

High-Fidelity Neural Human Motion Transfer from Monocular Video

Moritz Kappel¹ Vladislav Golyanik² Mohamed Elgharib² Jann-Ole Henningson¹
 Hans-Peter Seidel² Susana Castillo¹ Christian Theobalt² Marcus Magnor¹

¹ Computer Graphics Lab, TU Braunschweig, Germany {lastName}@graphics.tu-bs.de

² Max Planck Institute for Informatics, Saarland Informatics Campus, Germany {lastName}@mpi-inf.mpg.de

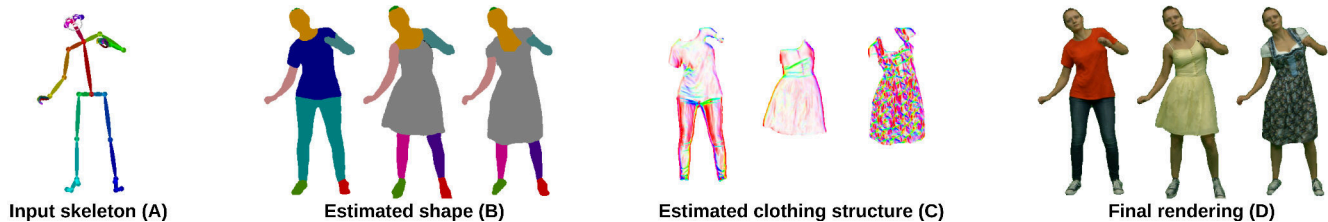


Figure 1. Our human motion transfer framework translates 2D pose skeletons (A) to photorealistic images (D) by explicitly estimating the actor’s 2D shape (B) and internal structure of clothing (C), enabling direct deformation handling and support of high-frequency details such as wrinkles and shadows.

Abstract

Video-based human motion transfer creates video animations of humans following a source motion. Current methods show remarkable results for tightly-clad subjects. However, the lack of temporally consistent handling of plausible clothing dynamics, including fine and high-frequency details, significantly limits the attainable visual quality. We address these limitations for the first time in the literature and present a new framework which performs high-fidelity and temporally-consistent human motion transfer with natural pose-dependent non-rigid deformations, for several types of loose garments. In contrast to the previous techniques, we perform image generation in three subsequent stages: synthesizing human shape, structure, and appearance. Given a monocular RGB video of an actor, we train a stack of recurrent deep neural networks that generate these intermediate representations from 2D poses and their temporal derivatives. Splitting the difficult motion transfer problem into subtasks that are aware of the temporal motion context helps us to synthesize results with plausible dynamics and pose-dependent detail. It also allows artistic control of results by manipulation of individual framework stages. In the experimental results, we significantly outperform the state-of-the-art in terms of video realism. The source code is available at <https://graphics.tu-bs.de/publications/kappel2020high-fidelity>.

1. Introduction

Human motion transfer methods, also known as performance cloning or reenactment methods, can generate real-

istic video animations of an actor following a target motion specified by a user. This has several applications in AR/VR and video editing. Building upon new advances in machine learning, current motion transfer methods tackle this challenging problem by learning a direct mapping between an actor-independent motion space and the resulting target actor’s appearance space. These methods often require a training video of an actor performing a rich set of motions [2, 5, 17, 19, 27, 33].

Some recent motion transfer approaches parameterize motion as skeletal pose sequences that can be computed from videos with off-the-shelf pose detectors [5, 27]. Others use pre-captured template meshes or parameterized body models to provide additional guidance to the synthesis step [17, 19]. Acquisition of such templates [17], however, requires an extensive structure-from-motion reconstruction of the static target actor under constant lighting. Furthermore, existing human motion transfer approaches are likely to produce notable temporal and spatial artifacts when actors wear loose clothing, such as dresses, skirts and hoodies [2, 5, 19, 33]. On such garments, they struggle to realistically reproduce the appearance of fine-scale details like folds and wrinkles, as well as plausible dynamics.

In this paper, we present a new human motion transfer framework that generates visually plausible video animations of humans that are spatially and temporally coherent, and show natural dynamics, even for actors wearing loose garments (see Fig. 1). Given a single monocular video of an actor performing a rich set of motions, we train a stack of deep generative networks to learn a mapping from 2D pose to a silhouette with semantic part labels, and per-pixel

appearance of the actor. We model the person’s shape as a dense foreground silhouette mask with per-pixel labels encoding assignment to limbs and garments. We further encode the structure of wrinkles and texture patterns of garments as the orientation and strength of local image gradients. We extract this structure from images using a bank of oriented filter kernels [23, 31]. Encoding the actor’s appearance with these explicitly decoupled intermediate representations of silhouette and structure is key to enhance the temporal and spatial quality of synthesized videos comprising human actors in loose clothing.

