Dance In the Wild: Monocular Human Animation with Neural Dynamic Appearance Synthesis

Tuanfeng Y. Wang   Duygu Ceylan   Krishna Kumar Singh
Adobe Research

{yangtwan,ceylan,krishsin}@adobe.com

Niloy J. Mitra
Adobe Research, University College London

nimitra@adobe.com

Abstract

Synthesizing dynamic appearances of humans in motion plays a central role in applications such as AR/VR and video editing. While many recent methods have been proposed to tackle this problem, handling loose garments with complex textures and high dynamic motion still remains challenging. In this paper, we propose a video based appearance synthesis method that tackles such challenges and demonstrates high quality results for in-the-wild videos that have not been shown before. Specifically, we adopt a StyleGAN based architecture to the task of person specific video based motion retargeting. We introduce a novel motion signature that is used to modulate the generator weights to capture dynamic appearance changes as well as regularizing the single frame based pose estimates to improve temporal coherency. We evaluate our method on a set of challenging videos and show that our approach achieves state-of-the-art performance both qualitatively and quantitatively.

1. Introduction

Generating plausible video based animations of an actor has several applications in AR/VR and video editing. A common approach to tackle this generation task is to transfer or retarget a motion sequence extracted from a *source* video to the *target* actor. Such approaches often depend on learning an actor specific appearance model from a training performance of the actor which is later utilized to synthesize the actor in novel poses.

While many recent approaches that leverage advances in machine learning have been proposed to the video-based performance retargeting problem, extending them to *in-the-wild* scenarios still remains challenging. First of all, many methods make simplifying assumptions about the appear-
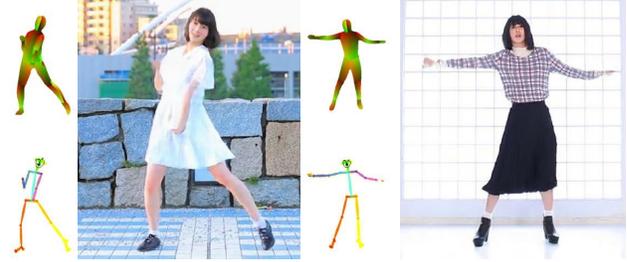
Figure 1. We present a method to synthesize the dynamic appearance of an actor given a target motion sequence. Our approach works well for actors wearing loose garments with complex textures and performing complex motions.

ance of the actor such as wearing tight clothing. Some recent work that extend to actors wearing loose garments [16], on the other hand, requires very long training videos of the actor which limit the applicability of the approach. Furthermore, such methods are often limited in terms of the complexity of the motions they can handle. We observe that with loose garments and complex motion sequences, state-of-the-art pose estimation methods [3, 10] often suffer from wrong predictions, missing parts, and temporal noise; making it particularly challenging to learn a consistent appearance model for the actor which can capture motion dependent appearance changes.

In this work, we propose a novel approach to learn the dynamic appearance of an actor, potentially wearing loose garments (which may be far away from the body and may have significant dynamics under motion), and synthesize unseen complex motion sequences (see Fig. 1). Our approach adopts the state-of-the-art photo-realistic generative models for human portraits and bodies, such as Style-GAN [18, 19] and StylePoseGAN [33], to the task of person-specific appearance modeling and motion retargeting. In order to disentangle the pose and appearance of the actor, we utilize off-the-shelf keypoint [3] or dense corre-

spondence [10] estimations to represent the target pose and provide as spatial features to the StyleGAN architecture. The core of our approach is to learn an effective motion representation that is used to demodulate the weights of the generator. This modulation not only helps to capture the appearance of loose garments that heavily depend on the underlying body motion but also captures plausible motion specific appearance changes (see Fig. 6). Since keypoint or dense correspondence estimation methods operate on single images, they often result in pose representations that are not perfect and temporally coherent. We show that we can also use the motion features to refine the per-frame pose features and provide significant improvements in terms of temporal coherency as well.

We evaluate our method on several challenging motion sequences as well as target actors with loose garments consisting of complex texture patterns. We show that our results can synthesize plausible garment deformations while also maintaining high quality visual results. We also perform extensive qualitative and quantitative comparisons with previous methods and demonstrate state of the art results. In summary, our contributions are as follows:

- We adopt the StyleGAN architecture to the task of actor specific video based motion retargeting. We demonstrate results that capture the dynamic appearance of loose garments under complex motion sequences that have not been shown before.

- We introduce an explicit motion representation that is used for demodulation of the generator weights. Such motion features are used to both capture motion specific appearance changes as well as generating temporally coherent results.

2. Related Work

Rendering of humans in different poses, from different viewpoints has attracted a lot of attention both from the graphics and vision communities. Below, we review the previous work most related to our method.

