

Pose with Style: Detail-Preserving Pose-Guided Image Synthesis with Conditional StyleGAN

BADOUR ALBAHAR, Virginia Tech, USA and Kuwait University, Kuwait
JINGWAN LU, Adobe Research, USA
JIMEI YANG, Adobe Research, USA
ZHIXIN SHU, Adobe Research, USA
ELI SHECHTMAN, Adobe Research, USA
JIA-BIN HUANG, Virginia Tech, USA and University of Maryland College Park, USA

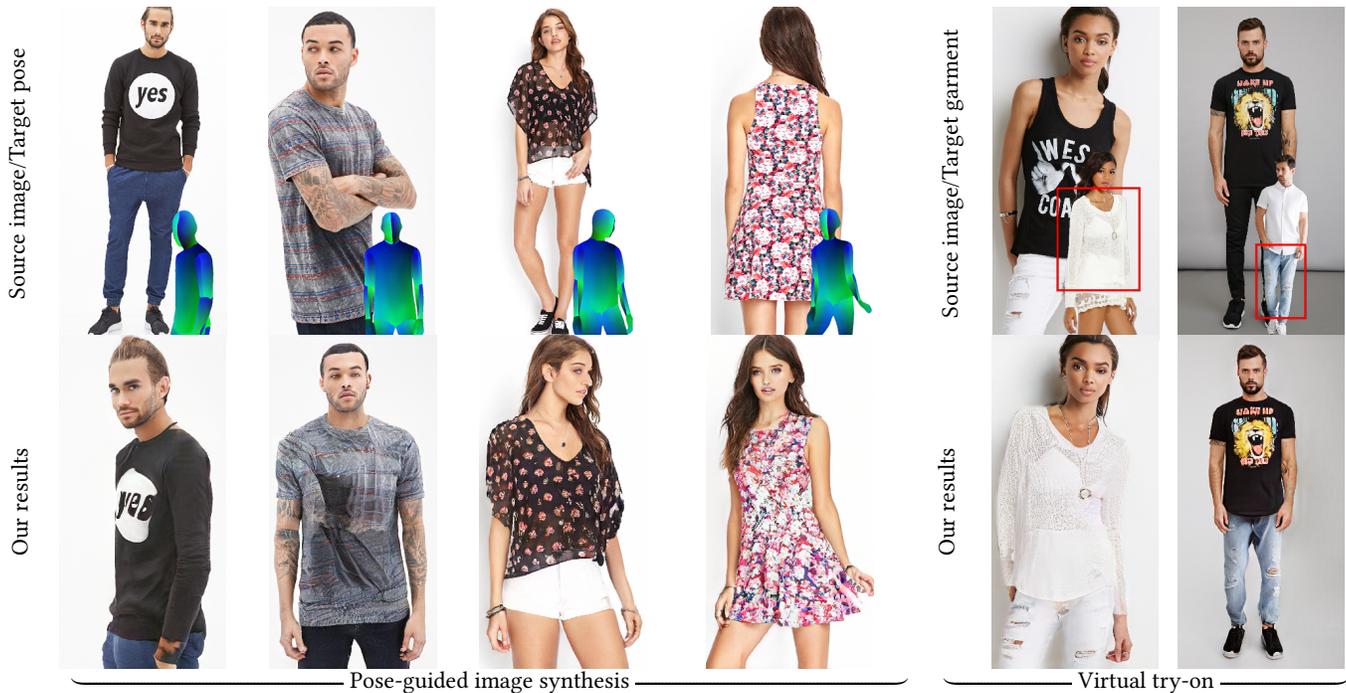


Fig. 1. **Detail-preserving pose-guided person image generation.** We present a single-image human reposing algorithm guided by arbitrary body shapes and poses. Our method first transfers local appearance features in the source image to the target pose with a human body symmetry prior. We then leverage a pose-conditioned StyleGAN2 generator with spatial modulation to produce photo-realistic reposing results. Our work enables applications of posed-guided synthesis (*left*) and virtual try-on (*right*). Thanks to spatial modulation, our result preserves the texture details of the source image better than prior work.

We present an algorithm for re-rendering a person from a single image under arbitrary poses. Existing methods often have difficulties in hallucinating occluded contents photo-realistically while preserving the identity and fine details in the source image. We first learn to inpaint the correspondence field between the body surface texture and the source image with a human body symmetry prior. The inpainted correspondence field allows us to transfer/warp local features extracted from the source to the target view even

under large pose changes. Directly mapping the warped local features to an RGB image using a simple CNN decoder often leads to visible artifacts. Thus, we extend the StyleGAN generator so that it takes pose as input (for controlling poses) and introduces a spatially varying modulation for the latent space using the warped local features (for controlling appearances). We show that our method compares favorably against the state-of-the-art algorithms in both quantitative evaluation and visual comparison.

