

Liquid Warping GAN: A Unified Framework for Human Motion Imitation, Appearance Transfer and Novel View Synthesis

Wen Liu^{1*} Zhixin Piao^{1*} Jie Min¹ Wenhan Luo² Lin Ma² Shenghua Gao¹

¹ShanghaiTech University ²Tencent AI Lab

{liuwen, piaozhx, minjie, gaoshh}@shanghaitech.edu.cn

{whluo.china, forest.linma}@gmail.com

Abstract

We tackle the human motion imitation, appearance transfer, and novel view synthesis within a unified framework, which means that the model once being trained can be used to handle all these tasks. The existing task-specific methods mainly use 2D keypoints (pose) to estimate the human body structure. However, they only express the position information with no abilities to characterize the personalized shape of the individual person and model the limbs rotations. In this paper, we propose to use a 3D body mesh recovery module to disentangle the pose and shape, which can not only model the joint location and rotation but also characterize the personalized body shape. To preserve the source information, such as texture, style, color, and face identity, we propose a Liquid Warping GAN with Liquid Warping Block (LWB) that propagates the source information in both image and feature spaces, and synthesizes an image with respect to the reference. Specifically, the source features are extracted by a denoising convolutional auto-encoder for characterizing the source identity well. Furthermore, our proposed method is able to support a more flexible warping from multiple sources. In addition, we build a new dataset, namely Impersonator (iPER) dataset, for the evaluation of human motion imitation, appearance transfer, and novel view synthesis. Extensive experiments demonstrate the effectiveness of our method in several aspects, such as robustness in occlusion case and preserving face identity, shape consistency and clothes details. All codes and datasets are available on <https://svip-lab.github.io/project/impersonator.html>.

1. Introduction

Human image synthesis, including human motion imitation [1, 19, 31], appearance transfer [26, 37] and novel

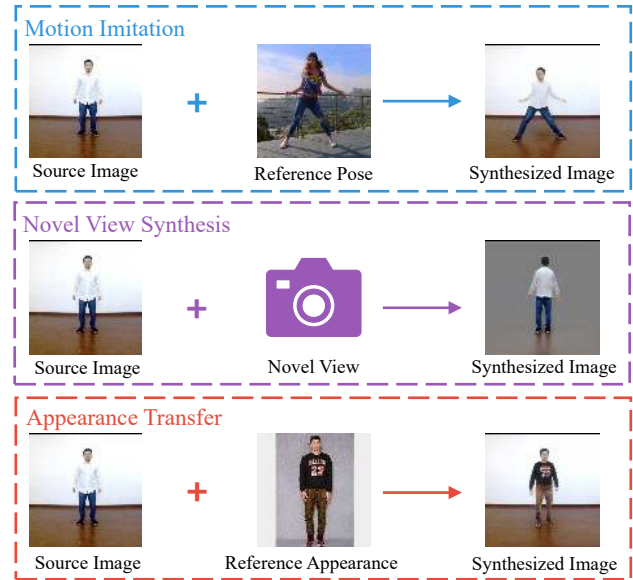


Figure 1. Illustration of human motion imitation, appearance transfer and novel view synthesis. The first column is the source image and the second column is reference condition, such as image or novel view of camera. The third column is the synthesized results.

view synthesis [40, 42], has huge potential applications in re-enactment, character animation, virtual clothes try-on, movie or game making and so on. The definition is that given a source human image and a reference human image, i) the goal of motion imitation is to generate an image with texture from source human and pose from reference human, as depicted in the top of Fig. 1; ii) human novel view synthesis aims to synthesize new images of the human body, captured from different viewpoints, as illustrated in the middle of Fig. 1; iii) the goal of appearance transfer is to generate a human image preserving reference identity with clothes, as shown in the bottom of Fig. 1 where different parts might come from different people.

In the realm of human image synthesis, previous works separately handle these tasks [19, 26, 42] with task-specific

*Contributed equally and work done while Wen Liu was a Research Intern with Tencent AI Lab.

pipeline, which seems to be difficult to extend to other tasks. Recently, generative adversarial network (GAN) [6] achieves great successes on these tasks. Taking human motion imitation as an example, we summarize recent approaches in Fig. 2. In an early work [19], as shown in Fig. 2 (a), source image (with its pose condition) and target pose condition are concatenated which thereafter is fed into a network with adversarial training to generate an image with desired pose. However, direct concatenation does not take the spatial layout into consideration, and it is ambiguous for the generator to place the pixel from source image into a right position. Thus, it always results in a blurred image and loses the source identity. Later, inspired by the spatial transformer networks (STN) [10], a texture warping method [1], as shown in Fig. 2 (b), is proposed. It firstly fits a rough affine transformation matrix from source and reference poses, uses an STN to warp the source image into reference pose and generates the final result based on the warped image. Texture warping, however, could not preserve the source information as well, in terms of the color, style or face identity, because the generator might drop out source information after several down-sampling operations, such as stride convolution and pooling. Meanwhile, contemporary works [4, 31] propose to warp the deep features of the source images into target pose rather than that in image space, as shown in Fig 2 (c), named as feature warping. However, features extracted by an encoder in feature warping cannot guarantee to accurately characterize the source identity and thus consequently produce a blur or low-fidelity image in an inevitable way.

The aforementioned existing methods encounter with challenges in generating unrealistic-looking images, due to three reasons: 1) diverse clothes in terms of texture, style, color, and high-structure face identity are difficult to be captured and preserved in their network architecture; 2) articulated and deformable human bodies result in a large spatial layout and geometric changes for arbitrary pose manipulations; 3) all these methods cannot handle multiple source inputs, such as in appearance transfer, different parts might come from different source people.