Rendering based approaches. Given very high quality geometry and appearance representations, a traditional approach to controllable synthesis of humans is to employ classical rendering techniques. Hence, several works have focused on reconstructing such geometry and appearance representations from multi-view captures of humans [40, 4, 7]. With the recent success of deep learning, more recent methods have replaced certain components of such a pipeline with neural networks. For example, given a person specific 3D template of an actor, some methods learn dynamic appearances of the actor with neural networks [23, 22, 11, 43]. Although providing explicit control, the 3D template requirement limits the application

of such methods. Another recent thread of work has explored the use of *neural textures* [39, 8] and *neural rendering* [38] to synthesize humans under different poses and viewpoints [35, 34, 32, 29]. Most of these approaches, however, assume humans wear tight clothing and cannot handle complex motion sequences as our method.

Image generation methods. Image generation methods have gained significant popularity in the human reposing problem, i.e., generating a new image of a person given a source image and a target pose. Approaches that synthesize images from target poses represented as keypoints [26, 31, 36, 45], dense uv coordinates [28, 9, 30], and heatmaps [27] have been presented. Some methods have explored the use of warping techniques along with GANs [24] to improve the quality of the results. In the context of synthesizing human performance videos, while some approaches have explored temporal cues as training losses [5, 1], others have also utilized the temporal information at inference time [41, 2]. In both single image and video-based generation methods, one problem that has been relatively less explored is to handle actors with loose clothing. Very recently, Kappel et al. [16] have presented a recurrent network architecture that performs video-based animation for actors wearing loose garments. Starting from a set of 2D keypoints, they predict semantic part labels, a structure map that encodes garment wrinkles and texture patterns, and the final appearance of the actor. While showing impressive results, our experiments show that the method requires relatively long training sequences and is limited in terms of the complexity of the motions it can handle.

In the context of face generation, StyleGAN [18, 19] has shown impressive results and many follow up work has focused on providing control in the latent space of the generator [37]. Compared to faces, however, controllable synthesis of full human bodies is more challenging. Very recently, StylePoseGAN [33], VOGUE [21], and TryOnGAN [20] present earlier examples of adopting the StyleGAN architecture to the tasks of reposing and virtual try on. While our work is inspired by such approaches, we focus on the different and challenging task of dynamic motion synthesis.

3. Methodology

3.1. Overview

Our framework learns the dynamic appearance of a specific actor $\mathbb{X}$ from a reference video of the actor. We represent the reference video as a sequence of RGB image and pose pairs $\{\mathbb{I}_i, \mathbb{P}_i\}, i = 1, \cdots N$. We assume the camera is fixed so the pose of the actor is represented by image-space pose representations such as 2D keypoints [3] and dense uv renderings [10]. At test time, given a motion sequence consisting of query poses $\{\mathbb{P}'_i\}, i = 1, \cdots M$, the network synthesizes the corresponding appearance of $\mathbb{X}$ performing
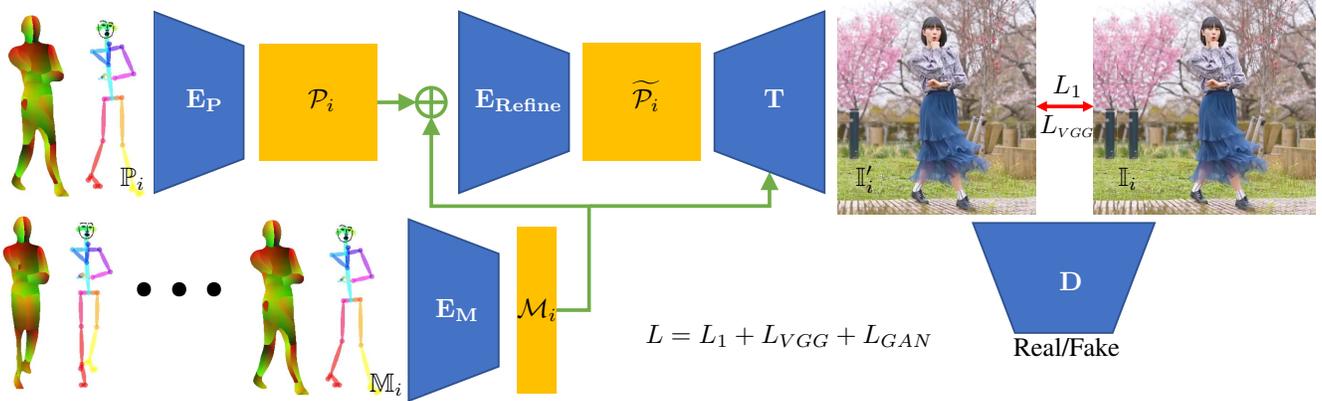
Figure 2. Pipeline overview. Our network takes the 2D body dense UV and keypoints as input and learns a pose feature for each frame. By concatenating the pose inputs for the past few frames, we also learn motion features. The learned motion features are used to refine the pose features to improve the temporal coherency. We synthesize the final motion-aware dynamic appearance of the character using a StyleGAN based generator conditioned on the refined pose features and modulated by the motion features.