Authors' addresses: Badour AlBahar, Virginia Tech, USA, Kuwait University, Kuwait, badour@vt.edu; Jingwan Lu, Adobe Research, USA, jlu@adobe.com; Jimei Yang, Adobe Research, USA, jimyang@adobe.com; Zhixin Shu, Adobe Research, USA, zshu@adobe.com; Eli Shechtman, Adobe Research, USA, elishe@adobe.com; Jia-Bin Huang, Virginia Tech, USA, University of Maryland College Park, USA, jbh Huang@vt.edu.

CCS Concepts: • **Computing methodologies;**

ACM Reference Format:

Badour AlBahar, Jingwan Lu, Jimei Yang, Zhixin Shu, Eli Shechtman, and Jia-Bin Huang. 2021. Pose with Style: Detail-Preserving Pose-Guided Image Synthesis with Conditional StyleGAN. *ACM Trans. Graph.* 40, 6, Article 218 (December 2021), 11 pages. <https://doi.org/10.1145/3478513.3480559>

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3478513.3480559>.

1 INTRODUCTION

Controllable, photo-realistic human image synthesis has a wide range of applications, including virtual avatar creation, reposing, virtual try-on, motion transfer, and view synthesis. Photo-realistic rendering of human images is particularly challenging through traditional computer graphics pipelines because it involves 1) designing or capturing 3D geometry and appearance of human and garments, 2) controlling poses via skeleton-driven deformation of 3D shape, and 3) synthesizing complicated wrinkle patterns for loose clothing. Recent learning-based approaches alleviate these challenges and have shown promising results. These methods typically take inputs 1) a single source image capturing the human appearance and 2) a target pose representation (part confidence maps, skeleton, mesh, or dense UV coordinates) and synthesize a novel human image with the appearance from source and the pose from the target.

Image-to-image translation based methods [Esser et al. 2018; Ma et al. 2017, 2018; Pumarola et al. 2018; Siarohin et al. 2018], building upon conditional generative adversarial networks [Isola et al. 2017], learn to predict the reposed image from the source image and the target pose. However, as human reposing involves significant spatial transformations of appearances, such approaches often require per-subject training using multiple images from the same persons [Chan et al. 2019; Wang et al. 2018a] or are incapable of preserving the person’s identity and the fine appearance details of the clothing in the source image.

Surface-based approaches [Grigorev et al. 2019; Lazova et al. 2019; Neverova et al. 2018; Sarkar et al. 2020] map human pixels in the source image to the canonical 3D surface of the human body (e.g., SMPL model [Loper et al. 2015]) with part segmentation and UV parameterization. This allows transferring pixel values (or local features) of visible human surfaces in the input image to the corresponding spatial location specified by the target pose. These methods thus retain finer-grained local details and identity compared to image-to-image translation models. However, modeling human appearance as a single UV texture map cannot capture view/pose-dependent appearance variations and loose clothing.

StyleGAN-based methods [Lewis et al. 2021; Men et al. 2020; Sarkar et al. 2021] very recently have shown impressive results for controllable human image synthesis [Men et al. 2020; Sarkar et al. 2021] or virtual try-on [Lewis et al. 2021]. The key ingredient is to extend the unconditioned StyleGAN network [Karras et al. 2020] to a *pose-conditioned* one. While their generated images are photo-realistic, it remains challenging to preserve fine appearance details (e.g., unique patterns/textures of garments) in the source image due to the global (spatially-invariant) modulation/demodulation of latent space.

We present a new algorithm for generating *detail-preserving* and *photo-realistic* re-rendering of human with novel poses from a *single source image*. Similar to the concurrent work [Lewis et al. 2021; Sarkar et al. 2021], we use a pose-conditioned StyleGAN network for generating pose-guided images. To preserve fine-grained details in the source image, we learn to inpaint the correspondence field between 3D body surface and the source image using a body symmetry prior. Using this inpainted correspondence field, we transfer local features from the source to the target pose and use the warped local features to modulate the StyleGAN generator network