In this paper, to preserve the source information, including details of clothes and face identity, we propose a Liquid Warping Block (LWB) to address the loss of source information from three aspects: 1) a denoising convolutional auto-encoder is used to extract useful features that preserve source information, including texture, color, style and face identity; 2) source features of each local part are blended into a global feature stream by our proposed LWB to further preserve the source details; 3) it supports multiple-source warping, such as in appearance transfer, warping the features of head from one source and those of body from another, and aggregating into a global feature stream. This will further enhance the local identity of each source part.

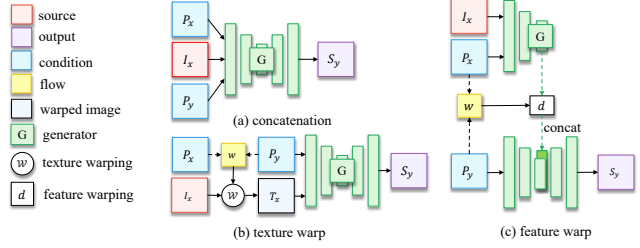


Figure 2. Three existing approaches of propagating source information into target condition. (a) is early concatenation, and it concatenates the source image and source condition, as well as target condition, into the color channel. (b) and (c) are texture and feature warping, respectively, and the source image or its features are propagated into target condition under a fitted transformation flow.

In addition, existing approaches mainly rely on 2D pose [1, 19, 31], dense pose [22] and body parsing [4]. These methods only take care of the layout locations and ignore the personalized shape and limbs (joints) rotations, which are even more essential than layout location in human image synthesis. For example, in an extreme case, a tall man imitates the actions of a short person and using the 2D skeleton, dense pose and body parsing condition will unavoidably change the height and size of the tall one, as shown in the bottom of Fig. 6. To overcome their shortcomings, we use a parametric statistical human body model, SMPL [2, 18, 12] which disentangles human body into pose (joint rotations) and shape. It outputs 3D mesh (without clothes) rather than the layouts of joints and parts. Further, transformation flows can be easily calculated by matching the correspondences between two 3D triangulated meshes, which is more accurate and results in fewer misalignments than previous fitted affine matrix from keypoints [1, 31].

Based on SMPL model and Liquid Warping Block (LWB), our method can be further extended into other tasks, including human appearance transfer and novel view synthesis for free and one model can handle these three tasks. We summarize our contributions as follows: 1) we propose a LWB to propagate and address the loss of the source information, such as texture, style, color, and face identity, in both image and feature space; 2) by taking advantages of both LWB and the 3D parametric model, our method is a unified framework for human motion imitation, appearance transfer, and novel view synthesis; 3) we build a dataset for these tasks, especially for human motion imitation in video, and all codes and datasets are released for further research convenience in the community.

2. Related Work

Human Motion Imitation. Recently, most methods are based on conditioned generative adversarial networks (CGAN) [1, 3, 19, 20, 22, 30] or Variational Auto-Encoder [5]. Their key technical idea is to combine target

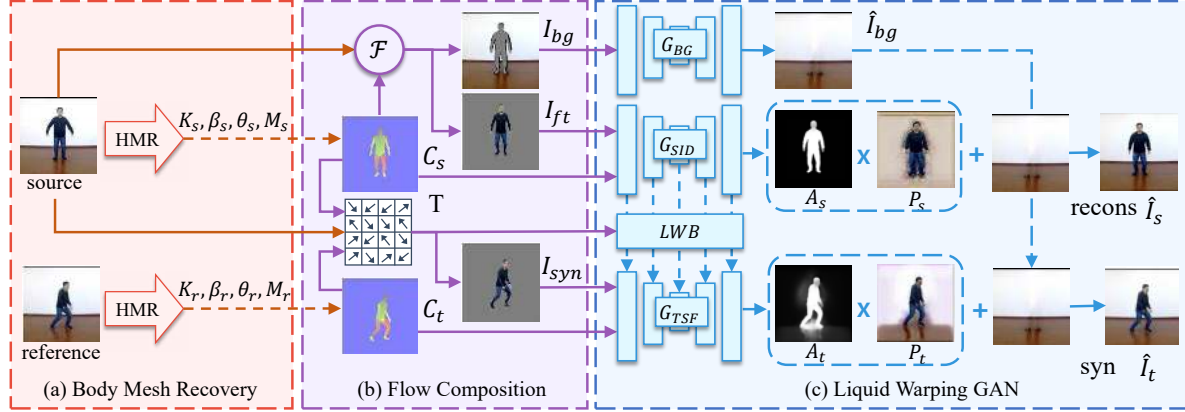


Figure 3. The training pipeline of our method. We randomly sample a pair of images from a video, denoting one of them as source image, named I_s and the other as reference image named I_r . (a) A body mesh recovery module will estimate the 3D mesh of each image, and render their correspondence map, C_s and C_t ; (b) The flow composition module will first calculate the transformation flow T based on two correspondence maps and their projected vertices in image space. Then it will separate source image I_s into foreground image I_{ft} and masked background I_{bg} . Finally it warps the source image based on transformation flow T , and produces a warped image I_{syn} ; (c) In the last GAN module, the generator consists of three streams, which separately generates the background image \hat{I}_{bg} by G_{BG} , reconstructs the source image \hat{I}_s by G_{SID} and synthesizes the target image \hat{I}_t under reference condition by G_{TSF} . To preserve the details of source image, we propose a novel Liquid Warping Block (LWB, shown in Fig. 4) which propagates the source features of G_{SID} into G_{TSF} at several layers and preserve the source information, in terms of texture, style and color.

image along with source pose (2D key-points) as inputs and generate realistic images by GANs using source pose. The difference of those approaches are merely in network architectures and adversarial losses. In [19], a U-Net generator is designed and a coarse-to-fine strategy is utilized to generate 256×256 images. Si *et al.* [1, 30] propose a multistage adversarial loss and separately generate the foreground (or different body parts) and background. Neverova *et al.* [22] replace the sparse 2D key-points with the dense correspondences between image and surface of the human body by DensePose [27]. Chan *et al.* [3] use pix2pixHD [35] framework together with a specialized Face GAN to learn a mapping from 2D skeleton to image and generate a more realistic target image. Furthermore, Wang *et al.* [34] extend it to video generation and Liu *et al.* [16] propose a neural renderer of human actor video. However, their works just train a mapping from 2D pose (or parts) to image of each person — in other words, every body need to train their own model. This shortcoming might limit its wide application.