the motion. We adopt a StyleGAN based generator which has recently been shown to be effective in synthesizing high quality full body images [33]. In order to guide the generation of the actor performing a specific pose, we extract spatial pose features for each frame given $\mathbb{P}_i$ similar to [21, 33]. Unlike static image generation, however, the appearance of the garments, specifically if they are loose, are heavily affected by the motion of the actor. In order to capture such dynamic appearance changes (e.g., folds and wrinkles and secondary movement of loose clothes and hair), we introduce an explicit motion feature representation that is extracted from the past K frames. We use such motion features as the latent style code to demodulate the generator. Furthermore, we observe two common limitations of the off-the-shelf pose estimators [3, 10] as shown in Fig. 3. First, being trained on single images, they often produce temporally jittery results when run on a video sequence. Second, in case of complex poses with self-occlusions, they are prone to missing parts and incorrect estimations. Since past frames provides strong cues to overcome these limitations, we propose a refinement module to regularize the spatial pose features based on the learned motion features. With the refined spatial pose features and the learned motion features, our generator synthesizes the final output image, $\mathbb{I}'_i$. Fig. 2 provides an overview of our pipeline.

3.2. Input

Given a sequence of RGB images $\{\mathbb{I}_i\}, i = 1 \cdots N$, extracted from a captured video stream, we run Dense-Pose [10] to predict dense body IUV maps which are three channel images of the same size of the input image. We also use OpenPose [3] to predict keypoints, i.e., skeleton, face, and hand landmarks, which are represented as RGB images of the same size (W, H) as the input image. We concatenate the two sources of information to form

a pose signature $\mathbb{P}_i \in \mathbb{R}^{6 \times W \times H}$ for each input frame i. In order to extract explicit motion features, as explained in the next section, we also provide the pose signature of $K = 10$ frames sampled unevenly from the past 20 frames as input to our network. Specifically, we use the frames $\{1, 2, 3, 4, 6, 8, 10, 13, 16, 20\}$ to build a motion signature $\mathbb{M}_i \in \mathbb{R}^{60 \times W \times H}$. The frames are sampled unevenly as closer frames have more impact on the current frame.

Figure 3. The current state of the art pose detection methods suffer from artifacts such as jitter, missing parts, and wrong detections. Given a reference image, we show the predictions for two consecutive frames. (Left) The dense body UV predicted by Dense-Pose [10] has mistakes in the left arm and leg regions, it also misses the right arm. (Right) The keypoints predicted by Open-Pose [3] miss the shoulder and the legs.

3.3. Learning Pose/Motion Features

Inspired by [33], we extract 2D spatial pose features $\mathcal{P}_i$ from the pose signature $\mathbb{P}_i$ to condition the generator to synthesize an image of the actor performing a specific pose. Since our generator is trained on a reference video of a specific actor, the global appearance of the actor is implicitly embedded into the generator weights. Unlike single-image generation case, however, the overall shape and appearance of the garments heavily depend on the motion being performed (e.g., the overall movement of a loose skirt, varying wrinkle and fold patterns). In order to capture such dynamic appearance changes, we propose to extract a 1D mo-

tion feature $\mathcal{M}_i$ from the motion signature $\mathbb{M}_i$ which are used to demodulate the generator via the AdaIN operations. To this end, we use a convolutional neural network $\mathbf{E_P}$ to encode the pose signature $\mathbb{P}_i \in \mathbb{R}^{6 \times W \times H}$ into a pose feature $\mathcal{P}_i \in \mathbb{R}^{512 \times W_s \times H_s}$, where $W_s = W/16, H_s = H/16$. Similarly, we extract the motion features with another convolutional neural network, $\mathbf{E_M}$, which consists of a reshape operation and some fully connected layers to produce a 1D motion feature with a dimension of 2048. We refer to the supplementary material for the detailed architecture.

Temporal coherent refinement. As shown in Fig. 3, the current state of the art pose detection methods suffer from artifacts such as jitter, missing parts, and wrong detections. We hypothesize that observing the past motion, although imperfect, provides strong cues to identify and regularize such artifacts. Therefore, we propose to use the motion features $\mathcal{M}_i$ to refine the pose feature of the current frame $\mathcal{P}_i$. To realize such a refinement operation, we first concatenate $\mathcal{M}_i$ to each spatial location of $\mathcal{P}_i \in \mathbb{R}^{512 \times W_s \times H_s}$ along the channel dimension to obtain an intermediate pose feature $\mathcal{P}_i^{int} \in \mathbb{R}^{2560 \times W_s \times H_s}$. $\mathcal{P}_i^{int}$ is then passed through another convolutional neural network $\mathbf{E_{Refine}}$ which produces a refined pose feature $\widetilde{\mathcal{P}_i} \in \mathbb{R}^{512 \times W_s \times H_s}$. In Fig. 7 and Fig. 8, we show that $\mathbf{E_{Refine}}$ learns to effectively regularize the input pose of the current frame based on the motion feature which is learned from the past frames.