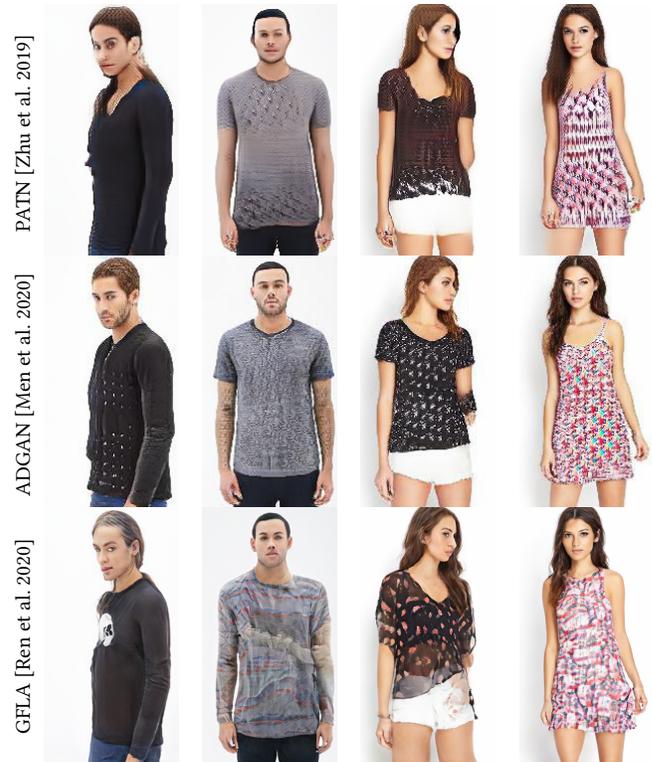


Fig. 2. **Limitations of existing methods.** Existing human reposing methods struggle to preserve details in the source image. Common issues include identity (1st and 2nd columns) and clothing textures (3rd, 4th columns) changes. Compare these results with ours in Figure 1.

at multiple StyleBlocks in a *spatially varying* manner. As we combine complementary techniques of photo-realistic image synthesis (from StyleGAN-based methods) and the 3D-aware detail transfers (from surface-based methods), our method achieves high-quality human-reposing and garment transfer results (Figure 1) and alleviates visible artifacts compared with the state-of-the-art (Figure 2). While StylePoseGAN [Sarkar et al. 2021] (concurrent work to ours) also combines pose-conditioned StyleGAN with the use of proxy geometry, its global modulation/demodulation scheme limits its ability to preserve fine appearance details. We evaluate the proposed algorithm visually and quantitatively using the DeepFashion dataset [Liu et al. 2016] and show favorable results compared to the current best-performing methods. *Our contributions* include:

- We integrate the techniques from surface-based and styleGAN-based methods to produce *detail-preserving* and *photo-realistic* controllable human image synthesis.
- We propose an explicit *symmetry prior* of the human body for learning to inpaint the correspondence field between human body surface and the source image which facilitates detail transfer, particularly for drastic pose changes.
- We present a *spatially varying* variant of latent space modulation in the StyleGAN network, allowing us to transfer local details while preserving photo-realism.

2 RELATED WORK

2.1 Pose-guided Person Image Synthesis

Pose-guided person image synthesis aims to transfer a person’s appearance from a source image to a target pose. Example applications include motion transfer [Aberman et al. 2019; Chan et al. 2019; Yoon et al. 2021], human reposing [Ma et al. 2017, 2018; Men et al. 2020], and virtual try-on [Lewis et al. 2021; Men et al. 2020; Sarkar et al. 2021]. These approaches typically encode the pose as either part confidence maps [Aberman et al. 2019; Ma et al. 2017, 2018] or skeleton [Chan et al. 2019; Esser et al. 2018; Men et al. 2020; Pumarola et al. 2018; Ren et al. 2020; Siarohin et al. 2018; Zhu et al. 2019] and use a conditional GAN to produce the reposed images. To handle large pose changes, existing methods leverage per-subject training [Aberman et al. 2019; Chan et al. 2019; Wang et al. 2018a], spatial transformation/deformation [Balakrishnan et al. 2018; Ren et al. 2020; Siarohin et al. 2018], and local attention [Ren et al. 2020]. To retain the source identity and appearance, surface-based methods first establish the correspondence between pixels from the source/target image to a canonical coordinate system of the 3D human body (with UV parameterization). These methods can then transfer pixels [Alldieck et al. 2019; Grigorev et al. 2019; Lazova et al. 2019; Neverova et al. 2018] or local features [Sarkar et al. 2020] to the target pose. As the commonly used UV parameterization only captures the surface of a tight human body [Loper et al. 2015], methods either explicitly predict garment labels [Yoon et al. 2021] or implicitly re-render the warped features [Sarkar et al. 2020]. Very recently, pose-conditioned StyleGAN networks have been proposed [Lewis et al. 2021; Sarkar et al. 2021]. To control the target appearance, pose-independent UV texture [Sarkar et al. 2021] is used to modulate the latent space. Our method builds upon pose-conditioned StyleGAN but differs from prior work in two critical aspects. First, instead of *global* latent feature modulation used in prior work, we propose to use a *spatially varying* modulation for improved local detail transfer. Second, we train a coordinate inpainting network for completing partial correspondence field (between the body surface and source image) using a human body symmetry prior. This allows us to directly transfer local features extracted from the source to the target pose.