Human Appearance Transfer. Human appearance modeling or transfer is a vast topic, especially in the field of virtual try-on applications, from computer graphics pipelines [24] to learning based pipelines [26, 37]. Graphics based methods first estimate the detailed 3D human mesh with clothes via garments and 3d scanners [38] or multiple camera arrays [15] and then human appearance with clothes is capable to be conducted from one person to another based on the detailed 3D mesh. Although these methods can produce high-fidelity result, their cost, size and controlled environment are unfriendly and inconvenient to customers. Recently, in the light of deep generative models,

SwapNet [26] firstly learns a pose-guided clothing segmentation synthetic network, and then the clothing parsing results with texture features from source image feed into an encoder-decoder network to generate the image with desired garment. In [37], the authors leverage a geometric 3D shape model combined with learning methods, swap the color of visible vertices of the triangulated mesh and train a model to infer that of invisible vertices.

Human Novel View Synthesis. Novel view synthesis aims to synthesize new images of the same object, as well as the human body, from arbitrary viewpoints. The core step of existing methods is to fit a correspondence map from the observable views to novel views by convolutional neural networks. In [41], the authors use CNNs to predict appearance flow and synthesize new images of the same object by copying the pixel from source image based on the appearance flow, and they have achieved decent results of rigid objects like vehicles. Following work [23] proposes to infer the invisible textures based on appearance flow and adversarial generative network (GAN) [6], while Zhu *et al.* [42] argue that appearance flow based method performs poorly on articulated and deformable objects, such as human bodies. They propose an *appearance-shape-flow* strategy for synthesizing novel views of human bodies. Besides, Zhao *et al.* [40] design a GAN based method to synthesize high-resolution views in a coarse-to-fine way.

3. Method

Our Liquid Warping GAN contains three stages, body mesh recovery, flow composition and a GAN module with

Liquid Warping Block (LWB). The training pipeline is the same for different tasks. Once the model has been trained on one task, it can deal with other tasks as well. Here, we use motion imitation as an example, as shown in Fig. 3. Denoting the source image as I_s and the reference image I_r . The first body mesh recovery module will estimate the 3D mesh of I_s and I_r , and render their correspondence maps, C_s and C_t . Next, the flow composition module will first calculate the transformation flow T based on two correspondence maps and their projected mesh in image space. The source image I_s is thereby decomposed as front image I_{ft} and masked background I_{bg} , and warped to I_{syn} based on transformation flow T . The last GAN module has a generator with three streams. It separately generates background image by G_{BG} , reconstructs the source image \hat{I}_s by G_{SID} and synthesizes the image \hat{I}_t under reference condition by G_{TSF} . To preserve the details of source image, we propose a novel Liquid Warping Block (LWB) and it propagates the source features of G_{SID} into G_{TSF} at several layers.

3.1. Body Mesh Recovery Module

As shown in Fig. 3 (a), given source image I_s and reference image I_r , the role of this stage is to predict the kinematic pose (rotation of limbs) and shape parameters, as well as 3D mesh of each image. In this paper, we use the HMR [12] as 3D pose and shape estimator due to its good trade-off between accuracy and efficiency. In HMR, an image is firstly encoded into a feature with \mathbb{R}^{2048} by a ResNet-50 [8] and then followed by an iterative 3D regression network that predicts the pose $\theta \in \mathbb{R}^{72}$ and shape $\beta \in \mathbb{R}^{10}$ of SMPL [18], as well as the weak-perspective camera $K \in \mathbb{R}^3$. SMPL is a 3D body model that can be defined as a differentiable function $M(\theta, \beta) \in \mathbb{R}^{N_v \times 3}$, and it parameterizes a triangulated mesh by $N_v = 6,890$ vertices and $N_f = 13,776$ faces with pose parameters $\theta \in \mathbb{R}^{72}$ and $\beta \in \mathbb{R}^{10}$. Here, shape parameters β are coefficients of a low-dimensional shape space learned from thousands of registered scans and the pose parameters θ are the joint rotations that articulate the bones via forward kinematics. With such process, we will obtain the body reconstruction parameters of source image, $\{K_s, \theta_s, \beta_s, M_s\}$ and those of reference image, $\{K_r, \theta_r, \beta_r, M_r\}$, respectively.

3.2. Flow Composition Module

Based on the previous estimations, we first render a correspondence map of source mesh M_s and that of reference mesh M_r under the camera view of K_s . Here, we denote the source and reference correspondence maps as C_s and C_t , respectively. In this paper, we use a fully differentiable renderer, Neural Mesh Renderer (NMR) [13]. We thereby project vertices of source V_s into 2D image space by weak-perspective camera, $v_s = Proj(V_s, K_s)$. Then, we calculate the barycentric coordinates of each mesh face, and

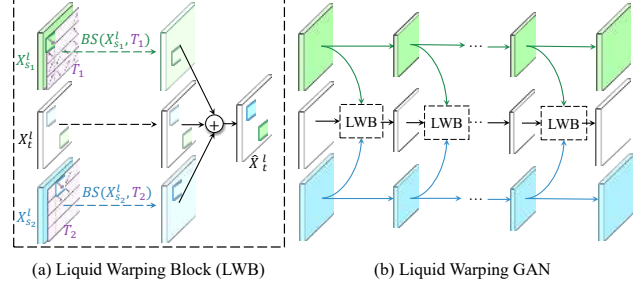


Figure 4. Illustration of Liquid Warping Block. (a) is the structure of LWB. X_{s1}^l and X_{s2}^l are the feature maps extracted by G_{SID} of different sources in l^{th} layers. X_t^l is the feature map of G_{TSF} at the l^{th} layer. Final output features \hat{X}_t^l aggregate the feature from G_{TSF} and warped source features by bilinear sampler (BS) with respect to the flow T_1 and T_2 . (b) is the architecture of LWB.