Our method improves over current motion transfer approaches in terms of visual fidelity using a single RGB camera. Furthermore, our representation provides an additional level of control over the final image generation. For example, for the same overall dynamic geometric outline (*i.e.*, the same garment geometry), color and appearance, including fold and wrinkle style, can be manipulated in a purely image-based way. Overall, our contributions can be summarized as follows: (1) A new motion transfer framework with an emphasis on visually-plausible fine-scale deformations and dynamics in the actor’s clothing. (2) For this, we propose to decompose the pose-to-image translation task into better conditioned cascaded processes, where the final appearance is conditioned on the predicted shape outline and internal structure of the clothing. (3) We show that our intermediate representations do not only help to provide more temporally coherent conditioning resulting in more appealing image synthesis, but also allow controlling individual aspects of the final rendering (*e.g.*, enhance wrinkles and transfer the clothing style).

2. Related Work

Apart from general-purpose image-translation techniques which can be used for novel human view synthesis [2, 12, 32, 33, 37], most specialized techniques can be classified into still image-based [8, 21, 22, 25, 28] or video-based [1, 5, 7, 9, 17, 18, 19]. Pix2pix [12] is one of the most recognized general-purpose image translation techniques in the literature. Other image-to-image translation variants include pix2pixHD [33] for processing high-resolution images, Zhu *et al.* [37] for learning from unpaired data, and Wang *et al.* [32] for video processing. The work of Bansal *et al.* [2] learns a mapping from unpaired video data with a focus on processing temporal information. While general-purpose image-to-image translation techniques [2, 12, 32, 33, 37] can in principle be applied in our scenario, their visual accuracy usually suffers from multiple types of artifacts such as spurious details (*e.g.*, jittering texture details), unnatural deformations, and temporal inconsistency, among others.

Neural human motion transfer methods are specifically designed to generate novel views of a person observed in a

single RGB image. Still-image-based approaches [7, 8, 21, 22, 25, 28] focus is different from the problem we solve. Their focus is on re-targeting a person in just a still image and do not produce temporally consistent and visually pleasing video results. Video-based human motion transfer techniques [1, 5, 9, 17, 18, 19], however, focus on producing video results by handling temporal information. Aberman *et al.* [1] combine a skeleton representation with an image-translation network for human performance cloning. Additionally, they introduce an actor-independent training branch for unpaired data to enlarge the set of representable poses. Recently, Chan *et al.* [5] introduced a translation method for human bodies using an intermediate, subject independent pose representation. Inspired by pix2pix [12], they learn a pose-to-appearance mapping from video clips of target identities for human performance transfer. To improve temporal consistency and overall image quality, they include a temporal discriminator and dedicated face prediction network. While their method produces visually appealing results, it frequently generates unnatural and inconsistent deformations and texture details for loose clothing, as it is underconstrained from 2D pose keypoints.

Some video-based techniques aim to improve temporal and spatial fidelity by incorporating 3D information through human body meshes [9, 17, 18, 19]. L. Liu *et al.* [17, 18] presented high-quality human reenactments. As a proxy, they use an explicitly rigged, textured and skinned 3D model of the target actor. The underlying deformation model, however, does not allow to capture small wrinkles and local deformations. W. Liu *et al.* [19] use a human mesh recovery technique to disentangle poses from shapes as well as a correspondence map for human appearance transfer and novel view synthesis. They propagate input information both in the image and feature space and advocate a warping-based module for the enhanced preservation of the source information. Along the same lines is the work of Guan *et al.* [9] for human action transfer. It shows generalization to different target persons without retraining, thanks to a texture extraction method and a parametric human body model. However, the resolution of the supported renderings does not capture fine details and wrinkles. Textured Neural Avatars [27] can be trained for rendering a target actor in tight clothes in arbitrary body and camera poses. Proceeding from monocular videos with extracted 3D poses and regions of interest (human bodies), they train a convolutional neural network to predict a dense UV-map under a given target pose. Explicit decoupling of the geometry and texture results in higher visual quality and temporal consistency; yet, the overall representation is still comparably coarse and lacks fine details as well as local pose-dependent deformations and shading. Thus, no body renderings with loose clothes are demonstrated. While mesh-based techniques impose stronger human shape priors for higher qual-