3.4. Style-based Generator

We use the style-based generator proposed by [19] for our image generation task. The original StyleGAN architecture takes a constant spatial tensor as input, and a latent noise vector passed through a mapping network to demodulate the intermediate layers to control the details of final generated image. As inspired by [21, 33], we use our learned spatial pose feature $\widetilde{\mathcal{P}_i}$ instead of the constant tensor to condition the generator. Different from previous work, we use the motion feature to control the dynamic appearance. Specifically, our generator, $\mathbf{T}(\widetilde{\mathcal{P}_i}, \mathcal{M}_i)$, takes the spatial pose feature $\widetilde{\mathcal{P}_i}$ as input and passes it through a set of convolutional layers. The convolutional weights are then demodulated by the 1D motion feature $\mathcal{M}_i$. Our generator converts the spatial feature from size $512 \times W_s \times H_s$ to $3 \times W \times H$ after four residual blocks and four upsampling residual blocks, with random noise injection at every layer similar to [19]. Therefore, our full pipeline can be written as:

$$\mathbf{T}(\mathbf{E_{Refine}}(\mathbf{E_P}(\mathbb{P}_i), \mathbf{E_M}(\mathbb{M}_i))|\mathbf{E_M}(\mathbb{M}_i)) = \mathbb{I}'_i \quad (1)$$

3.5. Loss Function

We train our framework in an end-to-end manner. For a specific frame, i, in the training sequence, we directly supervise the generated image, $\mathbb{I}'_i$, with the ground truth image

$\mathbb{I}_i$. Our loss function includes three terms. First, we use an $L1$ reconsturction loss:

$$L_1 = |\mathbb{I}'_i - \mathbb{I}_i|. \quad (2)$$

Second, we use the VGG-based perceptual loss [15] to encourage perceptual similarity:

$$L_{VGG} = \sum_k MSE(VGG_k(\mathbb{I}'_i) - VGG_k(\mathbb{I}_i)), \quad (3)$$

where $MSE(\cdot)$ is the element-wise Mean-Square-Error, $VGG_k(\cdot)$ is the k-th layer of a VGG network pre-trained on ImageNet. Finally, we use an adversarial loss L_{GAN} which utilizes a discriminator $\mathbf{D}$ identical to [19]. During training, we minimize our objective $L = L_1 + L_{VGG} + L_{GAN}$ w.r.t. the parameters of $\mathbf{E_P}$, $\mathbf{E_M}$, $\mathbf{E_{Refine}}$, and $\mathbf{T}$; and maximize L_{GAN} w.r.t. $\mathbf{D}$.

4. Experiments

We evaluate our approach on a set of Youtube videos with complex motion sequences as well as the datasets from existing work [5, 16] (Sec. 4.1). We provide both quantitative and qualitative comparisons with previous work (Sec. 4.2) and perform ablation studies to motivate the various design choices (Sec. 4.3).

We implement our network in Pytorch operating on images of resolution 512×512. We use the Adam optimizer with a learning rate of 0.02. For a training sequence of $6K$ frames, the training takes approximately 72 hours for $100K$ iterations when trained with a batch size of 16 on 4 NVIDIA V100 GPUS. Inference with our model takes about 40 ms/frame on a single NVIDIA V100 GPU.