obtain $f_s \in \mathbb{R}^{N_f \times 2}$. Next, we calculate the transformation flow $T \in \mathbb{R}^{H \times W \times 2}$ by matching the correspondences between source correspondence map with its mesh face coordinates f_s and reference correspondence map. Here $H \times W$ is the size of image. Consequently, a front image I_{ft} and a masked background image I_{bg} are derived from masking the source image I_s based on C_s . Finally, we warp the source image I_s by the transformation flow T , and obtain the warped image I_{syn} , as depicted in Fig. 3.

3.3. Liquid Warping GAN

This stage synthesizes high-fidelity human image under the desired condition. More specifically, it 1) synthesizes the background image; 2) predicts the color of invisible parts based on the visible parts; 3) generates pixels of clothes, hairs and others out of the reconstruction of SMPL.

Generator. Our generator works in a three-stream manner. One stream, named G_{BG} , works on the concatenation of the masked background image I_{bg} and the mask obtained by the binarization of C_s in color channel (4 channels in total) to generate the realistic background image \hat{I}_{bg} , as shown in the top stream of Fig. 3 (c). The other two streams are source identity stream, namely G_{SID} and transfer stream, namely G_{TSF} . G_{SID} is a denoising convolutional auto-encoder which aims to guide the encoder to extract the features that are capable to preserve the source information. Together with the \hat{I}_{bg} , it takes the masked source foreground I_{ft} and the correspondence map C_s (6 channels in total) as inputs, and reconstructs source front image \hat{I}_s . G_{TSF} stream synthesizes the final result, which receives the warped foreground by bilinear sampler and the correspondence map C_t (6 channels in total) as inputs. To preserve the source information, such as texture, style and color, we propose a novel Liquid Warping Block (LWB) that links the source with target streams. It blends the source features from G_{SID} and fuses them into transfer stream G_{TSF} , as shown in the bottom of Fig. 3 (c).

One advantage of our proposed Liquid Warping Block

(LWB) is that it addresses multiple sources, such as in human appearance transfer, preserving the head of source one, and wearing the upper outer garment from the source two, while wearing the lower outer garment from the source three. The different parts of features are aggregated into G_{TSF} by their own transformation flow, independently. Here, we take two sources as an example, as shown in Fig. 4. Denoting $X_{s_1}^l$ and $X_{s_2}^l$ as the feature maps extracted by G_{SID} of different sources in the l^{th} layer. X_t^l is the feature map of G_{TSF} at the l^{th} layer. Each part of source feature is warped by their own transformation flow, and aggregated into the features of G_{TSF} . We use bilinear sampler (BS) to warp the source features $X_{s_1}^l$ and $X_{s_2}^l$, with respect to the transformation flows, T_1 and T_2 , respectively. The final output feature is obtained as follows:

$$\hat{X}_t^l = BS(X_{s_1}^l, T_1) + BS(X_{s_2}^l, T_2) + X_t^l.$$

Please note that we only take two sources as an example, which can be easily extended to multiple sources.

G_{BG} , G_{SID} and G_{TSF} have the similar architecture, named ResUnet, a combination of ResNet [7] and U-Net [28] without sharing parameters. For G_{BG} , we directly regress the final background image, while for G_{SID} and G_{TSF} , we concretely generate an attention map A and a color map P , as illustrated in Fig. 3 (c). The final image can be obtained as follows:

$$\begin{aligned}\hat{I}_s &= P_s * A_s + \hat{I}_{bg} * (1 - A_s) \\ \hat{I}_t &= P_t * A_t + \hat{I}_{bg} * (1 - A_t).\end{aligned}$$

Discriminator. For discriminator, we follow the architecture of Pix2Pix [9]. More details about our network architectures are provided in supplementary materials.

3.4. Training Details and Loss Functions

In this part, we will introduce the loss functions, and how to train the whole system. For body recovery module, we follow the network architecture and loss functions of HMR [12]. Here, we use a pre-trained model of HMR.

For the Liquid Warping GAN, in the training phase, we randomly sample a pair of images from each video and set one of them as source I_s , and another as reference I_r . Note that our proposed method is a unified framework for motion imitation, appearance transfer and novel view synthesis. Therefore once the model has been trained, it is capable to be applied to other tasks and does not need to train from scratch. In our experiments, we train a model for motion imitation and then apply it to other tasks, including appearance transfer and novel view synthesis.

The whole loss function contains four terms and they are perceptual loss [11], face identity loss, attention regularization loss and adversarial loss.

Perceptual Loss. It regularizes the reconstructed source image \hat{I}_s and generated target image \hat{I}_t to be closer to the

ground truth I_s and I_r in VGG [32] subspace. Its formulation is given as follows:

$$L_p = \|f(\hat{I}_s) - f(I_s)\|_1 + \|f(\hat{I}_t) - f(I_r)\|_1.$$

Here, f is a pre-trained VGG-19 [32].

Face Identity Loss. It regularizes the cropped face from the synthesized target image \hat{I}_t to be similar to that from the image of ground truth I_r , which pushes the generator to preserve the face identity. It is shown as follows:

$$L_f = \|g(\hat{I}_t) - g(I_r)\|_1.$$

Here, g is a pre-trained SphereFaceNet [17].