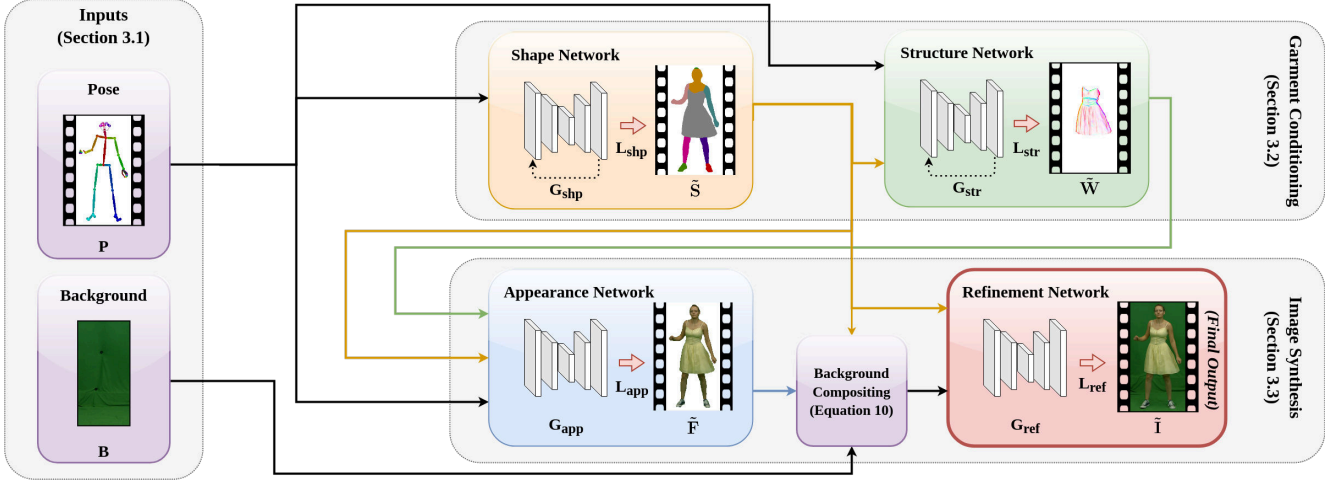


Figure 2. Framework Overview: Our framework synthesizes new output views of humans from a static background image and a pose input sequence extracted from the target (training) or source (testing) image sequence. Therefore, we apply two consecutive networks to explicitly pre-estimate human body parsing (*yellow*) and the internal gradient structure of clothing (*green*). The resulting outputs are used to condition final image generation, which is separated into the foreground (*blue*) and global (*red*) components to handle loose clothing in a temporally coherent way. Dashed lines indicate recurrent networks that feedback the last output to generate the next prediction.

ity results, they often assume water-tight clothing and a pre-captured template of the examined subject [9, 17, 18, 27]. This can lead to difficulties in capturing loosely-swinging garments or fine pose-dependent deformations. Furthermore, getting access to pre-captured templates can lead to additional assumptions about the scene (*e.g.*, that the target actor has to stay still for rigid structure-from-motion and brightness consistency for feature point matching) [17, 18].

3. Method

Our framework G_θ generates photorealistic videos of a target actor mimicking the motion of another source identity. Given only a single monocular RGB image sequence $I = (i_n)_{n=1}^N \in \mathbb{R}^{N \times 3 \times h \times w}$ containing N frames of width w and height h showing a target actor performing a rich set of motions, we extract an identity-independent 2D skeletal pose $P(i_n)$ for each frame and fit the network parameters θ to approximate the inverse mapping P^{-1} back to the original frames. Instead of performing the translation task within a single network, our framework can be thought of as a function composition of four generative neural networks (G_{ref} , G_{app} , G_{str} , G_{shp}) executed in a cascade fashion to progressively generate higher-level representations, as illustrated in Fig. 2. Next, we describe all individual framework components (Secs. 3.1-3.3), before providing details on the network architectures and implementation (Sec. 3.4).

3.1. Inputs

As inputs, our method takes the target image sequence I and a static image of the scene background $b \in \mathbb{R}^{3 \times h \times w}$

that is fused with the synthesized actor later on. From the given image sequence, we extract actor-independent pose representations $P : \mathbb{R}^{3 \times h \times w} \rightarrow H \times D$, where $H = \{0, 1\}^{C_1 \times h \times w}$ symbolizes the set of rasterized binary pose skeletons with C_1 input channels, and $D = \mathbb{R}^{C_2 \times h \times w}$ is the set of temporal derivatives for a single pose consisting of C_2 channels. Similar to recent performance cloning methods [5, 27], we apply an off-the-shelf pose estimator [3, 29, 35] that predicts 2D keypoints $k_n \in \mathbb{R}^{127 \times 2}$ for the body (including hands and face), to generate the skeleton $h_n \in H$ by connecting adjacent keypoints via binary lines. We further distribute $C_1 = 9$ limbs (face, head, torso, arms, legs, hands) over multiple channels, which helps the networks to distinguish between overlapping or symmetrical body parts.

