

Learning Motion-Dependent Appearance for High-Fidelity Rendering of Dynamic Humans from a Single Camera

Jae Shin Yoon^{†,‡} Duygu Ceylan[‡] Tuanfeng Y. Wang[‡]
 Jingwan Lu[‡] Jimei Yang[‡] Zhixin Shu[‡] Hyun Soo Park[†]

[†]University of Minnesota [‡]Adobe Research

Abstract

Appearance of dressed humans undergoes a complex geometric transformation induced not only by the static pose but also by its dynamics, i.e., there exists a number of cloth geometric configurations given a pose depending on the way it has moved. Such appearance modeling conditioned on motion has been largely neglected in existing human rendering methods, resulting in rendering of physically implausible motion. A key challenge of learning the dynamics of the appearance lies in the requirement of a prohibitively large amount of observations. In this paper, we present a compact motion representation by enforcing equivariance—a representation is expected to be transformed in the way that the pose is transformed. We model an equivariant encoder that can generate the generalizable representation from the spatial and temporal derivatives of the 3D body surface. This learned representation is decoded by a compositional multi-task decoder that renders high fidelity time-varying appearance. Our experiments show that our method can generate a temporally coherent video of dynamic humans for unseen body poses and novel views given a single view video.

1. Introduction

We express ourselves by moving our body that drives a sequence of natural secondary motion, e.g., dynamic movement of dress induced by dancing as shown in Figure 1. This secondary motion is the resultant of complex physical interactions with the body, which is, in general, *time-varying*. This presents a major challenge for plausible rendering of dynamic dressed humans in applications such as video based retargetting or social presence. Many existing approaches such as pose-guided person image generation [7] focus on static poses as a conditional variable. Despite its promising rendering quality, it fails to generate a physically plausible secondary motion, e.g., generating the same appearance for fast and slow motions.

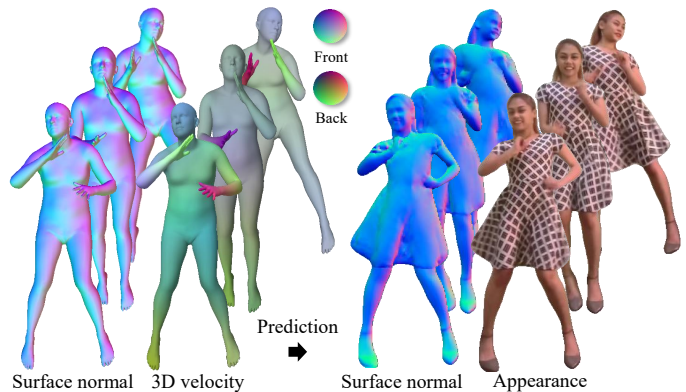


Figure 1. Given surface normal and velocity of a 3D body model, our method synthesizes subject-specific surface normal and appearance. We specifically focus on synthesis of plausible dynamic appearance by learning an effective 3D motion descriptor.

One can learn the dynamics of the secondary motion from videos. This, however, requires a tremendous amount of data, i.e., videos depicting all possible poses and associated motions. In practice, only a short video clip is available, e.g., the maximum length of videos in social media (e.g., TikTok) are limited to 15-60 seconds. The learned representation is, therefore, prone to overfitting.

In this paper, we address the fundamental question of “can we learn a representation for dynamics given a limited amount of observations?”. We argue that a meaningful representation can be learned by enforcing an equivariant property—a representation is expected to be transformed in the way that the body pose is transformed. With the equivariance, we model the dynamics of the secondary motion as a function of spatial and time derivative of the 3D body. We construct this representation by re-arranging 3D features in the canonical coordinate system of the body surface, i.e., the UV map, which is invariant to the choice of the 3D coordinate system. The UV map also captures the semantic meaning of body parts since each body part is represented by a UV patch. The resulting representation is compact and

discriminative compared to the 2D pose representations that often suffer from geometric ambiguity due to 2D projection.

We observe that two dominant factors significantly impact the physicality of the generated appearance. First, the silhouette of dressed humans is transformed according to the body movement and the physical properties (e.g., material) of the individual garment types (e.g., top and bottom garments might undergo different deformations). Second, the local geometry of the body and clothes is highly correlated, e.g., surface normals of T-shirt and body surface, which causes appearance and disappearance of folds and wrinkles. To incorporate these factors, we propose a compositional decoder that breaks down the final appearance rendering into modular subtasks. This decoder predicts the time-varying semantic maps and surface normals as intermediate representations. While the semantic maps capture the time-varying silhouette deformations, the surface normals are effective in synthesizing high quality textures, which further enables re-lighting. We combine these intermediate representations to produce the final appearance.

Our experiments show that our method can generate a temporally coherent video of an unseen secondary motion from novel views given a single view training video. We conduct thorough comparisons with various state-of-the-art baseline approaches. Thanks to the discriminative power, our representation demonstrates superior generalization ability, consistently outperforming previous methods when trained on shorter training videos. Furthermore, our method shows better performance in handling complex motion sequences including 3D rotations as well as rendering consistent views in applications such as free-viewpoint rendering. The intermediate representations predicted by our method such as surface normals also enable applications such as relighting which are otherwise not applicable.

2. Related Work

Two major rendering approaches have been used for high quality human synthesis: model based and retrieval based approaches. Model based approaches leverage the 3D shape models, e.g., 3DMM face model [4], that can synthesize a novel view based on the geometric transformation [11, 57]. Retrieval approaches synthesize an image by finding matches in both local and global shape and appearance [6, 31]. These are, now, combined with a deep representation to form neural rendering and generative networks.

Neural Rendering. As human appearances are modulated by their poses, it is possible to generate the high fidelity appearance by using a parametric 3D body model, e.g., deformable template models [35]. For example, some existing approaches [44, 67] learned a constant appearance of a person by mapping from per-vertex RGB colors defined on the SMPL body to synthesized images. Textured neural avatars [51] learned a person-specific texture map by pro-

jecting the image features to a body surface coordinate (invariant to poses) to model human appearances. These approaches, however, are limited to statics, i.e., the generated appearance is completely blind to pose and motion.