Adversarial Loss. It pushes the distribution of synthesized images to the distribution of real images. As shown in following, we use $LSGAN_{-110}$ [21] loss in a way like PatchGAN for the generated target image \hat{I}_t . The discriminator D regularizes \hat{I}_t to be more realistic-looking. We use conditioned discriminator, and it takes generated images and the correspondence map C_t (6 channels) as inputs.

$$L_{adv}^G = \sum D(\hat{I}_t, C_t)^2.$$

Attention Regularization Loss. It regularizes the attention map A to be smooth and to prevent them from saturating. Considering that there is no ground truth of attention map A , as well as color map P , they are learned from the resulting gradients of above losses. However, the attention masks can easily saturate to 1 which prevents the generator from working. To alleviate this situation, we regularize the mask to be closer to silhouettes S rendered by 3D body mesh. Since the silhouettes is a rough map and it contains the body mask without clothes and hair, we also perform a Total Variation Regularization over A like [25], to compensate the shortcomings of silhouettes, and further to enforce smooth spatial color when combining the pixel from the predicted background \hat{I}_{bg} and the color map P . It is shown as follows:

$$L_a = \|A_s - S_s\|_2^2 + \|A_t - S_t\|_2^2 + TV(A_s) + TV(A_t)$$

$$TV(A) = \sum_{i,j} [A(i,j) - A(i-1,j)]^2 + [A(i,j) - A(i,j-1)]^2.$$

For generator, the full objective function is shown in the following, and λ_p , λ_f and λ_a are the weights of perceptual, face identity and attention losses.

$$L^G = \lambda_p L_p + \lambda_f L_f + \lambda_a L_a + L_{adv}^G.$$

For discriminator, the full objective function is

$$L^D = \sum [D(\hat{I}_t, C_t) + 1]^2 + \sum [D(I_r, C_t) - 1]^2.$$

3.5. Inference

Once trained model on the task of motion imitation, it can be applied to other tasks in inference. The difference lies in the transformation flow computation, due to the different conditions of various tasks. The remaining modules, Body Mesh Recovery and Liquid Warping GAN modules are all the same. Followings are the details of each task of Flow Composition module in testing phase.

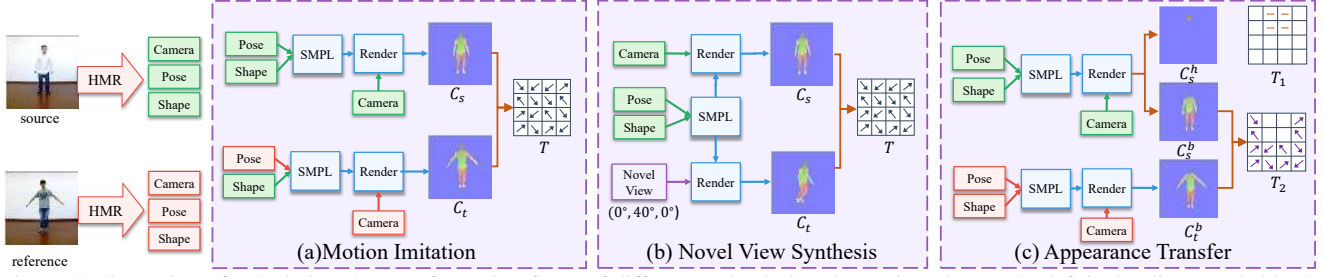


Figure 5. Illustration of calculating the transformation flows of different tasks during the testing phase. The left is the disentangled body parameters by Body Recovery module of both source and reference images. The right is the different implementations to calculate the transformation flow in different tasks.

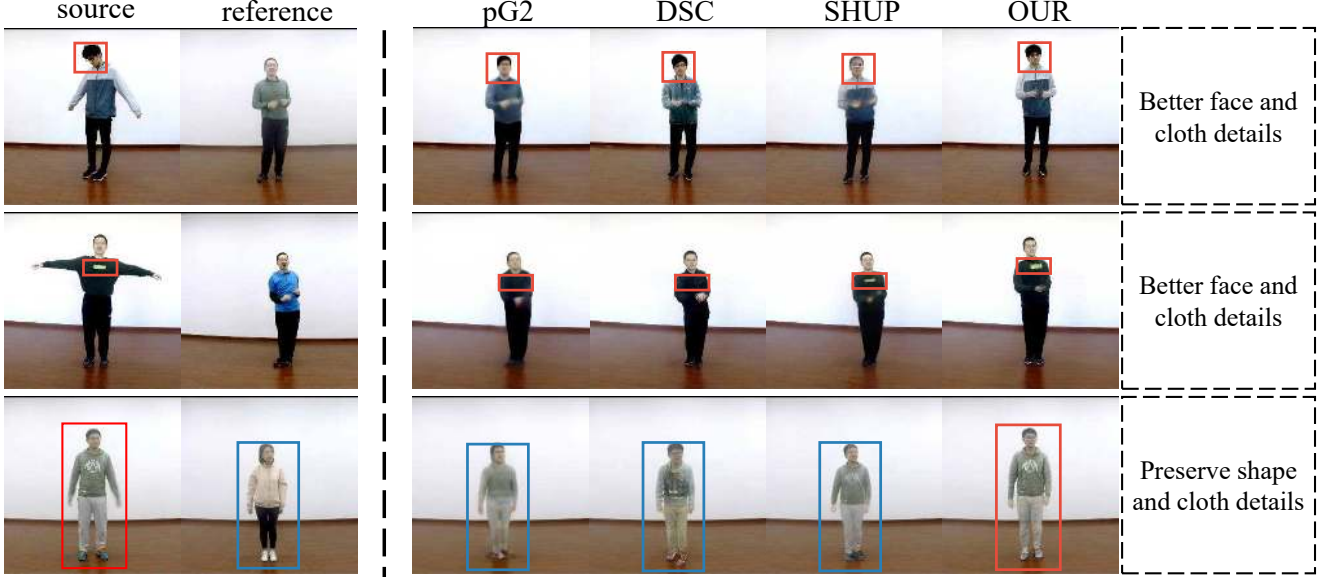


Figure 6. Comparison of our method with others of motion imitation on the iPER dataset (zoom-in for the best of view). 2D pose-guided methods pG2 [19], DSC [31] and SHUP [1] cannot preserve the clothes details, face identity and shape consistency of source images. We highlight the details by red and blue rectangles.