However, estimating soft body deformations from a single pose remains highly ambiguous as the states depend on the temporal order of poses. Thus, we calculate a temporal pose context $d_n \in D$ as the first and second temporal derivatives of k_n with respect to the image index n in x and y direction, respectively ($C_2 = 2 \cdot 2 \cdot 9$), similar to velocities and accelerations in classical dynamics simulation:

$$k'_n = \left(\frac{\partial k_n}{\partial n}, \frac{\partial^2 k_n}{\partial n^2} \right). \quad (1)$$

We create rasterizations d_n from k'_n in the same way as for the pose skeletons, and linearly interpolate their values along the bones, similar to the depth representation of Shysheya *et al.* [27]. During the reenactment, we apply the same procedure to the source image sequence, but additionally perform pose normalization as described by Chan *et*

al. [5] for *inter-target* appearance transfer before drawing the skeletons. Finally, our pose conditioning $\mathbf{P}(i_n) = (h_n, d_n)$ is provided to our framework as the concatenation of the frame-wise pose skeleton and temporal context. We find that the described procedure results in a compact, yet expressive pose representation, while slightly outperforming a simple sliding window approach in our experiments.

3.2. Garment Conditioning

In our framework, we use two dedicated networks that pre-estimate the actor’s shape and the internal gradient structure of clothing in 2D space to improve the consistent modeling of deforming garments and the fall of the folds.

3.2.1 Shape Estimator

As a first step, our shape estimator \mathbf{G}_{shp} transforms the provided pose representations \mathbf{P} into the actor’s current silhouette

$$\tilde{s}_n^* = \mathbf{G}_{shp}(\mathbf{P}(i_n), \tilde{s}_{n-1}^*), \quad (2)$$

where $\tilde{s}_i^* \in \mathbb{R}^{j \times h \times w}$ is represented as semantic body-part segmentation with j labels. We extract a final segmentation mask $\tilde{s}_n \in ([\mathbb{N}]_1^j)^{h \times w}$ from our network by extracting the indices of maximum elements per pixel over the channel dimension.

As knowledge about the preceding shape is necessary to achieve temporal coherence of deforming garments, we design our networks in a recurrent fashion where each execution is conditioned on the previous output. Thus, \mathbf{G}_{shp} is tasked to estimate a delta in the observed shape based on the current pose and the contextual temporal derivatives. The resulting segmentation mask is a significant prior for the final appearance network to produce visually-plausible renderings, and it is further used to segment the actor and parts of their clothing in later stages. For training, we use the human parsing method of Li *et al.* [15] trained on the ATR dataset [16] to extract pseudo-ground-truth segmentation maps $S = (s_n)_{n=1}^n$ from I . This dataset provides $j = 18$ labels for all body limbs and common clothing, enabling our method to handle various clothing styles including shirts, pants, skirts, dresses and scarfs. We formulate our shape training loss \mathbf{L}_{shp} as the cross-entropy between the network prediction and the training label:

$$\mathbf{L}_{shp} = \mathbb{E}_{n \sim N} \log \left(\sum_j \exp(\tilde{s}_n^*(j)) \right) - \tilde{s}_n^*(s), \quad (3)$$

where $\tilde{s}_n^*(j)$ is the output channel for a given label j .

3.2.2 Structure Estimator

While our shape representation is sufficient prior to infer the appearance of solid body parts like arms and faces, it does

not provide information about wrinkles and folds within the clothing, that are needed to generate temporally consistent shading and texture patterns. Thus, we apply a second recurrent network \mathbf{G}_{str} that estimates the internal gradient structure $\tilde{w}_n \in \mathbb{R}^{2 \times h \times w}$ for clothing regions as indicated by the segmentation map \tilde{s}_n :

$$\tilde{w}_n = \mathbf{G}_{str}(\mathbf{P}(i_n), \tilde{s}_n, \tilde{w}_{n-1}). \quad (4)$$

We model the clothing structure as the pixelwise gradient direction and strength extracted from the responses of 32 oriented Gabor filters, similar to recent work on neural hair synthesis [31]. However, as a floating garment comprises sparser gradients than human hair, we do not only use the maximum filter responses to smooth the angle field, but append a normalized version to the smoothed gradient directions to model the probability of a wrinkle or texture change at a specific location, as visualized in Fig. 1 (color and saturation encode the gradient direction and the strength, respectively). We train the structure estimator using the L_1 distance between estimated \tilde{w}_n and ground-truth w_n structure using a mask for garment labels extracted from \tilde{s}_n :

$$\mathbf{L}_{str} = \mathbb{E}_{n \sim N} \chi_C(\tilde{s}_n) |\tilde{w}_n - w_n|, \quad (5)$$

where C denotes the set of segmentation labels corresponding to the actor’s clothing, and χ_C is the indicator function:

$$\chi_C(x) = \begin{cases} 1, & \text{if } x \in C, \\ 0, & \text{else.} \end{cases} \quad (6)$$

3.3. Image Synthesis

The second stage of our framework synthesizes the final output image based on the provided pose and garment conditionings. Again, we use two dedicated networks to independently generate the actor in the foreground and fuse it with the provided background image.