To model pose-dependent appearances, Liu et al. [34] implicitly learned the texture variation over poses, which allowed them to refine the initial appearance obtained from texture map through a template model. On the other hand, Raj et al. [47] explicitly learned pose-dependent neural textures. To further enhance the quality of rendering, person-specific template models [2] were used by incorporating additional meshes representing the garments [65]. However, none of these approaches are capable of modeling the time-varying secondary motion. Habermann et al. [19] utilized a motion cue to model the motion-dependent appearances while requiring a pre-learned person-specific 3D template model. Zhang et al. [68] proposed a neural rendering approach to synthesize the dynamic appearance of loose garments assuming a coarse 3D garment proxy is provided. In contrast, our method uses the 3D body prior to model the dynamic appearance of both tight and loose garments.

The requirement of the parametric model can be relaxed by leveraging flexible neural rendering fields. For instance, neural volumetric representations [39] have been used to model general dynamic scenes [15, 32, 56, 58, 64] and humans [41, 42] using deformation fields. Nonetheless, the range of generated motion is still limited. Recent methods learned the appearance of a person in the canonical space of the coarse 3D body template and employ skinning and volume rendering approaches to synthesize images [8, 43, 46]. Liu et al. [33] extended such approaches by introducing pose-dependent texture maps to model pose-dependent appearance. Modeling time-varying appearance induced by the secondary motion with such volumetric approaches is still the uncharted area of study.

Generative Networks. Generative adversarial learning enforces a generator to synthesize photorealistic images that are almost indistinguishable from the real images. For example, image-to-image translation can synthesize pose-conditioned appearance of a person by using various pose representations such as 2D keypoints [13, 36, 45, 49, 52, 55, 70], semantic labels [3, 12, 14, 20, 38, 54, 66], or dense surface coordinate parametrizations [1, 17, 40, 50]. Despite remarkable fidelity, these were built upon the 2D synthesis of static images at every frame, in general, failing to generate physically plausible secondary motion. To address this challenge, several works have utilized temporal cues either in training time [7] to enable temporally smooth results or as input signals [59, 60] to model motion-dependent appearances. Kappel et al. [25] modeled the pose-dependent appearances of loose garments conditioned on 2D keypoints by learning temporal coherence using a recurrent network. However, due to the nature of 2D pose representations, the

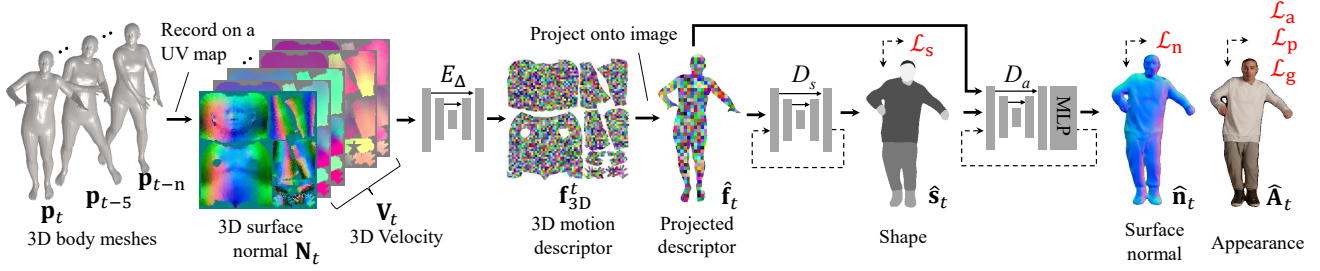


Figure 2. The overview of our human rendering pipeline. Given a set of time-varying 3D body meshes $\{\mathbf{P}_t, \dots, \mathbf{P}_{t-n}\}$ obtained from a monocular input video, we aim to synthesize high-fidelity appearance of a dressed human. We learn an effective 3D body pose and motion representation by recording the surface normal \mathbf{N}^t of the posed 3D mesh at time t and the body surface velocity \mathbf{V}^t over several past times in the spatially aligned UV space. We define an encoder E_Δ which is designed to reconstruct 3D motion descriptors \mathbf{f}_{3D}^t that encode the spatial and temporal relation of the 3D body meshes. Given a target 3D body configuration, we project \mathbf{f}_{3D}^t onto the image space which are then utilized by our compositional networks (D_s and D_a) to predict a shape with semantic labels, surface normal, and final appearance.

physics of the generated motion is limited, e.g., our experiments show that the method works well mostly for planar motions but is limited in expressing 3D rotations. Wang et al. [62] is the closest to our work, which maps a sequence of dense surface parameterization to motion features that are used to synthesize dynamic appearances using StyleGAN [26]. In contrast, our method is built on a new 3D motion representation which shows superior discriminative power, consistently outperforming [62] in terms of generalizing to unseen poses as demonstrated in our experiments.

3. Method

Given a monocular video of a person in motion and the corresponding 3D body fit estimates, we learn a motion representation to describe the time-varying appearance of the secondary motion induced by body movement (Section 3.1). We propose a multitask compositional renderer (Section 3.2) that uses this representation to render the subject-specific final appearance of moving dressed humans. Our renderer first predicts two intermediate representations including time-varying semantic maps that capture garment specific silhouette deformations and surface normals that capture the local geometric changes such as folds and wrinkles. These intermediate representations are combined to synthesize the final appearance. We obtain the 3D body fits estimates from the input video using a new model-based tracking optimization (Supplementary material). The overview of our rendering framework is shown in Figure 2.

3.1. Equivariant 3D Motion Descriptor

We cast the problem of human rendering as learning a representation via a feature encoder-decoder framework:

$$\mathbf{f} = E(\mathbf{p}), \quad \mathbf{A} = D(\mathbf{f}), \quad (1)$$

where an encoder E takes as an input a representation of a posed body, \mathbf{p} (e.g., 2D sparse or dense keypoints or 3D

body surface vertices), and outputs per-pixel features \mathbf{f} that can be used by the decoder D to reconstruct the appearance $\mathbf{A} \in [0, 1]^{w \times h \times 3}$ of the corresponding pose where w and h are the width and height of the output image (appearance). We first discuss how E can be modeled to render static appearance, then introduce our 3D motion descriptor to render time-varying appearance with secondary motion effects.