4.1. Datasets

To test the performance of our method on unconstrained videos, we collect seven dancing sequences from Youtube. The selected videos are captured with a fixed camera, and show variation in terms of scenes (indoor vs. outdoor), garment types (tight clothing, loose or multilayer garments) and appearances (plain color, grid or stochastic texture, text logo). We also test our method on the sequences provided by [5] and [16]. In order to quantify the complexity of the motion sequences, we also calculate the corresponding motion speed. Specifically, we calculate the average displacement of all keypoints between two consecutive frames which have been normalized to have a height of 1. In Table 1 , we provide the different characteristics of each sequence along with the length of the sequences, i.e., the number of frames sampled with a frame rate of 24 fps. The speed of the motion for the sequences in our dataset are significantly higher than previous datasets demonstrating the complexity of our examples. The duration of a typical Youtube dancing video is between 2-5 minutes, which
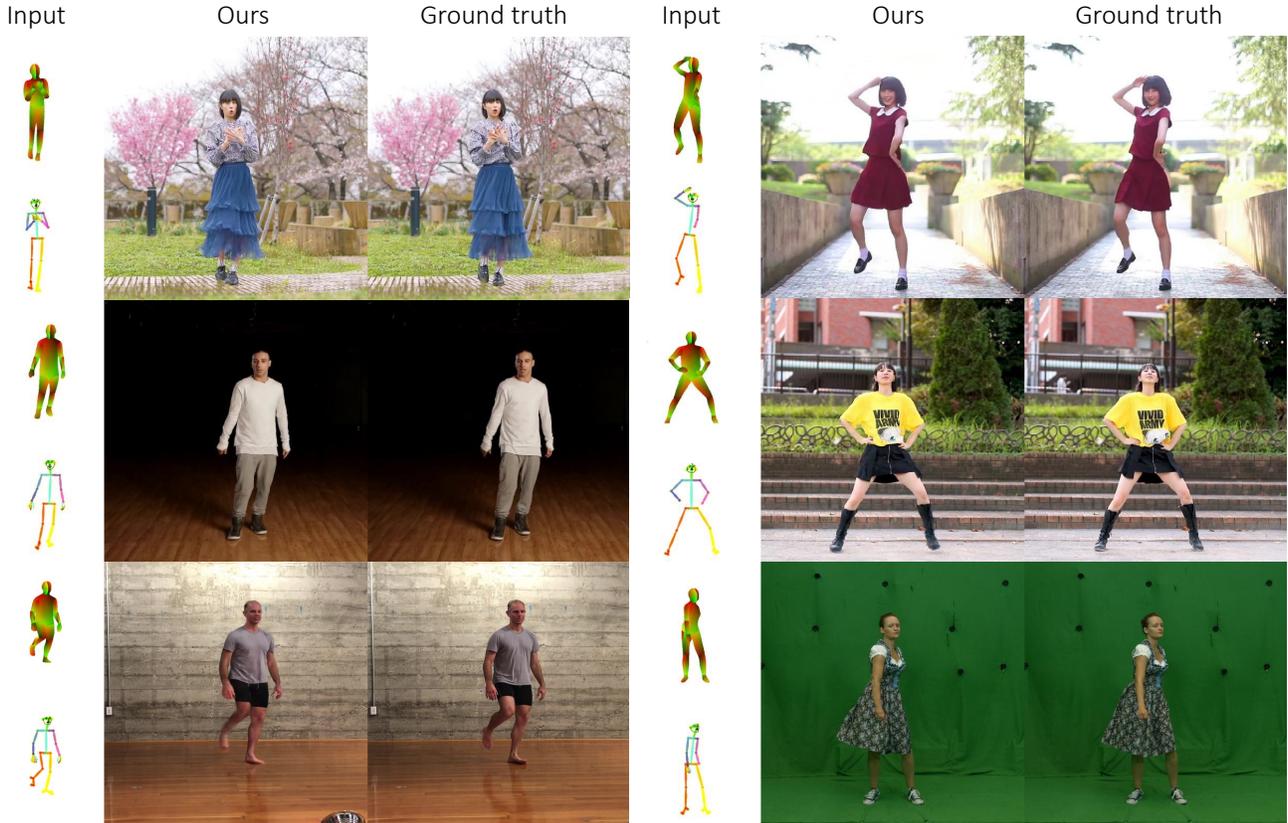
Figure 4. We train our network for character specific appearance synthesis. Here are the exemplary results on the test set for some sequences.

results in a sequence with 2k to 8k frames. This is significantly shorter than previous sequences captured in controlled lab settings, which often include more than 10k frames. We note that shorter sequences are more challenging test cases with respect to generalization to unseen poses.

	length	motion speed	clothes type	texture
Seq 1	7.5k	3.7×	loose	plain
Seq 2	3.4k	2.9×	loose	grid
Seq 3	6.0k	3.0×	loose	plain
Seq 4	6.0k	4.3×	loose	stochastic
Seq 5	3.2k	4.0×	tight	text
Seq 6	6.5k	1.4×	tight	plain
Seq 7	6.1k	4.5×	multi-layer	stochastic
Seq 8 ([16])	12.5k	0.006(1.0×)	loose	stochastic
Seq 9 ([5])	11.4k	2.7×	tight	plain

Table 1. Statistics of our dataset. For the motion speed of each sequence, we show the multiple of Seq 8.

4.2. Evaluation and Comparisons

For each sequence, we train our network with the first 85% of the frames and test with the last 10% of the frames. We skip the intermediate 5% of frames to make sure that the beginning of the test sequence is different from the end of the train sequence. In Fig 4, we show example frames from the testing sequences where we synthesize the appearance of the character given different target poses and motion signatures. We refer to the supplementary video for more examples.

We also compare our approach with several previous approaches: (i) pix2pixHD [14], the state-of-the-art image-to-image translation method that generates the full body appearance from dense uv based pose representations, (ii) its temporal extension vid2vid [41] which utilizes both keypoint and dense uv based posed representations, (iii) Everybody-Dance-Now (EDN) [5] which takes 2D keypoints as input, and (iv) the recently proposed recurrent based architecture that focuses on actors with loose garments (HF-NHMT) [16]. We also include the nearest neighbour (NN) baseline where for each query pose we select the frame from the training set with the most similar pose. We define the pose similarity as the L2 distance between the 2D keypoints. This baseline cannot capture the desired target poses since they are significantly different than the training poses. We train all the learning based alternatives on the same set of training frames as our method until convergence. In Fig 5, we provide visual results of our method, the

| | Ground truth | Ours | pix2pixHD [14] | vid2vid [41] | EDN [5] | HF-NHMT [16] | Nearest neighbor |