3.3.1 Appearance Network

Our first rendering network \mathbf{G}_{app} takes provided pose as well as the pre-estimated shape \tilde{s} and internal clothing structure \tilde{W} to synthesize the actor’s appearance $\tilde{f}_n \in \mathbb{R}^{3 \times h \times w}$:

$$\tilde{f}_n = \mathbf{G}_{app}(\mathbf{P}(i_n), \tilde{s}_n, \tilde{w}_n). \quad (7)$$

We train the appearance module using a combination of L_1 distance and perceptual reconstruction loss [13]. Again, an indicator function χ_B is used to mask the foreground pixels that are not assigned to the background label β :

$$\chi_B(x) = \begin{cases} 0, & \text{if } x = \beta, \\ 1, & \text{else.} \end{cases} \quad (8)$$

Thus, our appearance loss L_{app} reads

$$L_{app} = \mathbb{E}_{n \sim N} \chi_B(\tilde{s}_n) \left(\lambda_r |\tilde{f}_n - i_n| + \lambda_p |\phi(\tilde{f}_n) - \phi(i_n)| \right) \quad (9)$$

where $\phi(\cdot)$ denotes feature maps extracted from different layers of a pre-trained VGG19 network [30], and λ_r, λ_p are free hyperparameters for weighting.

3.3.2 Refinement Network

Finally, we apply a shallow refinement network to fuse the foreground prediction with the provided scene background. For this, we first paste the generated foreground \tilde{f}_n onto the static background image b using the masking function (8):

$$\tilde{i}_n^* = (\chi_B(\tilde{s}_n) \cdot \tilde{f}_n) + ((1 - \chi_B(\tilde{s}_n)) \cdot b). \quad (10)$$

Then, given foreground-background composition \tilde{i}_n^* and actor segmentation \tilde{s}_n , our refinement network G_{ref} performs simple transition smoothing and shadow generation to produce the final output image i_n :

$$\tilde{i}_n = G_{ref}(\tilde{i}_n^*, \tilde{s}_n). \quad (11)$$

We use a combination of structural and perceptual losses similar to our L_{app} (9), but over the entire image plane:

$$L_{ref} = \mathbb{E}_{n \sim N} \lambda_r |\tilde{i}_n - i_n| + \lambda_p |\phi(\tilde{i}_n) - \phi(i_n)|. \quad (12)$$

In contrast to recent pose-to-video translation methods [1, 5], we do not apply an adversarial loss for our generator network as we do not want to hallucinate high frequency details in the clothing based on statistics from the data sequence I , but rather encourage the network to stick to the predicted structure layout.

3.4. Implementation Details

We implement our framework in Python using PyTorch 1.5 [24]. Our four translation networks employ the local Pix2PixHD [33] generator architecture; we halve the number of residual blocks in our refinement network G_{ref} due to the similarity of the domains. We optimize our final loss

$$L_{total} = \lambda_1 L_{seg} + \lambda_2 L_{str} + \lambda_3 L_{app} + \lambda_4 L_{ref} \quad (13)$$

with hyperparameters ($\lambda_1 = 0.5, \lambda_{2,3,4} = 1.0, \lambda_r = 0.1, \lambda_p = 0.9$) in a sequential way for 30 epochs using Adam optimization [14] (momentum $\beta_1 = 0.999$ and $\beta_2 = 0.5$). On a single Quadro RTX 8000 GPU, training converges after approximately four days for $N \approx 23K$ frames, while processing a single testing-frame takes ~ 280 milliseconds at a resolution of 512×512 pixels. For more details, please refer to our official implementation and supplementary material.

4. Experiments

We compare our framework to several state-of-the-art approaches with subjects wearing different types of loose clothing with rich wrinkling and textures (Sec. 4.2), conduct a user study (Sec. 4.3) and perform an ablation study for individual framework stages (Sec. 4.4). Since our results improve upon the temporal consistency and fine texture details, they are best viewed in our supplementary video.

4.1. Datasets, Methods and Metrics

Existing datasets for the evaluation of recent appearance transfer methods usually comprise actors in monochromatic body-tight clothing, as loose garments are a frequent cause of visual artifacts and temporal irregularities. For our experiments, we capture new sequences of three subjects, each in up to three clothing styles with a length of seven to ten minutes at a resolution of 512×512 pixels, and select the last 5% of the frames for testing. Our dataset features various clothing styles including jeans, loose t-shirts, hoodies, and two dresses with rich wrinkle and texture patterns.

We compare our framework to current video-based appearance transfer approaches, *i.e.*, a general-purpose pix2pixHD network [33] that synthesizes an image from a rendering of the extracted pose, the Everybody-Dance-Now (EDN) method by Chan *et al.* [5], which adds a temporal discriminator and specialized face GAN, and RecycleGAN [2] that performs direct unsupervised retargeting between two videos. Additionally, we test our approach on the dataset of Liu *et al.* [17] comprising various actors in tight clothes. All results were generated using either the official implementations or data directly provided by the authors. Moreover, we have explored pose editing using NHRR [25] and Liquid Warping GAN [19] but the results showed poor visual quality on our sequences. Kindly note that these methods predict pixel correspondences to the SMPL mesh [20] using DensePose [10], and hence struggle with handling loose clothes. We provide exemplary results showing limitations of mesh-based approaches in our supplementary material. Furthermore, the code of Textured Neural Avatars [27] is, unfortunately, not publicly available, and we were not able to obtain results on our sequences directly from the authors. For quantitative analysis, we utilize the widely-used SSIM [34], LPIPS [36] and FID [11] metrics that assess the pixel-space structural similarity, perceptual distance based on neural network features, and the Fréchet distance between two data collections, respectively.