Learning a representation from Equation (1) from a limited amount of data is challenging because both encoder and decoder need to memorize every appearance in relation to the corresponding pose, $\mathbf{A} \leftrightarrow \mathbf{p}$. To address the data challenge, one can use an *equivariant* geometric transformation, \mathcal{W} , such that a feature is expected to be transformed in the way that the body pose is transformed:

$$E(\mathcal{W}\mathbf{x}) = \mathcal{W}E(\mathbf{x}). \quad (2)$$

where \mathbf{x} is an arbitrary pose. A naive encoder that satisfies this equivariance learns a constant feature \mathbf{f}_0 by warping any \mathbf{p} to a neutral pose \mathbf{p}_0 :

$$\mathbf{f}_0 = E(\mathcal{W}^{-1}\mathbf{p}) = \text{const.}, \quad \mathbf{A} = D(\mathcal{W}\mathbf{f}_0), \quad (3)$$

where $\mathbf{p} = \mathcal{W}\mathbf{p}_0$. Figure 3(a) and (b) illustrate cases where \mathcal{W} is defined as image warping in 2D or skinning in 3D respectively. \mathbf{f} can be derived by warping \mathbf{p} to the T-pose, $\mathcal{W}^{-1}\mathbf{p}$ of which feature can be warped back to the posed feature before decoding, $D(\mathcal{W}E(\mathcal{W}^{-1}\mathbf{p}))$. Since \mathbf{f}_0 is constant, the encoder E does not need to be learned. One can only learn the decoder D to render a static appearance.

To model the time-varying appearance for the secondary motion that depends on both body pose and motion, one can extend Equation (3) to encode the spatial and temporal gradients as a residual feature encoding:

$$\begin{aligned} \mathbf{f} &= E\left(\mathbf{p}, \frac{\partial \mathbf{p}}{\partial x}, \frac{\partial \mathbf{p}}{\partial t}\right) \approx E(\mathbf{p}) + E_\Delta\left(\frac{\partial \mathbf{p}}{\partial x}, \frac{\partial \mathbf{p}}{\partial t}\right) \\ \iff \mathbf{f}_0 &= \underbrace{E(\mathcal{W}^{-1}\mathbf{p})}_{\text{const.}} + E_\Delta\left(\mathcal{W}^{-1}\frac{\partial \mathbf{p}}{\partial x}, \mathcal{W}^{-1}\frac{\partial \mathbf{p}}{\partial t}\right), \quad (4) \end{aligned}$$




Figure 3. We apply equivariance to learn a compact representation. (a) In 2D, the feature $\mathbf{f} = E(\mathbf{p})$ is expected to be transformed to the feature of the neutral pose, $\mathbf{f}_0 = E(\mathcal{W}^{-1}\mathbf{p})$ by a coordinate transform \mathcal{W} , e.g., image warping. This eliminates the necessity of learning the encoder E , i.e., the appearance of the pose \mathbf{p} is generated by warping the appearance of the neutral pose. (b) Equivariance in 3D can be applied by incorporating 3D body reconstruction Π^{-1} where the feature is expected to be transformed by the 3D warping \mathcal{W} , e.g., skinning. (c) We use the canonical body surface coordinate (UV coordinate) to represent the feature coordinate transformation.

where $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial t}$ are the spatial and temporal derivatives of the posed body, respectively. The spatial derivatives essentially represent the pose corrective deformations [30, 35]. The temporal derivatives denote the body surface velocity which results in secondary motion. Since these spatial and temporal gradients are no longer constant, one needs to learn an encoder E_Δ to encode the residual features.

In this paper, we use a 3D representation of the posed body lifted from an image by leveraging recent success in single view pose reconstruction [9, 24, 28]. Hence, spatial and temporal derivatives of the body pose correspond to the surface normals and body surface velocities, respectively:

$$\mathbf{f}_{3D} = E_\Delta(\mathcal{W}^{-1}\mathbf{N}, \mathcal{W}^{-1}\mathbf{V}), \quad \mathbf{A} = D(\Pi\mathcal{W}\mathbf{f}_{3D}), \quad (5)$$

where $\mathbf{N} = \frac{\partial \mathbf{p}}{\partial X} \in \mathbb{R}^{m \times 3}$ is the 3D surface normal, and $\mathbf{V} = \frac{\partial \mathbf{p}}{\partial t} \in \mathbb{R}^{m \times 3}$ represents the instantaneous velocities of the m vertices in the body surface. We model the geometric transformation function \mathcal{W} to warp an arbitrary 3D pose \mathbf{p} to a canonical representation, \mathbf{p}_0 . We record \mathbf{f}_{3D} in a spatially aligned 2D positional map, specifically the UV map of the 3D body mesh where each pixel contains the 3D information of a unique point on the body mesh surface. This enables us to leverage 2D CNNs to apply local convolution operations to capture the relationship between neighboring body parts [37]. Therefore, $\mathbf{f}_{3D} \in \mathbb{R}^{m \times d}$ called *3D motion descriptor* is the feature defined in the UV coordinates where d is the dimension of the per-vertex 3D feature. $\mathbf{f} = \Pi\mathcal{W}\mathbf{f}_{3D}$ is the projected 3D feature in the image coordinates where Π is a coordinate transformation that transports the features defined in the UV space to the image plane via the dense UV coordinates of the body mesh.