Figure 5. We qualitatively compare our method with other related approaches. For the examples on the second row, we provide a close-up view for the text pattern on the shirt to highlight the synthesis quality.

baseline methods, and the corresponding ground truth. In Table 2, we provide quantitative evaluations with respect to the ground truth using several metrics: (i) the mean square error (MSE) of the pixel value normalized to the range $[-1, 1]$, (ii) the structural similarity index (SSIM) [42], (iii) the perceptual similarity metric (LPIPS) [44], (iv) the Fréchet Inception Distance (FID) [12], and (v) the tOF [6], pixel-wise difference of the estimated optical flow between each sequence and the ground truth. Our method outperforms the baselines with respect to all the metrics. Since our motion features $\mathcal{M}_i$ are learned from a short clip of past frames (a.k.a., motion window), it naturally encodes the temporal information. Therefore, we observe that our method generates more temporally smooth results without the need of an extra temporal discriminator. We also observe that conditioning the synthesis of the current frame on the motion features captures motion dependent appearance changes while avoiding the problem of error accumulation that is apparent in recurrent based approaches [16].

	MSE ↓	SSIM ↑	LPIPS ↓	FID ↓	tOF ↓
pix2pixHD [14]	0.0212	0.9807	0.0474	36.1323	7.6282
vid2vid [41]	0.0276	0.9795	0.2318	58.2085	5.5523
EDN [5]	0.0201	0.9811	0.0423	30.6969	6.1260
HF-NHMT [16]	0.0743	0.9629	0.1498	53.6912	8.5114
Ours	**0.0199**	**0.9813**	**0.0398**	**21.1877**	**5.4122**

Table 2. Quantitative comparison to the related approaches. We evaluate the MSE, SSIM, LPIPS, FID, and tOF score of each method.

4.3. Ablation study

We perform various ablation studies to motivate our design choices. First, in order to demonstrate the effect of the motion features on capturing dynamic appearance changes, given the same pose signature $\mathbb{P}$, we synthesize frames with different motion signatures $\mathbb{M}$. Specifically, given a testing motion sequence, we synthesize a target frame by (i) computing the motion signature from the original sequence of past poses, i.e., *forward* motion, (ii) hallucinate a still motion, i.e., a *frozen* motion signature by using the same pose for all the past frames, and (iii) hallucinate the *backward* motion, i.e., a motion signature computed from the future

Figure 6. We show that our model captures motion specific appearance changes. We synthesize the same target frame by using a motion signature computed from (i) the original past frames, i.e., forward motion, (ii) from the same pose repeated for each past frame, i.e., frozen motion, and (iii) from the future frames in the reverse order, i.e., backward motion. In the forward motion, the character is swinging from her left side to right resulting in the skirt being dragged towards her left. While there is no significant dragging of the skirt with frozen motion signatures, the dragging is reversed with the backward motion.
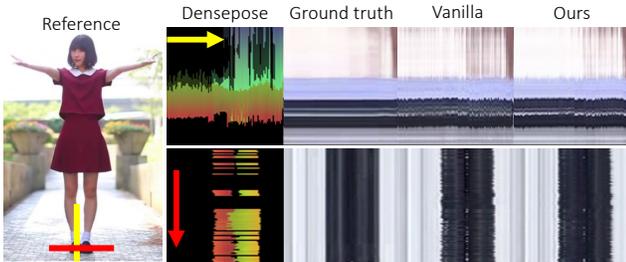
Figure 7. We cut a slice of vertical and horizontal pixels from a sequence of 100 frames and concatenate them to form slice plots. Due to the artifacts in the dense pose estimation, the vanilla model generates jittery results. Ours refines the pose features and generates smoother results visually more similar to the ground truth.

frames in the reverse order. As shown in Fig 6, we show that our method is able to synthesize the overall shape of the garment that agree with the motion signature. While the flow of the skirt is reversed when simulating the backward motion, with frozen motion signatures, the skirt tends to keep a more stable rest shape.

Second, we study the effect of the length of the past frames used to compute the motion signature. Specifically, we compare four different cases where we use the (i) past 4 frames, $\{1, 2, 3, 4\}$; (ii) past 10 frames, $\{1, 2, 3, 4, 6, 8, 10\}$; (iii) past 20 frames, $\{1, 2, 3, 4, 6, 8, 10, 13, 16, 20\}$ (the setting used in other experiments); and (iv) past 40 frames, $\{1, 2, 3, 4, 6, 8, 10, 13, 16, 20, 24, 29, 34, 40, 47, 56\}$. We calculate mean square error (MSE) and the structural similarity index (SSIM) [42] for the generated test sequences with respect to the ground truth. In Table 3, we show that using poses sampled in the past 20 frames provide a good balance. Shorter motion windows are not sufficient to capture motion dependent dynamic appearance changes. The

Figure 8. We use [13] to compute the flow map between the consecutive frames in the ground truth and generated sequences with ours and the vanilla model. Our approach better captures the motion of the ground truth sequence.

trend shows the performance will decrease if the motion window gets shorter. When training without motion feature and pose refinement, the MSE and SSIM drops to 0.0213 and 0.9801 respectively. With longer motion windows, we observe no significant improvement in terms of capturing motion dependent appearance changes but training becomes more difficult due to the increase in the network size. Please also refer to the supplementary video for a qualitative comparison.