4.2. Results and Comparisons

We show exemplary results of every network component in Fig. 3. As can be seen, our framework is capable of estimating pronounced body part segmentations, the corresponding wrinkle patterns and resulting target actor appear-




Figure 3. Exemplary results of our motion transfer framework. The first column of each block illustrates the driving sequence, while the last one shows the final result. They anchor the outputs for the different intermediate networks (shape, structure, and appearance). For clothing gradient structure, we encode the direction and confidence as color and saturation, respectively, similar to optical flow visualization.

ance for all types of complex clothing in our dataset, independent of the source actor’s posture or appearance. We further compare our method to recent state-of-the-art approaches in Fig. 4. Therefore, we use a bidirectional reenactment scenario with two different types of clothing, *i.e.*, a loose monochromatic t-shirt with jeans, and a dress featuring a delicate arrangement of folds. While all of the tested methods are able to produce strong wrinkle patterns, our framework manages to achieve the most plausible temporal and spatial results due to the explicit handling of deforming garments. Please note that the temporal consistency can only be evaluated by watching our supplemental video. Also, we provide results and comparisons on the video sequences of Liu *et al.* [17] in our supplemental material.

We further assess the visual quality by numerically comparing our method to related pose-to-video translation approaches in Table 1. To calculate the metrics, we employ the target actor sequences shown in the first column of Fig. 4 using the corresponding self-reenactment scenarios, where ground-truth images are available from the testing sets. Recycle-GAN [2] focuses on learning from unpaired data and is upper-bounded by paired translation techniques

Metric		pix2pixHD [33]	EdN [5]	ours
SSIM [34]	↑	0.9288	0.9043	0.9358
LPIPS [36]	↓	0.0401	0.0303	0.0289
FID [11]	↓	21.7724	13.5421	17.9573

Table 1. Quantitative comparison to related methods. We assess the quantitative quality according to SSIM, LPIPS and FID metrics in a self-reenactment scenario.

when trained on paired data. Hence, we exclude it from our quantitative analysis. Our method outperforms the current state of the art according to the SSIM and LPIPS metrics, which perform a direct structural and perceptual comparison on a per-frame basis, indicating that the physical behaviour of clothing synthesized by our method better reflects ground truth. On the other hand, FID states that the pixel distribution generated by EDN is closer to the training set, presumably due to strong adversarial losses.

4.3. Perceptual Experiment

We find that quantitative metrics do not reveal all aspects of the visual video quality, and complement our evaluation by two user studies in which videos are graded according to human perception. We compare our method against recent image-based techniques using the video sequences shown in Fig. 4, including wide clothes with significant deformations. We also conduct an additional experiment on the dataset of Liu *et al.* [17] to evaluate on videos of actors in tight clothes. In both studies, we show the participants video reenactments of ~ 20 seconds in length with results of two methods at a time and ask them which video looks the most realistic. To measure both the performance of each method and the agreement between participants, we follow the linked-paired comparison design [6]. We rank each method according to the number of times they are preferred over the rest and perform a significance test of the votes differences [4, 26]. The total number of votes a method received in each experiment is displayed in Table 2. As can be seen, our method significantly outperforms recent image-based pose-to-video-translation methods in terms of video




Figure 4. Comparison to state-of-the-art motion transfer methods. We compare our method against pix2pixHD [33], EDN [5] and Recycle-GAN [2] on two sequences of an actress in a loose t-shirt and a dress. Please note that results of Recycle-GAN are only available at half of the original resolution, as training on full resolution produced highly implausible results.

realism for the clothing and is even preferred over Recycle-GAN, which uses a richer input domain for motion transfer. Our second experiment on the dataset of Liu *et al.* [17] further confirms that our monocular video-based approach also achieves considerably more realistic results than the related EDN method when exchanging motion between actors in tight clothing. However, in that scenario, image-based approaches cannot quite reach the quality of priors provided by expensively recorded template meshes, which excel for tight and static shapes (Liu *et al.* [17]). We provide more details on the experimental setup and evaluation in our supplementary material.

