantee the alignment of the sampled textural feature. The texture maps are gathered from available views, which involve detailed human appearances and can provide effective conditions to decode fine-grained details.

Also, notice that the inverse skinning often includes residuals when sampled points are not on the SMPL mesh. To further ensure the alignment of the sampled textural feature, we do not use the canonical coordinate x_c , but use x_t in the observation space of pose Θ_t instead.

4.5. Training

We use pre-trained resnet34 [16] as the backbone of the CNN that extracts the features of texture image, and the parameters are fixed during the training stage. Also, we do not regress the density value σ_t in Eq. (1) directly. Instead, we output SDF value s_t and convert it to density value following VolSDF [67]. The total loss \mathcal{L} includes the color loss, the mask loss, the eikonal loss, and the perceptual loss, which is defined as:

$$\mathcal{L} = \lambda_{color} \mathcal{L}_{color} + \lambda_{mask} \mathcal{L}_{mask} + \lambda_{perceptual} \mathcal{L}_{perceptual} + \lambda_{eikonal} \mathcal{L}_{eikonal} \quad (10)$$

where λ_{color} , $\lambda_{perceptual}$, λ_{mask} and $\lambda_{eikonal}$ stand for the loss weights. \mathcal{L}_{color} measures the MSE between the

Method	Training Poses				Novel Poses			
	PSNR \uparrow	SSIM \uparrow	LPIPS* \downarrow	FID \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS* \downarrow	FID \downarrow
TAVA [24]	25.57	0.9624	29.58	64.59	26.61	0.9597	31.14	79.96
ARAH [64]	22.78	0.9335	77.93	126.69	22.13	0.9241	93.06	113.88
AniNeRF [44]	25.19	0.9592	31.21	81.57	23.85	0.9486	32.07	95.41
PoseVocab [26]	34.06	0.9852	14.43	22.88	30.13	0.9806	16.32	28.10
Ours	36.52	0.9896	10.83	12.31	32.09	0.9832	13.40	18.79

Table 2. Quantitative comparisons against TAVA, ARAH, AniNeRF, and PoseVocab on sequence “subject00” of THuman4.0 dataset. We evaluate the numerical results of each method on both training poses and novel poses. Here LPIPS* = 1000 \times LPIPS.

Method	PSNR \uparrow	LPIPS \downarrow	FID \downarrow
NeuralActor [29]	23.531	0.066	19.714
Ours	26.325	0.017	17.836

Table 3. Quantitative comparisons against NeuralActor on sequence “S2” of “DeepCap” dataset. Results of NeuralActor are borrowed from [29].

rendered and ground-truth pixel colors, \mathcal{L}_{mask} is an MAE loss which supervises the occupancy values of the rendered pixels, and $\mathcal{L}_{eikonal}$ is the Eikonal loss encouraging the geometry fields to approximate a true signed distance function [67]. The perceptual loss $\mathcal{L}_{perceptual}$ is widely used in NeRF training, which leads to better recovery of high-frequency details like the clothed wrinkles and thin lines [70]. We choose VGGNet as the backbone to calculate the learned perceptual image patch similarity (LPIPS).

5. Experiments

5.1. Datasets and Metrics.

Datasets. We use 6 sequences of multi-view videos for experiments. 3 sequences with 24 views are from THuman4.0 dataset [71], 2 sequences with 21 or 23 views are from ZJU-MoCap dataset [44] and 1 with 11 views is from DeepCap dataset [15]. All the datasets provide parameters of cameras. Deepcap and THuman4.0 provide SMPL-X [43] registrations, and ZJU-MoCap provides SMPL [31]. We split each sequence into 2 consecutive parts for training and testing, where the training part accounts for 40% \sim 80% of the multi-view sequences, and the novel poses include the remaining part of the videos and other poses like poses from AIST++ dataset [23].

Metrics. We report quantitative results using four standard metrics: Peak Signal-to-Noise Ratio(PSNR), Structure Similarity Index Measure(SSIM) [66], Learned Perceptual Image Patch Similarity(LPIPS) [70] and Frechet Inception Distance(FID) [17].

5.2. Results and Comparisons.

We train the networks for each multi-view video sequence individually. The qualitative results shown in Fig. 1 demonstrate that our approach represents fine-grained dynamic details under various novel poses. For more novel pose synthesis results, please refer to the supplementary materials. Then we compare our method against other approaches, such as TAVA [24], ARAH [64], AniNeRF [44], PoseVocab [26] and NeuralActor [29]. We do not compare our results with TexDVA [49] because TexDVA takes 2 or 3 images to render appearances, which is different from our approach that reconstructs avatar driven by poses.