	0	5	10	20 (ours)	40
MSE $\downarrow$	0.0213	0.0208	0.0201	**0.0199**	0.0205
SSIM $\uparrow$	0.9801	0.9809	0.9812	**0.9813**	0.9811

Table 3. We evaluate the MSE and SSIM score for different length of the motion signature. Taking the past 20 frames is a good trade of between complexity and high-fidelity.

Finally, we study the effect of our motion driven pose refinement on temporal coherency. Specifically, we compare our proposed network with a vanilla motion driven pipeline which directly feeds the generator $\mathbf{T}$ with the learned pose feature $\mathcal{P}_i$, instead of the refined features $\widetilde{\mathcal{P}_i}$. We find that the vanilla motion driven approach achieves similar per-frame image synthesis quality but suffers from significant temporal coherency artifacts. In Fig. 7, we identify a vertical and a horizontal slice of pixels in a particular frame and concatenate such slices over time (for 100 frames) along the dimension denoted by the arrows to form a slice plot. The corresponding video clip does not contain significant movements around the leg resulting in smooth ground truth slice plots along the time axis. However, the dense body UV predictions [10] are not temporally stable producing high frequency signals in the corresponding densepose slice plots. The vanilla motion driven pipeline is affected by such artifacts and reproduces the high frequency signals, while ours with pose refinement produces a significantly smoother slice plot. We also present the predicted flow map [13] between the consecutive frames in the ground truth and generated sequences in Fig. 8 to show that our approach best captures the motion of the ground truth regardless the im-

Figure 9. Our approach allows motion transfer from a source sequence to a target actor.

perfect pose input. The average mean square error between the ground truth and generated flow maps also show the advantage of our method, ours and the vanilla version achieve an error of 5.41 and 6.11 respectively.

Please refer to our supplementary material for more evaluation and experiments.

5. Application: Motion transfer

Once we have an actor specific network, a typical application is to retarget a source motion extracted from a video of another actor. Since the two actors might have different body proportions, before retargeting, we first perform a simple alignment. We adjust the height and width of the detected skeletons and move them vertically so that they stand on the ground of the target background image [5]. In Fig. 9, we show some examples, please refer to our supplementary video for more results.

6. Conclusion, Limitations and Future Work

In this paper, we present a StyleGAN based approach for video synthesis of an actor performing complex motion sequences. We propose a novel motion signature which is used to modulate the generator to capture dynamic appearance changes as well as refining the pose features to obtain temporally coherent results. Our approach outperforms the previous alternatives both quantitatively and qualitatively.

While showing impressive results, our method has some limitations. We use image based pose representations which are limited to capture poses such as 3D rotations. Incorporating 3D representations, e.g., SMPL [25], is an interesting direction. The quality of the results can degrade with very complex motions significantly different than training sequences (see the supplementary material). Improving motion generalization is a promising direction. We want to extend our work to videos captured with moving cameras and incorporate the motion of the camera into the generation process. We present a simple retargeting and alignment process for motion transfer. More sophisticated retargeting strategies are needed to transfer motion between actors

with significantly different body shapes. While our method generates reasonable face and hand regions, further refining these body parts with specialized modules is worth investigating. Currently the model is trained for a specific identity, supporting multiple identities by conditioning our network on an extra channel that denotes the character identity is an interesting direction for future work. Finally, adopting advanced GAN architectures, e.g., Alias-Free GAN [17], may further improve the performance without much change in the overall design.

Acknowledgement. We would like to thank the anonymous reviewers for their constructive comments; as well as Jae Shin Yoo, Moritz Kappel, and Erika Lu for their kind help with the experiments and discussions.

References

[1] K. Aberman, M. Shi, J. Liao, D. Lischinski, B. Chen, and D. Cohen-Or. Deep video-based performance cloning. *Computer Graphics Forum*, 38(2):219–233, 2019. 4322

[2] Aayush Bansal, Shugao Ma, Deva Ramanan, and Yaser Sheikh. Recycle-gan: Unsupervised video retargeting. In *ECCV*, 2018. 4322

[3] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 4321, 4322, 4323

[4] Dan Casas, Marco Volino, John Collomosse, and Adrian Hilton. 4d video textures for interactive character appearance. *CGF*, 33(2):371–380, May 2014. 4322

[5] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *ICCV*, 2019. 4322, 4324, 4325, 4326, 4328

[6] Mengyu Chu, You Xie, Jonas Mayer, Laura Leal-Taixé, and Nils Thuerey. Learning temporal coherence via self-supervision for gan-based video generation. *ACM Transactions on Graphics (TOG)*, 39(4):75–1, 2020. 4326