4.4. Ablation Study

We next conduct an ablation study to assess the influence of different framework components. Therefore, we retrain our method on one of our most challenging sequences featuring a swinging dress with complicated texture patterns, while incrementally dropping the garment conditioning net-

Exp.	Method	Input	#Votes	Ranking
1	Ours	Video (Pose)	257	1
	Recycle-GAN [2]	Video (RGB)	194	2
	EDN [5]	Video (Pose)	143	3
	pix2pixHD [33]	Video (Pose)	54	4
2	Liu [17]	Textured mesh + Video (Pose)	103	1
	Ours	Video (Pose)	76	2
	EDN [5]	Video (Pose)	13	3

Table 2. Perceptual ranking of the compared methods for our user studies. Experiment 1 was conducted online with 54 participants with various backgrounds, while experiment 2 was in-situ with 16 CG/CV experts. The rankings are statistically significant.




Figure 5. Ablation study evaluating the influence of our framework stages in a self-reenactment scenario. For an exemplary ground-truth frame, we show a plain appearance network conditioned on pose only (P), pose and shape (PS), and our full framework without (PSS-R) and with (PSS/ours) temporal recurrence.

works. In our study, we consider four framework variations, *i.e.*, the plain appearance network without additional inputs (P), the appearance conditioned on our shape estimates (PS) and our full framework including the internal structure (PSS/ours); additionally, we examine the effect of removing the recurrent back-feeding from all network modules (PSS-R). As shown in Fig. 5, we find that every component contributes to improving the final image quality.

While our plain pix2pix generator fails to produce sharp outlines of body parts and garments (P), providing a pre-estimated shape mask results in distinct transitions and more realistically-floating garments (PS). Adding information about the internal structure of the clothing further enhances the generation of high-frequency details within the given shape (PSS). This effect is also observable in quantitative evaluation using structural and perceptive metrics, as shown in Table 3. Furthermore, we observe that dropping the temporal back-feeding of previous network outputs (PSS-R)—while achieving similar quality on a per-frame basis—significantly decreases temporal consistency, as can be seen in our supplemental video.

Metric		P	PS	PSS-R	PSS (ours)
SSIM [34]	↑	0.8959	0.9045	0.9045	0.9046
LPIPS [36]	↓	0.044	0.036	0.036	0.035
FID [11]	↓	15.046	12.282	12.925	12.176

Table 3. Quantitative ablation analysis according to SSIM, LPIPS and FID metrics. We compare different versions of our framework in a self-reenactment scenario. The numbers indicate that every component contributes a structural and perceptual improvement to the synthesized output.



Figure 6. Applications: Our approach allows novel video editing through style transfer and wrinkle intensity manipulation. For style transfer, we use the source shape and/or structure while keeping the target appearance.

5. Application: Material Editing

On top of high-fidelity motion transfer, our framework offers an additional level of control over the image generation, as it incorporates explicit specifications of the actor’s shape and garment structure. This enables editing and manipulating the generated videos. For actors wearing similar clothing, pseudo-ground-truth segmentation or structure maps extracted from the source sequence can significantly improve the final image quality and temporal consistency. Alternatively, after training multiple subjects wearing clothing of the similar type, garment conditioning networks can be interchanged to generate clothing with new combinations of shape, structure and appearance. Furthermore, manually-edited or completely handcrafted shape and structure patterns can be used to influence the physical behavior of clothing. Fig. 6 shows two examples of automatic and manual material and appearance editing. We change the style of the actress’s dress and the actor’s top while keeping the original appearance by applying the shape (left), or both shape and structure, from the source sequence and feed it to the appearance network (middle). Also, magnifying the second channel of structure estimate (gradient confidence map) can be used to magnify or minify the fall of the folds during reenactment (right). While material property exchange is still currently limited to garments of comparable type or texture, the proposed approach is the first neural motion transfer method enabling controllable clothing composition for videos, which can pave the way for many applications in future (such as virtual try-on).

6. Discussion and Limitations

We have demonstrated high-fidelity results for multiple challenging human appearances with various deformation patterns of garments, and were able to manipulate several properties of the clothing during the reenactment. The ablation study has shown that all stages of our framework are justified and contribute to the final result. Although our

method is a step forward in the neural rendering of humans, it has limitations and can provoke further research. First, the space of depictable poses is currently determined by the composition and length of the training sequences. Next, inaccuracies and missing 3D information from the pose key-point estimator can result in incorrect occlusion maps and temporal behavior of garments. At the same time, the propagation of errors in the estimated human body parsing can lead to missing or disconnected body parts and unnatural transitions between the frames. These aspects can be improved in future by incorporating stronger constraints on the shape and body structure and enlarging the pose space by training on unpaired data.

7. Conclusion

We introduce a novel multi-stage framework for high-fidelity human motion transfer from monocular video. The tests show that our method performs qualitatively appealing motion transfer for different clothes, visually improving over previous image-based techniques. The performance is also confirmed by a comprehensive perceptual study, indicating superior temporal and visual quality compared to current image-based state-of-the-art motion transfer approaches. All in all, we conclude that explicitly handling deformations of clothing in a temporally-coherent way is crucial for high-fidelity neural human motion transfer.