Motion Imitation. We firstly copy the value of pose parameters of reference θ_r into that of source, and get synthetic parameters of SMPL, as well as the 3D mesh, $M_t = M(\theta_r, \beta_s)$. Next, we render a correspondence map of source mesh M_s and that of synthetic mesh M_t under the camera view of K_s . Here, we denote the source and synthetic correspondence map as C_s and C_t , respectively. Then, we project vertices of source into 2D image space by weak-perspective camera, $v_s = Proj(V_s, K_s)$. Next, we calculate the barycentric coordinates of each mesh face, and have $f_s \in \mathbb{R}^{N_f \times 2}$. Finally, we calculate the transformation flow $T \in \mathbb{R}^{H \times W \times 2}$ by matching the correspondences between source correspondence map with its mesh face coordinates f_s and synthetic correspondence map. This procedure is shown in Fig. 5 (a).

Novel View Synthesis. Given a new camera view, in terms of rotation R and translation t . We firstly calculate the 3D mesh under the novel view, $M_t = M_s R + t$. The flowing operations are similar to motion imitation. We render a correspondence map of source mesh M_s and that of

novel mesh M_t under the weak-perspective camera K_s and calculate the transformation flow $T \in \mathbb{R}^{H \times W \times 2}$ in the end. This is illustrated in Fig. 5 (b).

Appearance Transfer. It needs to “copy” the clothes of torso or body from the reference image while keeping the head (face, eye, hair and so on) identity of source. We split the transformation flow T into two sub-transformation flow, source flow T_1 and referent flow T_2 . Denoting head mesh as $M^h = (V^h, F^h)$ and body mesh as $M^b = (V^b, F^b)$. Here, $M = M^h \cup M^b$. For T_1 , We firstly project the head mesh M_s^h of source into image space, and thereby obtain the silhouettes, S_s^h . Then, we create a mesh grid, $G \in \mathbb{R}^{H \times W \times 2}$. Then, we mask G by S_s^h , and derive $T_1 = G * S_s^h$. Here, $*$ represents element-wise multiplication. For T_2 , it is similar to motion imitation. We render the correspondence map of source body M_s^b and that of reference M_t^b , denoting as C_s^b and C_t^b , respectively. Finally, we calculate the transformation flow T_2 based on the correspondences between C_s^b and C_t^b . We illustrate it in Fig. 5 (c).

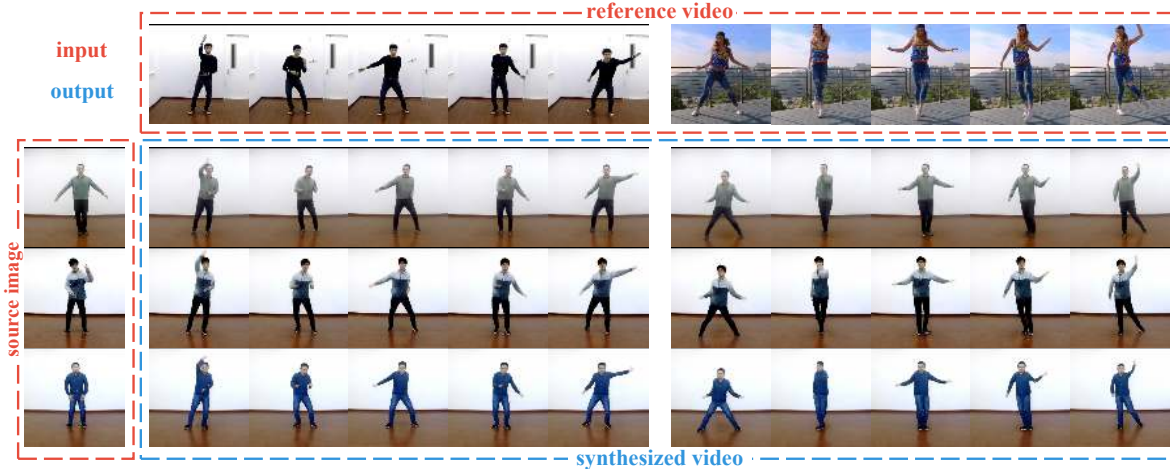


Figure 7. Examples of motion imitation from our proposed methods on the iPER dataset (zoom-in for the best of view). Our method could produce high-fidelity images that preserve the face identity, shape consistency and clothes details of source, even there are occlusions in source images such as the middle and bottom rows. We recommend accessing the supplementary material for more results in videos.

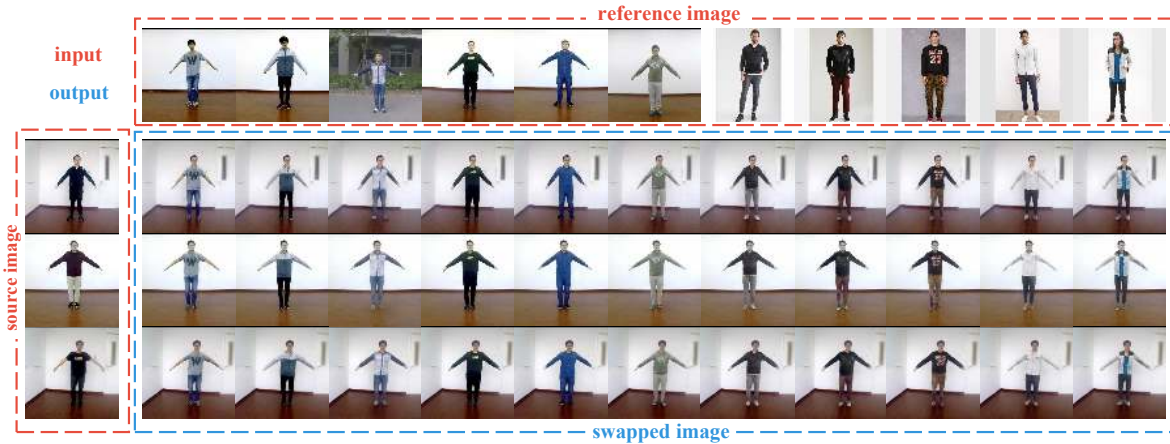


Figure 8. Examples of our method of human appearance transfer in the testing set of iPER (zoom-in for the best of view). Our method could produce high-fidelity and decent images that preserve the face identity and shape consistency of the source image, and keep the clothes details of reference image. We recommend accessing the supplementary material for more results.