[7] 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. 4322

[8] 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 (CVPR)*, pages 5151–5160, June 2021. 4322

[9] 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. 4322

[10] Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. 2018. 4321, 4322, 4323, 4327

[11] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM TOG*, 40(4), aug 2021. 4322

[12] 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. 4326

[13] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017. 4327

[14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017. 4325, 4326

[15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 4324

[16] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1541–1550, June 2021. 4321, 4322, 4324, 4325, 4326

[17] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *CoRR*, abs/2106.12423, 2021. 4329

[18] 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. 4321, 4322

[19] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, volume abs/1912.04958, 2019. 4321, 4322, 4324

[20] Kathleen M Lewis, Srivatsan Varadharajan, and Ira Kemelmacher-Shlizerman. Tryongan: Body-aware try-on via layered interpolation. *ACM SIGGRAPH*, 40(4), 2021. 4322

[21] Kathleen M Lewis, Srivatsan Varadharajan, and Ira Kemelmacher-Shlizerman. Vogue: Try-on by stylegan interpolation optimization. *arXiv preprint arXiv:2101.02285*, 2021. 4322, 4323, 4324

[22] Lingjie Liu, Weipeng Xu, Marc Habermann, Michael Zollhöfer, Florian Bernard, Hyeongwoo Kim, Wenping Wang, and Christian Theobalt. Neural human video rendering by learning dynamic textures and rendering-to-video translation. *IEEE TVCG*, PP:1–1, 05 2020. 4322

[23] Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Hyeongwoo Kim, Florian Bernard, Marc Habermann, Wenping Wang, and Christian Theobalt. Neural rendering and reenactment of human actor videos. *ACM TOG*, 2019. 4322

[24] Wen Liu, Zhixin Piao, Min Jie, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework

for human motion imitation, appearance transfer and novel view synthesis. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 4322

[25] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multiperson linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. 4328

[26] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *NeurIPS*, pages 405–415, 2017. 4322

[27] Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. In *CVPR*, June 2018. 4322

[28] Natalia Neverova, Riza Alp Güler, and Iasonas Kokkinos. Dense Pose Transfer. In *ECCV*, Munich, Germany, Sept. 2018. 4322

[29] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 4322

[30] Sergey Prokudin, Michael J. Black, and Javier Romero. SM-PLpix: Neural avatars from 3D human models. In *Winter Conference on Applications of Computer Vision (WACV)*, pages 1810–1819, Jan. 2021. 4322

[31] Albert Pumarola, Antonio Agudo, Alberto Sanfeliu, and Francesc Moreno-Noguer. Unsupervised person image synthesis in arbitrary poses. In *CVPR*, pages 8620–8628. IEEE Computer Society, 2018. 4322

[32] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. Anr-articulated neural rendering for virtual avatars. In *arXiv:2012.12890*, 2021. 4322

[33] Kripasindhu Sarkar, Vladislav Golyanik, Lingjie Liu, and Christian Theobalt. Style and pose control for image synthesis of humans from a single monocular view, 2021. 4321, 4322, 4323, 4324

[34] Kripasindhu Sarkar, Dushyant Mehta, Weipeng Xu, Vladislav Golyanik, and Christian Theobalt. Neural rerendering of humans from a single image. In *European Conference on Computer Vision (ECCV)*, 2020. 4322

[35] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliev, Renat Bashirov, Egor Burkov, Karim Iskakov, Aleksei Ivakhnenko, Yury Malkov, Igor Pasechnik, Dmitry Ulyanov, Alexander Vakhitov, and Victor Lempitsky. Textured neural avatars. In *CVPR*, June 2019. 4322

[36] Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, and Nicu Sebe. Deformable gans for pose-based human image generation. In *CVPR)*, June 2018. 4322

[37] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zöllhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images, cvpr 2020. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, june 2020. 4322

[38] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, K. Sunkavalli, R. Martin-Brualla, T. Simon, J. Saragih, M. Nießner, R. Pandey, S. Fanello, G. Wetzstein, J.-Y. Zhu, C. Theobalt, M. Agrawala, E. Shechtman, D. B Goldman, and M. Zollhöfer. State of the art on neural rendering. *CGF*, 39(2):701–727, 2020. 4322

[39] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM TOG*, 38(4):1–12, 2019. 4322

[40] M. Volino, D. Casas, J.P. Collomosse, and A. Hilton. Optimal representation of multiple view video. 2014. 4322

[41] 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. 4322, 4325, 4326

[42] 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. 4326, 4327

[43] Meng Zhang, Duygu Ceylan, Tuanfeng Wang, and Niloy J Mitra. Dynamic neural garments. *arXiv preprint arXiv:2102.11811*, 2021. 4322

[44] 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. 4326

[45] Zhen Zhu, Tengteng Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive pose attention transfer for person image generation. In *CVPR*, pages 2347–2356, 2019. 4322