The key advantage of the 3D motion descriptor over commonly used 2D sparse [25] or dense [62] keypoint representations is discriminativity. Consider a toy example of a person rotating his body left to right multiple times. Given

one cycle (i.e., $0-T$) of such a motion as input (Figure 4(a)), assume we want to synthesize the appearance of the person performing the repetitions of the same motion (i.e., cycles $T-5T$). As a proof of concept, we model D using a nearest neighbor classifier to retrieve the relevant image patches (top two patches) for each body part from the reference motion based on the correlation of the motion descriptors as shown in Figure 4(c). Due to the inherent depth ambiguity, multiple 3D motion trajectories yield the same 2D projected trajectory [23]. Hence, the 2D motion descriptors using sparse (Figure 4(b), top) and dense (Figure 4(b), middle) 2D keypoints confuse the directions of out-of-plane body rotation, resulting in erroneous nearest neighbor retrievals as shown in Figure 4(d). Furthermore, 2D representations entangle the viewpoint and pose into a common feature. This not only avoids a compact representation (e.g., the same body motion maps to different 2D trajectories with respect to different viewpoints and yields different features) but also suffers from occlusions in the image space. In our example in Figure 4(c), top and bottom, the upper arm is occluded in portions of the input video denoted with the purple blocks, hence no reliable local motion descriptors can be computed at those time instances. In contrast, our 3D motion descriptor is highly discriminative, which does not confuse the direction of body rotations, resulting in accurate image patch retrievals.

3.2. Multitask Compositional Decoder

Given 3D motion features, the decoder D still needs to learn to generate diverse and plausible secondary motion, which is prone to overfitting given a limited amount of data. We integrate the following properties that can mitigate this challenge. (1) Composition: we design the decoder using a composition of modular functions where each modular function is learned to generate physically and semantically meaningful intermediate representations. Learning each modular function is more accurate than learning

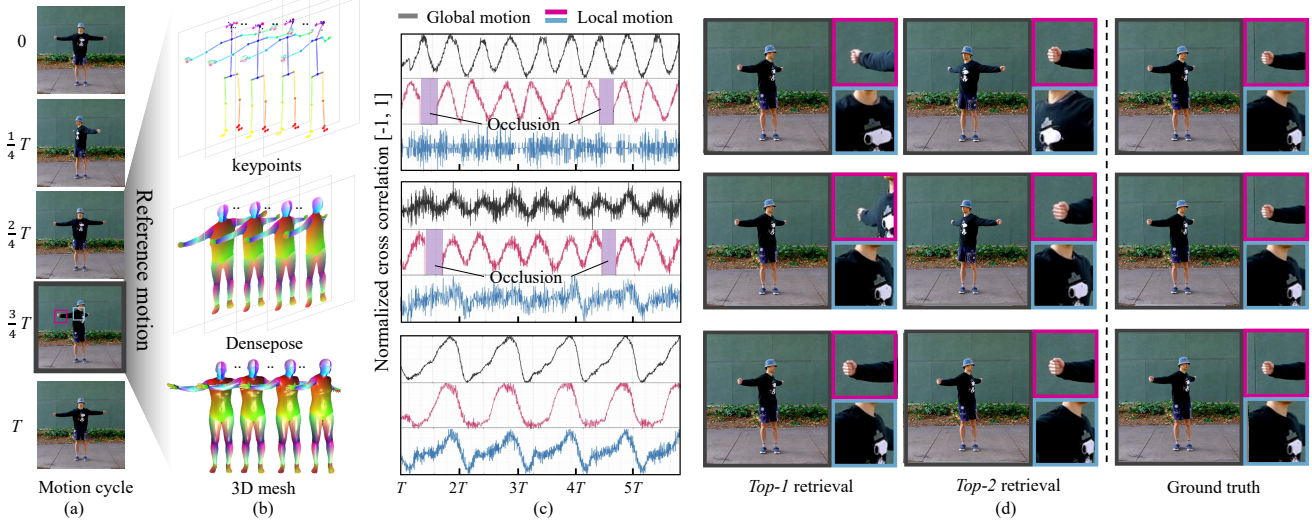


Figure 4. We show the strength of our 3D motion descriptor using a toy example. Given a video of a person rotating his body from left to right multiple times, we associate the first cycle of the motion (*i.e.*, $0 \sim T$) to the remaining cycles ($T \sim 6T$). As a proof-of-concept, we use a nearest neighbor classifier to model D . (b) We represent the motion descriptor using (top) 2D keypoints [5], (middle) 2D dense UV coordinates [18], (bottom) and 3D body mesh [35]. (c) We measure the similarity in motion descriptor for entire body (gray), local hand (pink) and upper torso (blue) using normalized cross correlation (NCC) where multiple peaks within a cycle indicate ambiguity of the descriptor. (d) Given the motion descriptors, we retrieve relevant image patches. While the 3D motion descriptors identify the image patches similar to the ground truth, due to the depth ambiguity, the 2D motion descriptors result in ambiguous matches. Furthermore, the 2D motion descriptors are not well defined in case of occlusions.

an end-to-end decoder as a whole as shown in our ablation study (Section 4.1); (2) Multi-task: each intermediate representation receives its own supervision signals resulting in multi-task learning. The motion features, \mathbf{f}_{3D} , are shared by all intermediate modules resulting in a compact representation; (3) Recurrence: each module is modeled as an autoregressive network, which allows learning the dynamics rather than memorizing the pose-specific appearance.

Our decoder is a composition of two modular functions:

$$D = D_a \circ D_s, \quad (6)$$

where D_s and D_a are the functions that generate the shape with semantic maps and the appearance.

D_s learns the dynamics of the 2D shape:

$$\hat{\mathbf{s}}_t = D_s(\hat{\mathbf{s}}_{t-1}; \hat{\mathbf{f}}_t) \quad (7)$$

where $\hat{\mathbf{f}}_t = \Pi \mathcal{W}_t \mathbf{f}_{3D}^t$ is the projected features onto the image at time t , and $\hat{\mathbf{s}}_t \in \{0, \dots, L\}^{w \times h}$ is the predicted shape with semantics where L is the number of semantic categories. In our experiments we set $L = 7$ (background, top clothing, bottom clothing, face, hair, skin, shoes).

D_a learns the dynamics of appearance given the shape and 3D motion descriptor:

$$\hat{\mathbf{A}}_t, \hat{\mathbf{n}}_t = D_a(\hat{\mathbf{A}}_{t-1}, \hat{\mathbf{n}}_{t-1}; \hat{\mathbf{s}}_t, \hat{\mathbf{f}}_t), \quad (8)$$

where $\hat{\mathbf{A}}_t \in \mathbb{R}^{w \times h \times 3}$ and $\hat{\mathbf{n}}_t \in \mathbb{R}^{w \times h \times 3}$ are the generated appearance and surface normals at time t .