TAVA, ARAH, AniNeRF and PoseVocab. Fig. 5 and Tab. 2 show the qualitative and quantitative results against TAVA, ARAH, AniNeRF and PoseVocab on DeepCap dataset [15] and THuman4.0 dataset [71]. The results synthesized by TAVA, ARAH and AniNeRF are blurry, probably because neither pose-dependent shading in TAVA nor per-frame latent codes in ARAH and AniNeRF provide effective conditions to decode fine-grained details. Although PoseVocab outperforms the other three approaches, there are still some fine-grained details it cannot represent because of the limited capability of feature lines. In contrast, our method not only reconstructs more details in terms of the edges and fine-grained garment wrinkles under the training poses, but also renders more realistic dynamics when giving novel poses benefiting from the realistic human appearance conditions served by pixel-aligned pose features. Overall, our method outperforms the other four approaches both qualitatively and quantitatively with the construction of TexVocab which adequately utilizes the existing image evidence.

NeuralActor. NeuralActor utilizes a 2D vid2vid model [65] and tries to predict the SMPL-defined texture maps from normal maps, thus extracting features and using a neural radiance field to decode the 3D human characters. We compare our method with NeuralActor [29] qualitatively and quantitatively on “S2” sequence of DeepCap dataset in Fig. 6 and Tab. 3. We follow the same training/testing splits and metric computation as NeuralActor, and the numerical and visual results are borrowed from [29].

Method	Training Poses				Novel Poses			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow
Global Pose	33.75	0.9824	18.65	22.59	29.59	0.9731	22.70	34.26
SMPL Joint	33.43	0.9813	19.40	21.83	29.23	0.9732	23.54	37.58
Body Part (Ours)	34.05	0.9878	14.51	17.30	30.95	0.9786	19.06	28.11

Table 4. Quantitative results of ablation study on encoding strategy. We assign the texture maps to global poses, SMPL joints and body parts respectively, and report the numerical results on 3 sequences of THuman4.0 dataset [71]. Here LPIPS* = $1000 \times$ LPIPS.




Figure 5. Qualitative comparisons against TAVA, ARAH, AniNeRF and PoseVocab. We evaluate methods on THuman4.0 dataset and DeepCap dataset and show the animation results on both training poses and novel poses respectively.




Figure 6. Qualitative comparisons against NeuralActor on novel pose synthesis. Results of NeuralActor are borrowed from [29]. (a) results of NeuralActor, (b) our results, (c) ground truth.

Although the vid2vid model is powerful enough to predict texture maps solely from SMPL-derived attributes, the low frequency of normal maps still limits the performance. In contrast, directly exploiting texture maps gathered from image evidence can eliminate such problem.

5.3. Ablation Studies

In this subsection, we conduct ablation studies to demonstrate the improvement brought by our contributions.

Ablation Study on Body-Part-Wise Embedding. To prove the effectiveness of the body-part-wise embedding,




Figure 7. Qualitative results of ablation study on embedding strategy. We show synthesized images of (a)global pose, (b)joint-structured, (c)body-part-wise embedding , and (d)ground truth.

we compare it against another two embedding strategy, *i.e.*, global pose embedding and joint-structured embedding. Specifically, we sample key global poses, key SMPL joints and key body parts and assign texture maps to them, respectively. Fig. 7 and Tab. 4 show the qualitative and quantitative results on 3 sequences of THUman4.0 dataset [71]. Although all the three methods can reconstruct details under training poses benefiting from the pixel-aligned features given by texture maps, global pose embedding cannot disentangle the effects of SMPL joints, and ends up with poor generation. While joint-structured embedding encodes features per joint and does not retain any information on kinematic chains, which makes it too local to generalize to challenging poses, so the wrinkles are often messy. In contrast, the proposed body-part-aware embedding not only encodes fine-grained details, but also displays realistic results by disentangling the effects of different joints and maintaining information along the kinematic chains.

Ablation Study on Texture Map. To prove the effectiveness of texture maps obtained from existing views, we take SMPL normal map as input, which is similar to NeuralActor [29]. Then we use an image generator to generate texture maps from normal maps, and take ground-truth texture maps as monitoring signals. Fig. 8 shows the qualitative results of novel pose synthesis on THUman4.0 dataset [71]. Although the 2D convolutional network is powerful enough to generate realistic texture maps




Figure 8. Qualitative results on novel pose synthesis of ablation study on texture maps.(a) Texture maps from an image generator, (b) texture maps from multi-view images, (c)ground truth.

from SMPL-derived normal maps, it still suffers from the low representation ability of driving signals. While our method can not only reconstruct fine-grained details but also achieve pose generalization with the guidance of pixel-aligned features.