Our supplementary material contains additional details on training, our garment conditioning representation/visualization, and experiments, as well as further comparison with other techniques. The full implementation is available at our project page <https://graphics.tu-bs.de/publications/kappel2020high-fidelity>.

Acknowledgements: This work was partially funded by the DFG (MA2555/15-1 “Immersive Digital Reality”) and the ERC Consolidator Grant 4DRepLy (770784). We thank Jalees Nehvi and Navami Kairanda for helping with comparisons.

References

- [1] Kfir Aberman, Mingyi Shi, Jing Liao, Dani Lischinski, Baoquan Chen, and Daniel Cohen-Or. Deep video-based performance cloning. *Comput. Graph. Forum*, 38:219–233, 2018.
- [2] Aayush Bansal, Shugao Ma, Deva Ramanan, and Yaser Sheikh. Recycle-gan: Unsupervised video retargeting. In *European Conference on Computer Vision (ECCV)*, pages 119–135, 2018.
- [3] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 43(1):172–186, 2021.
- [4] Susana Castillo, Tilke Judd, and Diego Gutierrez. Using eye-tracking to assess different image retargeting methods. In *Symposium on Applied Perception in Graphics and Visualization (APGV)*, page 7–14, 2011.
- [5] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *International Conference on Computer Vision (ICCV)*, 2019.
- [6] Herbert Aron David. *The method of paired comparisons*, volume 12 of *Griffin’s statistical monographs & courses*. Charles Griffin, London, 1 1963.
- [7] Patrick Esser, Ekaterina Sutter, and Björn Ommer. A variational u-net for conditional appearance and shape generation. In *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [8] Artur Grigorev, Artem Sevastopolsky, Alexander Vakhitov, and Victor Lempitsky. Coordinate-based texture inpainting for pose-guided human image generation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 12127–12136, 2019.
- [9] Shanyan Guan, Shuo Wen, Dexin Yang, Bingbing Ni, Wendong Zhang, Jun Tang, and Xiaokang Yang. Human action transfer based on 3D model reconstruction. In *AAAI Conference on Artificial Intelligence*, pages 8352–8359, 2019.
- [10] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Computer Vision and Pattern Recognition (CVPR)*, pages 7297–7306, 2018.
- [11] 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. In *Neural Information Processing Systems (NIPS)*, pages 6626–6637, 2017.
- [12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1125–1134, 2017.
- [13] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, pages 694–711, 2016.
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [15] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing. *arXiv preprint arXiv:1910.09777*, 2019.
- [16] Xiaodan Liang, Si Liu, Xiaohui Shen, Jianchao Yang, Luoqi Liu, Jian Dong, Liang Lin, and Shuicheng Yan. Deep human parsing with active template regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(12):2402–2414, 2015.
- [17] 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 Transactions on Visualization and Computer Graphics (TVCG)*, 2020.
- [18] 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 Transactions on Graphics (TOG)*, 2019.
- [19] 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 *International Conference on Computer Vision (ICCV)*, 2019.
- [20] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015.
- [21] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Neural Information Processing Systems (NIPS)*, 2017.
- [22] Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [23] Sylvain Paris, Will Chang, Oleg I. Kozhushnyan, Wojciech Jarosz, Wojciech Matusik, Matthias Zwicker, and Frédo Durand. Hair photobooth: Geometric and photometric acquisition of real hairstyles. *ACM Trans. Graph.*, 27(3), 2008.
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Neural Information Processing Systems (NIPS)*, pages 8024–8035, 2019.
- [25] Kripasindhu Sarkar, Dushyant Mehta, Weipeng Xu, Vladislav Golyanik, and Christian Theobalt. Neural re-rendering of humans from a single image. In *European Conference on Computer Vision (ECCV)*, 2020.
- [26] Iwan Setyawan and Reginald L. Lagendijk. Human perception of geometric distortions in images. In *Security, Steganography, and Watermarking of Multimedia Contents VI*, volume 5306, pages 256 – 267, 2004.

- [27] 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 *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [28] Aliaksandr Siarohin, Enver Sangineto, Stephane Lathuiliere, and Nicu Sebe. Deformable GANs for pose-based human image generation. In *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [29] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [31] Zhentao Tan, Menglei Chai, Dongdong Chen, Jing Liao, Qi Chu, Lu Yuan, Sergey Tulyakov, and Nenghai Yu. Michigan: multi-input-conditioned hair image generation for portrait editing. *ACM Transactions on Graphics (TOG)*, 39(4):95–1, 2020.
- [32] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Neural Information Processing Systems (NIPS)*, pages 1152–1164, 2018.
- [33] 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 *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [34] 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.
- [35] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [36] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [37] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *International Conference on Computer Vision (ICCV)*, 2017.