We learn the 3D motion descriptor as well as the modular decoder functions by minimizing the following loss:

$$\mathcal{L} = \sum_{\mathbf{P}, \mathbf{A} \in \mathcal{D}} \mathcal{L}_a + \lambda_s \mathcal{L}_s + \lambda_n \mathcal{L}_n + \lambda_p \mathcal{L}_p + \lambda_g \mathcal{L}_g, \quad (9)$$

where \mathcal{L}_a , \mathcal{L}_s , \mathcal{L}_n , \mathcal{L}_p , \mathcal{L}_g are the appearance, shape, surface normal, perceptual similarity, and generative adversarial losses, and λ_s , λ_n , λ_p , and λ_g are their weights, respectively. We set $\lambda_s = 10$, $\lambda_n = \lambda_p = 1$ and $\lambda_g = 0.01$ in our experiments. \mathcal{D} is the training dataset composed of the ground truth 3D pose \mathbf{P} and its appearance \mathbf{A} .

$$\begin{aligned} \mathcal{L}_a(\mathbf{P}, \mathbf{A}) &= \|\hat{\mathbf{A}} - \mathbf{A}\|, \\ \mathcal{L}_s(\mathbf{P}, \mathbf{A}) &= \|\hat{\mathbf{s}} - S(\mathbf{A})\|, \\ \mathcal{L}_n(\mathbf{P}, \mathbf{A}) &= \|\hat{\mathbf{n}} - N(\mathbf{A})\|, \\ \mathcal{L}_p(\mathbf{P}, \mathbf{A}) &= \sum_i \|VGG_i(\hat{\mathbf{A}}) - VGG_i(\mathbf{A})\|, \\ \mathcal{L}_g(\mathbf{P}, \mathbf{A}) &= \mathbb{E}_{S(\mathbf{A}), \mathbf{A}} [\log(D^*(S(\mathbf{A}), \mathbf{A}))] + \\ &\quad \mathbb{E}_{S(\mathbf{A}), \hat{\mathbf{A}}} [\log(1 - D^*(S(\mathbf{A}), \hat{\mathbf{A}}))], \end{aligned}$$

where $\hat{\mathbf{A}}$, $\hat{\mathbf{s}}$, and $\hat{\mathbf{n}}$ are the generated appearance, shape, and surface normal, respectively. S and N are the shape [16] and surface normal estimates [22], and VGG is the feature extractor that computes perceptual features from conv- i -2 layers in VGG-16 networks [53], D^* is the PatchGAN discriminator [21] that validates the plausibility of the synthesized image conditioned on the shape mask.




Figure 5. We compare our method to several baselines (EDN [7], V2V [61], HFMT [25], DIW [62]) on various sequences. For each example, we show the ground truth (GT) target appearance, the synthesized appearance by each method, and a color map of the error between the two. For our method, we also visualize the predicted surface normal.

3.3. Model-based Monocular 3D Pose Tracking

While there has been significant improvements in monocular 3D body estimation [9, 28, 29], we observe that predicting accurate and temporally coherent 3D body sequences is still challenging, which inhibit to reconstruct high-quality 3D motion descriptors. Hence, we devise a new optimization framework that learns a tracking function to address this challenge. We describe the full methodology and evaluation of our 3D tracking pipeline in the supplementary materials.

4. Experiments

We validate the performance of our method across various examples and perform extensive qualitative and quantitative comparisons with previous work.

Implementation Details. We utilize the Adam optimizer [27] to train our model with a learning rate of 1×10^{-3} . Given an input video ($\sim 10K$ frames), we train our model for roughly 72 hours using 4 NVIDIA V100 GPUs using a batch size of 4. Our motion features are learned from the body surface normals in the current frame and the body surface velocities in the past $t = 10$ frames. The features are recorded in a UV map of size 128×128 . We synthesize final renderings and surface normal maps of size 512×512 . We implement our model in Pytorch and utilize the Pytorch3D differentiable rendering layers [48]. Details of the network designs and 3D tracking pipeline are given in the supplementary materials.

The coordinate transformation that transports the motion

features from the UV space to the image plane, i.e., Π in Equation (5), can be implemented either by using image-based dense UV estimates [18] or by directly rendering the UV coordinates of the 3D body fits. To provide a fair comparison with previous work which also utilize dense UV estimates, we use the former option. When demonstrating our method on applications where we do not have corresponding ground truth frames to estimate dense UV maps (e.g., novel viewpoint synthesis), we use the latter option.

Baselines. We compare ours to four prior methods that focused on synthesizing dressed humans in motion. 1) Everybody dance now (EDN) [7] uses image-to-image translation to synthesize human appearance conditioned on 2D keypoints and uses a temporal discriminator to enforce plausible dynamic appearance. 2) Video-to-video translation (V2V) [61] is a sequential video generator that synthesizes high-quality human renderings from 2D keypoints and dense UV maps where the motion is modeled with optical flow in the image space. 3) High-fidelity motion transfer (HFMT) [25] is a compositional recurrent network which predicts plausible motion-dependent shape and appearance from 2D keypoints. 4) Dance in the wild (DIW) [62] synthesizes dynamic appearance of humans based on a motion descriptor composed of time-consecutive 2D keypoints and dense UV maps. We evaluate only on foreground by removing the background synthesized by the methods EDN, V2V, and DIW using a human segmentation method [16]. HFMT predicts a foreground mask similar to ours. In the supplementary material, we also compare our method to the 3D based approach [8]