4. Experiments

Dataset. To evaluate the performances of our proposed method of motion imitation, appearance transfer and novel view synthesis, we build a new dataset with diverse styles of clothes, named as Impersonator (iPER) dataset. There are 30 subjects of different conditions of shape, height and gender. Each subject wears different clothes and performs an A-pose video and a video with random actions. Some subjects might wear multiple clothes, and there are 103 clothes in total. The whole dataset contains 206 video sequences with 241,564 frames. We split it into training/testing set at the ratio of 8:2 according to the different clothes.

Implementation Details. To train the network, all images are normalized to $[-1, 1]$ and resized to 256×256 . We randomly sample a pair of images from each video. The mini-batch size is 4 in our experiments. λ_p , λ_f and λ_a are set to 10.0, 5.0 and 1.0, respectively. Adam [14] is used for parameter optimization of both generator and discriminator.

4.1. Evaluation of Human Motion Imitation.

Evaluation Metrics. We propose an evaluation protocol of testing set of iPER dataset and it is able to indicate the performance of different methods in terms of different aspects. The details are listed in followings: 1) In each video, we select three images as source images (frontal, sideways and occlusive) with different degrees of occlusion. The frontal image contains the most information, while the sideways will drop out some information, and occlusive image will introduce ambiguity. 2) For each source image, we perform self-imitation that actors imitate actions from themselves. SSIM [36] and Learned Perceptual Similarity (LPIPS) [39] are the evaluation metrics in self-imitation setting. 3) Besides, we also conduct cross-imitation that actors imitate actions from others. We use Inception Score (IS) [29] and Fréchet Distance on a pre-trained person-reid model [33], named as FReID, to evaluate the quality of generated images.

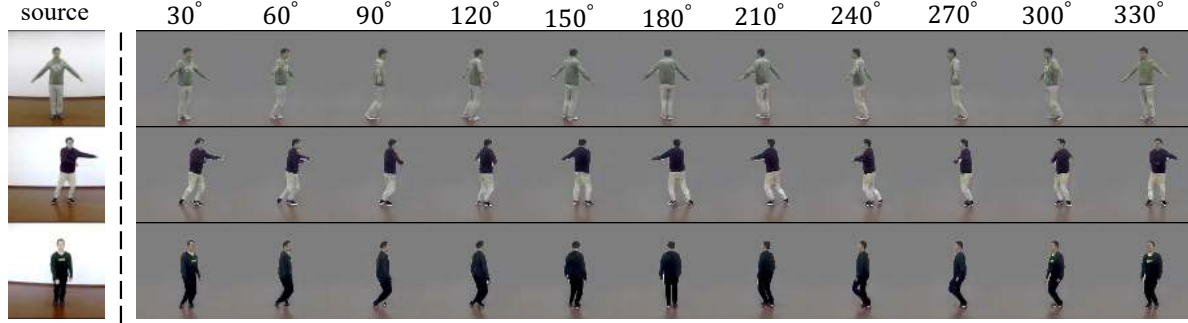


Figure 9. Examples of novel view synthesis from our method on the iPER dataset (zoom-in for the best of view). Our method could generate realistic-looking results under different camera views, and it is capable to preserve the source information, even in the self-occlusion case, such as the middle and bottom rows.

Table 1. Results of motion imitation by different methods on iPER dataset. \uparrow means the larger is better, and \downarrow represents the smaller is better. A higher SSIM may not mean a better quality of image [39].

	Self-Imitation		Cross-Imitation	
	SSIM \uparrow	LPIPS \uparrow	IS \uparrow	FRID \downarrow
PG2 [19]	0.854	0.865	3.242	0.353
SHUP [1]	0.832	0.901	3.371	0.324
DSC [31]	0.829	0.871	3.321	0.342
W_C	0.821	0.872	3.213	0.341
W_T	0.822	0.887	3.353	0.347
W_F	0.830	0.897	3.358	0.325
Ours-W_{LWB}	0.840	0.913	3.419	0.317

Comparison with Other Methods. We compare the performance of our method with that of existing methods, including PG2 [19], SHUP [1] and DSC [31]. We train all these methods on iPER dataset, and the evaluation protocol mentioned above is applied to these methods. The results are reported in Table 1. It can be seen that our method outperforms other methods. In addition, we also analyze the generated images and make comparisons between ours and above methods. From Fig. 6, we find that 1) the 2D pose-guided methods, including PG2 [19], SHUP [1] and DSC [31], change the body shape of source. For example, in the 3rd row of Fig. 6, a tall person imitates motion from a short person and these methods change the height of source body. However, our method is capable to keep the body shape unchanged. 2) When source image exhibits self-occlusion, such as invisible face in the 1st row of Fig. 6, our method could generate more realistic-looking content of the ambiguous and invisible parts. 3) Our method is more powerful in terms of preserving source identity, such as the face identity and cloth details of source than other methods, as shown in the 2nd and 3rd row of Fig. 6. 4) Our method also produces high-fidelity images in the cross-imitation setting (imitating actions from others) and we illustrate it in Fig. 7.