6. Discussions

Conclusion. We present TexVocab, a texture vocabulary that adequately utilizes explicit image evidence to guide the implicit NeRF to learn the dynamic details from expressive texture conditions. To further represent the dynamic wrinkles of clothed humans, we propose body-part-wise embedding to decompose all the SMPL skeletons into several body parts and assign texture maps to them, which both disentangles the effects of SMPL joints and maintains information of kinematic chains. The proposed approach not only reconstructs fine-grained details in terms of garment wrinkles and edges, but also achieves pose generalization that displays realistic dynamic appearances under novel poses. Experiments on different multi-view video datasets indicate that our approach outperforms other state-of-the-art methods both qualitatively and quantitatively, showing its enormous potential in different kinds of interactive applications.

Limitation. Since our avatar representation relies on the inverse skinning by SMPL skeletons, it cannot handle loose clothes like long dresses which do not follow the topological structure of the SMPL model. Also, the gathering of texture maps relies on dense views. For monocular or sparse view datasets, we cannot gather complete texture images to guide avatar construction.

Potential Social Impact. Since our method enables automatic creation of animatable human avatars, it can be misused to re-target individuals with actions they do not perform. To prevent the risks, it is critical to evaluate the caution before developing such kind of technology.

Acknowledgment. This research was funded through National Key Research and Development Program of China (Project No. 2022YFB36066). The project was also supported by the NSFC project No. 62125107.

References

- [1] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Pbnns: physically based neural simulator for unsupervised garment pose space deformation. *arXiv preprint arXiv:2012.11310*, 2020. [3](#)
- [2] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017. [2](#)
- [3] Andrei Burov, Matthias Nießner, and Justus Thies. Dynamic surface function networks for clothed human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10754–10764, 2021. [2](#)
- [4] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16123–16133, 2022. [2](#)
- [5] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. [2](#)
- [6] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11594–11604, 2021. [2](#)
- [7] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. [2](#)
- [8] Junting Dong, Qi Fang, Yudong Guo, Sida Peng, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Totalseffscan: Learning full-body avatars from self-portrait videos of faces, hands, and bodies. *Advances in Neural Information Processing Systems*, 35:13654–13667, 2022. [2](#)
- [9] Zijian Dong, Chen Guo, Jie Song, Xu Chen, Andreas Geiger, and Otmar Hilliges. Pina: Learning a personalized implicit neural avatar from a single rgb-d video sequence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20470–20480, 2022. [2](#)
- [10] Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8649–8658, 2021. [2](#)
- [11] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5712–5721, 2021. [2](#)
- [12] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14346–14355, 2021. [2](#)
- [13] Artur Grigorev, Karim Iskakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. Stylepeople: A generative model of fullbody human avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5151–5160, 2021. [3](#)
- [14] Kan Guo, Dongqing Zou, and Xiaowu Chen. 3d mesh labeling via deep convolutional neural networks. *ACM Transactions on Graphics (TOG)*, 35(1):1–12, 2015. [2](#)
- [15] Marc Habermann, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular human performance capture using weak supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5052–5063, 2020. [6](#)
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [5](#)
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. [6](#)
- [18] Wenbo Hu, Yuling Wang, Lin Ma, Bangbang Yang, Lin Gao, Xiao Liu, and Yuewen Ma. Tri-miprf: Tri-mip representation for efficient anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19774–19783, 2023. [2](#)
- [19] Du Q Huynh. Metrics for 3d rotations: Comparison and analysis. *Journal of Mathematical Imaging and Vision*, 35:155–164, 2009. [4](#)
- [20] Mustafa Işık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Nießner. Humanrf: High-fidelity neural radiance fields for humans in motion. *arXiv preprint arXiv:2305.06356*, 2023. [2](#)
- [21] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. Selfrecon: Self reconstruction your digital avatar from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5605–5615, 2022. [2](#)
- [22] Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16922–16932, 2023. [2](#)
- [23] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021. [6](#)
- [24] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. In *European Conference on Computer Vision*, pages 419–436. Springer, 2022. [1](#), [2](#), [6](#)
- [25] Zhe Li, Zerong Zheng, Hongwen Zhang, Chaonan Ji, and Yebin Liu. Avatarecap: Animatable avatar conditioned