Method	YouTube 1 (6K)	YouTube 2 (10K)	YouTube 3 (4K)	MPI (10K)	Custom 1 (15K)	Custom 2 (15K)	Avg.
EDN [7]	0.954 / 3.06 / 0.356	0.943 / 4.39 / 0.465	0.871 / 6.23 / 0.467	0.824 / 4.59 / 0.287	0.916 / 5.26 / 0.450	0.928 / 5.06 / 0.423	0.906 / 4.76 / 0.408
V2V [61]	0.960 / 2.23 / 0.235	0.958 / 3.33 / 0.405	0.880 / 4.47 / 0.401	0.824 / 3.58 / 0.298	0.935 / 3.52 / 0.306	0.943 / 4.15 / 0.385	0.916 / 3.54 / 0.338
HFMT [25]	0.944 / 4.19 / 0.412	0.923 / 6.63 / 0.775	0.862 / 7.16 / 0.456	0.826 / 5.03 / 0.291	0.905 / 6.24 / 0.321	0.915 / 6.63 / 0.390	0.895 / 5.98 / 0.440
DIW [62]	0.966 / 2.21 / 0.275	0.960 / 3.03 / 0.370	0.894 / 4.69 / 0.396	0.825 / 2.94 / 0.359	0.939 / 3.23 / 0.304	0.944 / 3.95 / 0.412	0.921 / 3.34 / 0.336
Ours	0.973 / 2.01 / 0.240	0.964 / 2.83 / 0.338	0.897 / 4.50 / 0.412	0.825 / 2.82 / 0.203	0.942 / 3.12 / 0.279	0.946 / 3.81 / 0.404	0.925 / 3.18 / 0.312

Table 1. Quantitative results. The number of training frames in each sequence is given in the top row. The three numbers are the SSIM (\uparrow), LPIPS (\downarrow) $\times 100$, and tLPIPS (\downarrow) $\times 100$ metrics, respectively. The red represents the best performer, and the blue second best.




Figure 6. Perceptual quality of a synthesized image over motion similarity between training and testing sequences.

for neural avatar modeling from a single camera, which explicitly reconstruct the geometry of animatable human.

Datasets. We perform experiments on video sequences that demonstrate a wide range of motion sequences and clothing types, which include non-trivial secondary motion. Specifically, we select three dance videos (e.g., hip-hop and salsa) from YouTube and one sequence from prior work [25] that shows a female subject in a large dress. We also capture two custom sequences showing a male and a female subject respectively performing assorted motions (e.g., walking, running, punching, jumping etc.) including 3D rotations.

Metrics. We measure the quality of the synthesized frames with two metrics: 1) Structure similarity (SSIM) [63] compares the local patterns of pixel intensity in the normalized luminance and contrast space. 2) Perceptual distance (LPIPS) [69] evaluates the cognitive similarity of a synthesized image to ground truth by comparing their perceptual features extracted from a deep neural network. We evaluate the temporal plausibility by comparing the perceptual change across frames [10]: $tLPIPS = \|LPIPS(s_t, s_{t-1}) - LPIPS(g_t, g_{t-1})\|$ where s and g are the synthesized and ground truth images.

4.1. Evaluation

Comparisons We provide quantitative evaluation in Table 1 and show qualitative results in Figure 5 (see Supplementary Video). Similar to our method, we train each baseline for roughly 72 hours until convergence. Both qualitative and quantitative results show that sparse 2D keypoint based pose representation used in EDN is not as effective as other baselines or our method. HFMT is successful in modeling dynamic appearance changes for mostly planar motions (i.e., MPI sequence), but shows inferior performance in remaining sequences that involve 3D rotations. This is due to the depth ambiguity inherent in sparse 2D keypoint based representation. While V2V performs well in terms of quan-

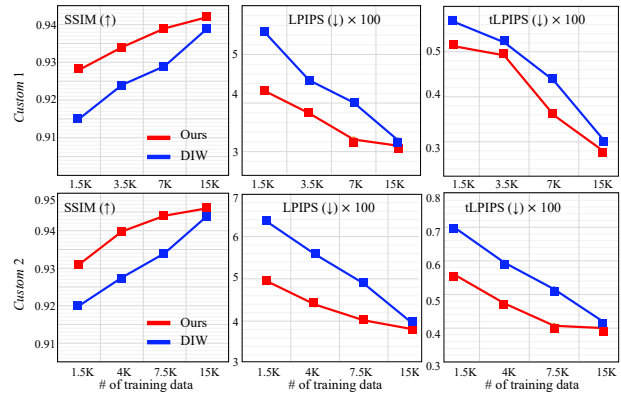


Figure 7. Performance depending on the amount of training data.

tative numbers, it suffers from significant texture drifting issues as shown in Figure 5, second row. We speculate that this is due to the errors in the optical flow estimation, especially in case of loose clothing, which is used as a supervisory signal. DIW uses dense UV coordinates to model the dynamic appearance changes of loose garments and is the strongest baseline. While it performs consistently well, we observe that the performance gap between DIW and our method increases for motion segments consisting of 3D rotations. This gap is magnified when the testing motion deviates from the training data. In Figure 6, we plot how the perceptual error changes along the motion similarity between the training and testing data which is computed by NCC between two sets of the time-varying 3D meshes similar to Figure 4. We observe bigger increase in the error for DIW as testing frames deviate more from the training data.

We next perform further comparisons with DIW evaluating the generalization ability of each approach.

Generalization An effective motion representation that encodes the dynamic appearance change of dressed humans should be discriminative to distinguish all possible deformations induced by a pose transformation given the current state of the body and the garments. In order to compare the discriminative power of our motion descriptor and the dense keypoint based representation proposed by DIW, we evaluate how well each representation generalizes to unseen poses. Specifically, we train each model using only 10% of the original training sequences by subsampling the training frames while ensuring training and testing pose sequences are sufficiently distinct. Considering the reduced amount of data, we limit the training time to 24 hours for both approaches. As shown in Table 2, the performance gap be-




Figure 8. Application. Our method enables several applications such as motion transfer with background composition, bullet time effects with novel view synthesis, and image-based relighting with the predicted surface normal.