Ablation Study. To verify the impact of our proposed Liquid Warping Block (LWB), we design three baselines with aforementioned ways to propagate the source information, including early concatenation, texture warping and feature warping. All modules and loss functions are the

same except the propagating strategies among our method and other baselines. Here, we denote early concatenation, texture warping, feature warping, and our proposed LWB as W_C , W_T , W_F and W_{LWB} . We train all these under the same setting on the iPER dataset, then evaluate their performances on motion imitation. From Table 1, we can see that our proposed LWB is better than other baselines. More details are provided in supplementary materials.

4.2. Results of Human Appearance Transfer.

It is worth emphasizing that once model has been trained, it is able to directly to be applied in three tasks, including motion imitation, appearance transfer and novel view synthesis. We randomly pick some examples displayed in Fig. 8. The face identity and clothes details, in terms of texture, color and style, are preserved well by our method. It demonstrates that our method can achieve decent results in appearance transfer, even when the reference image comes from Internet and is out of the domain of iPER dataset, such as the last five columns in Fig. 8.

4.3. Results of Human Novel View Synthesis.

We randomly sample source images from the testing set of iPER, and change the views from 30° to 330°. The results are illustrated in Fig. 9. Our method is capable to predict reasonable content of invisible parts when switching to other views and keep the source information, in terms of face identity and clothes details, even in the self-occlusion case, such as the middle and bottom rows in Fig. 9.

5. Conclusion

We propose a unified framework to handle human motion imitation, appearance transfer, and novel view synthesis. It employs a body recovery module to estimate the 3D body mesh which is more powerful than 2D Pose. Furthermore, in order to preserve the source information, we design a novel warping strategy, Liquid Warping Block (LWB), which propagates the source information in both image and feature spaces, and supports a more flexible warping from multiple sources. Extensive experiments show that our framework outperforms others and produce decent results.

References

- [1] Guha Balakrishnan, Amy Zhao, Adrian V. Dalca, Frdo Durand, and John Guttag. Synthesizing images of humans in unseen poses. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [2] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578. Springer, 2016.
- [3] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. *arXiv preprint arXiv:1808.07371*, 2018.
- [4] Haoye Dong, Xiaodan Liang, Ke Gong, Hanjiang Lai, Jia Zhu, and Jian Yin. Soft-gated warping-gan for pose-guided person image synthesis. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 472–482, 2018.
- [5] Patrick Esser, Ekaterina Sutter, and Björn Ommer. A variational u-net for conditional appearance and shape generation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8857–8866, 2018.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778, 2016.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, pages 630–645, 2016.
- [9] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976, 2017.
- [10] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2017–2025, 2015.
- [11] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II*, pages 694–711, 2016.
- [12] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [13] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3907–3916, 2018.
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, volume abs/1412.6980, 2015.
- [15] Vincent Leroy, Jean-Sébastien Franco, and Edmond Boyer. Multi-view dynamic shape refinement using local temporal integration. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 3113–3122, 2003.
- [16] 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. *ACM Transactions on Graphics 2019 (TOG)*, 2019.
- [17] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6738–6746, 2017.
- [18] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, oct 2015.
- [19] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Advances in Neural Information Processing Systems*, pages 405–415, 2017.
- [20] Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [21] Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, Zhen Wang, and Stephen Paul Smolley. On the effectiveness of least squares generative adversarial networks. *CoRR*, abs/1712.06391, 2017.
- [22] Natalia Neverova, Rıza Alp Güler, and Iasonas Kokkinos. Dense pose transfer. In *European Conference on Computer Vision (ECCV)*, 2018.
- [23] Eunbyung Park, Jimei Yang, Ersin Yumer, Duygu Ceylan, and Alexander C. Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [24] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J. Black. Clothcap: seamless 4d clothing capture and retargeting. *ACM Trans. Graph.*, 36(4):73:1–73:15, 2017.
- [25] Albert Pumarola, Antonio Agudo, Aleix M. Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: Anatomically-aware facial animation from a single image. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part X*, pages 835–851, 2018.

- [26] Amit Raj, Patsorn Sangkloy, Huiwen Chang, James Hays, Duygu Ceylan, and Jingwan Lu. Swapnet: Image based garment transfer. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XII*, pages 679–695, 2018.
- [27] Iasonas Kokkinos Rıza Alp Güler, Natalia Neverova. Densepose: Dense human pose estimation in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [28] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015 - 18th International Conference Munich, Germany, October 5 - 9, 2015, Proceedings, Part III*, pages 234–241, 2015.
- [29] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2226–2234, 2016.
- [30] Chenyang Si, Wei Wang, Liang Wang, and Tieniu Tan. Multistage adversarial losses for pose-based human image synthesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [31] Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuillière, and Nicu Sebe. Deformable gans for pose-based human image generation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 2015*.
- [33] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and A strong convolutional baseline). In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, pages 501–518, 2018.
- [34] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Nikolai Yakovenko, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 1152–1164, 2018.
- [35] 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [36] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Processing*, 13(4):600–612, 2004.
- [37] Mihai Zanfir, Alin-Ionut Popa, Andrei Zanfir, and Cristian Sminchisescu. Human appearance transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [38] Chao Zhang, Sergi Pujades, Michael J. Black, and Gerard Pons-Moll. Detailed, accurate, human shape estimation from clothed 3d scan sequences. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5484–5493, 2017.
- [39] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [40] Bo Zhao, Xiao Wu, Zhi-Qi Cheng, Hao Liu, Zequn Jie, and Jiashi Feng. Multi-view image generation from a single-view. In *2018 ACM Multimedia Conference on Multimedia Conference, MM 2018, Seoul, Republic of Korea, October 22-26, 2018*, pages 383–391, 2018.
- [41] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A. Efros. View synthesis by appearance flow. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, pages 286–301, 2016.
- [42] Hao Zhu, Hao Su, Peng Wang, Xun Cao, and Ruigang Yang. View extrapolation of human body from a single image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.