- monocular human volumetric capture. In *European Conference on Computer Vision*, pages 322–341. Springer, 2022. [2](#)
- [26] Zhe Li, Zerong Zheng, Yuxiao Liu, Boyao Zhou, and Yebin Liu. Posevocab: Learning joint-structured pose embeddings for human avatar modeling. *arXiv preprint arXiv:2304.13006*, 2023. [1](#), [2](#), [4](#), [6](#)
- [27] Siyou Lin, Hongwen Zhang, Zerong Zheng, Ruizhi Shao, and Yebin Liu. Learning implicit templates for point-based clothed human modeling. In *European Conference on Computer Vision*, pages 210–228. Springer, 2022. [2](#)
- [28] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. [2](#)
- [29] 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 transactions on graphics (TOG)*, 40(6):1–16, 2021. [2](#), [6](#), [7](#), [8](#)
- [30] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. [2](#)
- [31] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: A skinned multi-person linear model. *Acm Transactions on Graphics*, 34(6cd):248, 2015. [2](#), [6](#)
- [32] Qianli Ma, Shunsuke Saito, Jinlong Yang, Siyu Tang, and Michael J Black. Scale: Modeling clothed humans with a surface codec of articulated local elements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16082–16093, 2021. [2](#)
- [33] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J Black. The power of points for modeling humans in clothing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10974–10984, 2021. [2](#)
- [34] Ricardo Martin-Brualla, Rohit Pandey, Shuoran Yang, Pavel Pidlypenskyi, Jonathan Taylor, Julien Valentin, Sameh Khamis, Philip Davidson, Anastasia Tkach, Peter Lincoln, et al. Lookingood: Enhancing performance capture with real-time neural re-rendering. *arXiv preprint arXiv:1811.05029*, 2018. [2](#)
- [35] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 922–928. IEEE, 2015. [2](#)
- [36] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019. [2](#)
- [37] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [1](#), [2](#), [4](#)
- [38] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. [2](#)
- [39] Xiaoyu Pan, Jiaming Mai, Xinwei Jiang, Dongxue Tang, Jingxiang Li, Tianjia Shao, Kun Zhou, Xiaogang Jin, and Dinesh Manocha. Predicting loose-fitting garment deformations using bone-driven motion networks. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. [3](#)
- [40] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. [1](#), [2](#)
- [41] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. [2](#)
- [42] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7365–7375, 2020. [3](#)
- [43] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10975–10985, 2019. [6](#)
- [44] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14314–14323, 2021. [1](#), [2](#), [3](#), [4](#), [6](#)
- [45] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021. [2](#)
- [46] Charles R Qi, Hao Su, Matthias Nießner, Angela Dai, Mengyuan Yan, and Leonidas J Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5648–5656, 2016. [2](#)
- [47] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. [2](#)
- [48] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. [2](#)

- [49] Edoardo Remelli, Timur Bagautdinov, Shunsuke Saito, Chenglei Wu, Tomas Simon, Shih-En Wei, Kaiwen Guo, Zhe Cao, Fabian Prada, Jason Saragih, et al. Drivable volumetric avatars using texel-aligned features. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–9, 2022. 2, 6
- [50] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2886–2897, 2021. 2, 5
- [51] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Learning-based animation of clothing for virtual try-on. In *Computer Graphics Forum*, pages 355–366. Wiley Online Library, 2019. 3
- [52] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Snug: Self-supervised neural dynamic garments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8140–8150, 2022. 3
- [53] Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16632–16642, 2023. 2
- [54] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [55] Shuran Song and Jianxiong Xiao. Deep sliding shapes for amodal 3d object detection in rgb-d images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 808–816, 2016. 2
- [56] Zhaoqi Su, Liangxiao Hu, Siyou Lin, Hongwen Zhang, Shengping Zhang, Justus Thies, and Yebin Liu. Caphy: Capturing physical properties for animatable human avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14150–14160, 2023. 3
- [57] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. 4
- [58] Gusi Te, Xiu Li, Xiao Li, Jinglu Wang, Wei Hu, and Yan Lu. Neural capture of animatable 3d human from monocular video. In *European Conference on Computer Vision*, pages 275–291. Springer, 2022. 3
- [59] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *Acm Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 2
- [60] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12959–12970, 2021. 2
- [61] Raquel Vidas, Igor Santesteban, Elena Garces, and Dan Casas. Fully convolutional graph neural networks for parametric virtual try-on. In *Computer Graphics Forum*, pages 145–156. Wiley Online Library, 2020. 3
- [62] Pengyu Wang, Yuan Gan, Panpan Shui, Fenggen Yu, Yan Zhang, Songle Chen, and Zhengxing Sun. 3d shape segmentation via shape fully convolutional networks. *Computers & Graphics*, 76:182–192, 2018. 2
- [63] Shaofei Wang, Marko Mihajlovic, Qianli Ma, Andreas Geiger, and Siyu Tang. Metaavatar: Learning animatable clothed human models from few depth images. *Advances in Neural Information Processing Systems*, 34:2810–2822, 2021. 2
- [64] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdfs. In *European conference on computer vision*, pages 1–19. Springer, 2022. 1, 2, 4, 6
- [65] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. *arXiv preprint arXiv:1808.06601*, 2018. 2, 6
- [66] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [67] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 5, 6
- [68] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021. 2
- [69] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 1, 4
- [70] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [71] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15893–15903, 2022. 1, 2, 6, 7, 8
- [72] Zerong Zheng, Xiaochen Zhao, Hongwen Zhang, Boning Liu, and Yebin Liu. Avatarrex: Real-time expressive full-body avatars. *arXiv preprint arXiv:2305.04789*, 2023. 2
- [73] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12588–12597, 2023. 1