Method	YouTube 1 (0.6K)	YouTube 2 (1K)	YouTube 3 (0.4K)	MPI (1K)	Custom 1 (1.5K)	Custom 2 (1.5K)	Avg.
DIW [62]	0.939 / 4.12 / 0.330	0.940 / 5.00 / 0.463	0.869 / 6.75 / 0.513	0.824 / 4.45 / 0.472	0.915 / 5.46 / 0.566	0.920 / 6.36 / 0.698	0.901 / 5.36 / 0.507
Ours	0.949 / 3.36 / 0.457	0.951 / 4.16 / 0.402	0.883 / 5.35 / 0.546	0.824 / 4.09 / 0.327	0.928 / 4.25 / 0.512	0.931 / 4.95 / 0.558	0.911 / 4.36 / 0.467

Table 2. We train DIW [62] and our method on a reduced training set (10% of the original training set) and test on the same testing set as in Table 1. The three numbers in each box represent the SSIM (\uparrow), LPIPS (\downarrow) $\times 100$, and tLPIPS (\downarrow) $\times 100$ metrics, respectively.

Method	Full data (15K)	10% data (1.5K)
w/o shape	0.945 / 4.31 / 0.401	0.929 / 5.28 / 0.565
w/o surface normal	0.945 / 3.89 / 0.418	0.929 / 5.17 / 0.602
w/o 3D motion	0.942 / 4.17 / 0.584	0.928 / 5.43 / 0.760
Full	0.946 / 3.81 / 0.404	0.931 / 4.95 / 0.558

Table 3. Ablation study. The three metrics are SSIM (\uparrow), LPIPS (\downarrow) $\times 100$, and tLPIPS (\downarrow) $\times 100$ respectively. The number in the top row denotes the amount of training data.

tween the two methods increases. For the *Custom 1* and *2* sequences, we further repeat the same experiment using 10%, 25%, and 50% of the original training data as shown in Figure 7 where the performance of our method shows slower degradation than that of DIW. These quantitative results as well as visual results provided in the supplementary materials demonstrate the superiority of our 3D motion descriptor in terms of generalizing to novel poses.

Ablation Study Using the *Custom 2* sequence, we train a variant **w/o 3D motion descriptors** by providing dense uv renderings as input directly to the decoder. We also disable the shape (**w/o shape**) and surface normal (**w/o surface normal**) prediction components. We repeat these trainings with subsampled data (10%). As shown in Table 3, the use of 3D motion descriptors and compositional rendering improves the perceptual quality of the synthesized images. The performance gap between our full model and w/o surface normal is larger with limited training data, implying that our multi-task framework helps with generalization. Qualitative results are given in the supplementary materials.

4.2. Applications

Our method enables several additional applications as shown in Figure 8. Since our method works with 3D body based motion representation, it can be easily used to transfer motion from a source to a target character by simply transferring the joint rotations between the characters. We can also create bullet time effects by creating a target mo-

tion sequence by globally rotating the 3D body. Thanks to the surface normal prediction, we can also perform relighting which is otherwise not applicable. Please refer to the supplementary material for more details and results.

5. Conclusion

We presented a method to render the dynamic appearance of a dressed human given a reference monocular video. Our method utilizes a novel 3D motion descriptor that encodes the time varying appearance of garments to model effects such as secondary motion. Our experiments show that our 3D motion descriptor is effective in modeling complex motion sequences involving 3D rotations. Our descriptor also demonstrates superior discriminator power compared to state-of-the-art alternatives enabling our method to better generalize to novel poses.

While showing impressive results, our method still has limitations. Highly articulated hand regions can appear blurry, hence refining the appearance of such regions with specialized modules is a promising direction. Our current model is subject specific, extending different parts of the model, e.g., 3D motion descriptor learning, to be universal is also an interesting future direction.

Broader Impact While our goal is to enable content creation for tasks such as video based motion retargeting or social presence, there can be some misuse of our technology to fabricate fake videos or news. We hope that parallel advances in deep face detection and image forensics can help to mitigate such concerns.

Acknowledgement

We would like to thank Julien Philip for providing useful feedback on our paper draft. Jae Shin Yoon is supported by Doctoral Dissertation Fellowship from University of Minnesota. This work is partially supported by NSF CNS-1919965.

References

- [1] Badour AlBahar, Jingwan Lu, Jimei Yang, Zhixin Shu, Eli Shechtman, and Jia-Bin Huang. Pose with Style: Detail-preserving pose-guided image synthesis with conditional stylegan. *ACM Transactions on Graphics*, 2021. 2
- [2] Timur Bagautdinov, Chenglei Wu, Tomas Simon, Fabian Prada, Takaaki Shiratori, Shih-En Wei, Weipeng Xu, Yaser Sheikh, and Jason Saragih. Driving-signal aware full-body avatars. *SIGGRAPH*, 2021. 2
- [3] Guha Balakrishnan, Amy Zhao, Adrian V Dalca, Fredo Durand, and John Guttag. Synthesizing images of humans in unseen poses. In *CVPR*, 2018. 2
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH*, page 187–194, 1999. 2
- [5] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017. 5
- [6] Dan Casas, Marco Volino, John Collomosse, and Adrian Hilton. 4d video textures for interactive character appearance. *Comput. Graph. Forum*, 33(2):371–380, May 2014. 2
- [7] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *ICCV*, 2019. 1, 2, 6, 7
- [8] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, and Huchuan Lu. Animatable neural radiance fields from monocular rgb video. *ICCV*, 2021. 2, 6
- [9] Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J Black. Monocular expressive body regression through body-driven attention. In *ECCV*, 2020. 4, 6
- [10] Mengyu Chu, You Xie, Jonas Mayer, Laura Leal-Taixé, and Nils Thuerey. Learning temporal coherence via self-supervision for gan-based video generation. *SIGGRAPH*, 2020. 7
- [11] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Trans. Graph.*, 34(4), July 2015. 2
- [12] Haoye Dong, Xiaodan Liang, Xiaohui Shen, Bochao Wang, Hanjiang Lai, Jia Zhu, Zhiting Hu, and Jian Yin. Towards multi-pose guided virtual try-on network. In *ICCV*, 2019. 2
- [13] Patrick Esser, Ekaterina Sutter, and Björn Ommer. A variational u-net for conditional appearance and shape generation. In *CVPR*, 2018. 2
- [14] Oran Gafni, Oron Ashual, and Lior Wolf. Single-shot freestyle dance reenactment. In *CVPR*, 2021. 2
- [15] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. *ICCV*, 2021. 2
- [16] Ke Gong, Xiaodan Liang, Yicheng Li, Yimin Chen, Ming Yang, and Liang Lin. Instance-level human parsing via part grouping network. In *ECCV*, 2018. 5, 6
- [17] Artur Grigorev, Artem Sevastopolsky, Alexander Vakhitov, and Victor Lempitsky. Coordinate-based texture inpainting for pose-guided human image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12135–12144, 2019. 2
- [18] Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *CVPR*, pages 7297–7306, 2018. 5, 6
- [19] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *SIGGRAPH*, 2021. 2
- [20] Xintong Han, Xiaojun Hu, Weilin Huang, and Matthew R Scott. Clothflow: A flow-based model for clothed person generation. In *ICCV*, 2019. 2
- [21] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 5
- [22] Yasamin Jafarian and Hyun Soo Park. Learning high fidelity depths of dressed humans by watching social media dance videos. In *CVPR*, pages 12753–12762, June 2021. 5
- [23] Eakta Jain, Yaser Sheikh, Moshe Mahler, and Jessica Hodgins. Augmenting Hand Animation with Three-dimensional Secondary Motion. In MZoran Popovic and Miguel Otaduy, editors, *SCA*. The Eurographics Association, 2010. 4
- [24] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018. 4
- [25] Moritz Kappel, Vladislav Golyanik, Mohamed Elgharib, Jann-Ole Henningson, Hans-Peter Seidel, Susana Castillo, Christian Theobalt, and Marcus Magnor. High-fidelity neural human motion transfer from monocular video. In *CVPR*, 2021. 2, 4, 6, 7
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410. Computer Vision Foundation / IEEE, 2019. 3
- [27] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2013. 6
- [28] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020. 4, 6
- [29] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 6
- [30] J. P. Lewis, Matt Corder, and Nickson Fong. Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. In *SIGGRAPH*, page 165–172, 2000. 4
- [31] Kun Li, Jingyu Yang, Leijie Liu, Ronan Boulic, Yu-Kun Lai, Yebin Liu, Yubin Li, and Eray Molla. Spa: Sparse photorealistic animation using a single rgb-d camera. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(4):771–783, 2017. 2
- [32] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 2
- [33] 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. 2021. 2

- [34] Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Hyeonwoo Kim, Florian Bernard, Marc Habermann, Wenping Wang, and Christian Theobalt. Neural rendering and reenactment of human actor videos. *SIGGRAPH*, 2019. 2
- [35] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: A skinned multi-person linear model. *ACM TOG*, 2015. 2, 4, 5
- [36] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *NeurIPS*, 2017. 2
- [37] 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 *CVPR*, 2021. 4
- [38] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed gan. In *CVPR*, 2020. 2
- [39] 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. 2
- [40] Natalia Neverova, Riza Alp Guler, and Iasonas Kokkinos. Dense pose transfer. In *ECCV*, 2018. 2
- [41] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Deformable neural radiance fields. *ICCV*, 2021. 2
- [42] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *SIGGRAPH Asia*, 2021. 2
- [43] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Animatable neural radiance fields for human body modeling. In *ICCV*, 2021. 2
- [44] Sergey Prokudin, Michael J Black, and Javier Romero. Smpix: Neural avatars from 3d human models. In *WACV*, 2021. 2
- [45] Albert Pumarola, Antonio Agudo, Alberto Sanfeliu, and Francesc Moreno-Noguer. Unsupervised person image synthesis in arbitrary poses. In *CVPR*, 2018. 2
- [46] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 2
- [47] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. Anr: Articulated neural rendering for virtual avatars. In *CVPR, year=2021*. 2
- [48] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 6
- [49] Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H Li, and Ge Li. Deep image spatial transformation for person image generation. In *CVPR*, 2020. 2
- [50] Kripasindhu Sarkar, Dushyant Mehta, Weipeng Xu, Vladislav Golyanik, and Christian Theobalt. Neural re-rendering of humans from a single image. 2
- [51] 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*, 2019. 2
- [52] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *NeurIPS*, 2019. 2
- [53] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. 5
- [54] Sijie Song, Wei Zhang, Jiaying Liu, and Tao Mei. Unsupervised person image generation with semantic parsing transformation. In *CVPR*, 2019. 2
- [55] Hao Tang, Song Bai, Li Zhang, Philip HS Torr, and Nicu Sebe. Xinggan for person image generation. In *European Conference on Computer Vision*, 2020. 2
- [56] 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. *ICCV*, 2021. 2
- [57] M. Volino, D. Casas, J.P. Collomosse, and A. Hilton. Optimal representation of multiple view video. In *BMVC*, 2014. 2
- [58] Chaoyang Wang, Ben Eckart, Simon Lucey, and Orazio Gallo. Neural trajectory fields for dynamic novel view synthesis. *arXiv*, 2021. 2
- [59] Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. Few-shot video-to-video synthesis. *NeurIPS*, 2019. 2
- [60] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. *NeurIPS*, 2018. 2
- [61] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, 2018. 6, 7
- [62] Tuanfeng Y. Wang, Duygu Ceylan, Krishna Kumar Singh, and Niloy J. Mitra. Dance in the wild: Monocular human animation with neural dynamic appearance synthesis, 2021. 3, 4, 6, 7, 8
- [63] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 13(4):600–612, 2004. 7
- [64] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *CVPR*, 2021. 2
- [65] Donglai Xiang, Fabian Andres Prada, Timur Bagautdinov, Weipeng Xu, Yuan Dong, He Wen, Jessica Hodgins, and Chenglei Wu. Explicit clothing modeling for an animatable full-body avatar. *arXiv preprint arXiv:2106.14879*, 2021. 2
- [66] Ceyuan Yang, Zhe Wang, Xinge Zhu, Chen Huang, Jianping Shi, and Dahua Lin. Pose guided human video generation. In *ECCV*, 2018. 2
- [67] 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
- [68] Meng Zhang, Tuanfeng Y. Wang, Duygu Ceylan, and Niloy J. Mitra. Dynamic neural garments. *ACM TOG*, 2021. 2
- [69] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 7
- [70] 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