2.2 Neural Rendering

Neural rendering methods first render a coarse RGB image or neural textures that is then mapped to an RGB image using a translation network [Kim et al. 2018; Liu et al. 2020; Meshry et al. 2019; Raj et al. 2021; Tewari et al. 2020c; Thies et al. 2019]. Recent research focuses on learning volumetric neural scene representations for view synthesis [Lombardi et al. 2019; Mildenhall et al. 2020]. This has been extended to handle dynamic scenes (e.g., humans) [Gao et al. 2021; Li et al. 2021; Park et al. 2021; Treitschk et al. 2021; Xian et al. 2021]. Recent efforts further enable controls over viewpoints [Gafni et al. 2021; Gao et al. 2020], pose [Noguchi et al. 2021; Peng et al. 2021], expressions [Gafni et al. 2021], illumination [Zhang et al. 2021] of human face/body. However, most of these approaches often require computationally expensive *per-scene/per-person training*. Our method also uses body surface mesh as our geometry proxy for re-rendering. Instead of using a simple CNN translation network,

we integrate the rendered latent texture with StyleGAN through spatially varying modulation. In contrast to volumetric neural rendering techniques, our method does *not* require per-subject training.

2.3 Deep Generative Adversarial Networks

Deep generative adversarial networks have shown great potentials for synthesizing high-quality photo-realistic images [Brock et al. 2019; Goodfellow et al. 2014; Karras et al. 2018, 2020; Zhang et al. 2019]. Using the pre-trained model, several works discover directions in the latent space that correspond to spatial or semantic changes [Härkönen et al. 2020; Jahanian et al. 2020; Peebles et al. 2020; Shen and Zhou 2021; Shoshan et al. 2021]. In the context of portrait images, some recent methods provide 3D control for the generated samples [Abdal et al. 2021; Tewari et al. 2020b] or real photographs [Tewari et al. 2020a]. Our work focuses on designing a pose-conditioned GAN with precise control on the localized appearance (for virtual try-on) and pose (for reposing).

2.4 Image-to-Image Translation

Image-to-image translation provides a general framework for mapping an image from one visual domain to another [Isola et al. 2017; Park et al. 2019; Wang et al. 2018b]. Recent advances include learning from unpaired dataset [Huang et al. 2018; Lee et al. 2018; Zhu et al. 2017], extension to videos [Wang et al. 2019, 2018a], and talking heads [Wang et al. 2021; Zakharov et al. 2019]. Similar to many existing human reposing methods [Chan et al. 2019; Ma et al. 2017, 2018; Men et al. 2020; Pumarola et al. 2018; Ren et al. 2020; Zhu et al. 2019], our work maps an input target pose to an RGB image with the appearance from a source image. Our core technical novelties lie in 1) spatial modulation in StyleGAN for detail transfer and 2) a body symmetry prior for correspondence field inpainting.

2.5 Localized Manipulation

When editing images, localized manipulation is often preferable over global changes. Existing work addresses this via structured noise [Alharbi and Wonka 2020], local semantic latent vector discovery [Chai et al. 2021], latent space regression [Collins et al. 2020], and explicit masking [Shocher et al. 2020]. Our spatial feature modulation shares high-level similarity with approaches that add spatial dimensions to the latent vectors in unconditional StyleGAN [Alharbi and Wonka 2020; Kim et al. 2021] and conditional GANs [AlBahar and Huang 2019; Park et al. 2019]. Our work differs in that our spatial modulation parameters are predicted from the warped appearance features extracted from the source image instead of being generated from random noise using a mapping network.

2.6 Symmetry Prior

Symmetry prior (in particular reflective symmetry) has been applied for learning deformable 3D objects [Wu et al. 2020], 3D reconstruction of objects [Sinha et al. 2012], and human pose regression [Yeh et al. 2019]. Our work applies left-right reflective symmetry to facilitate the training of coordinate-based texture inpainting network. The symmetry prior allows us to reuse local appearance features from the source and leads to improved results when source and target poses are drastically different.

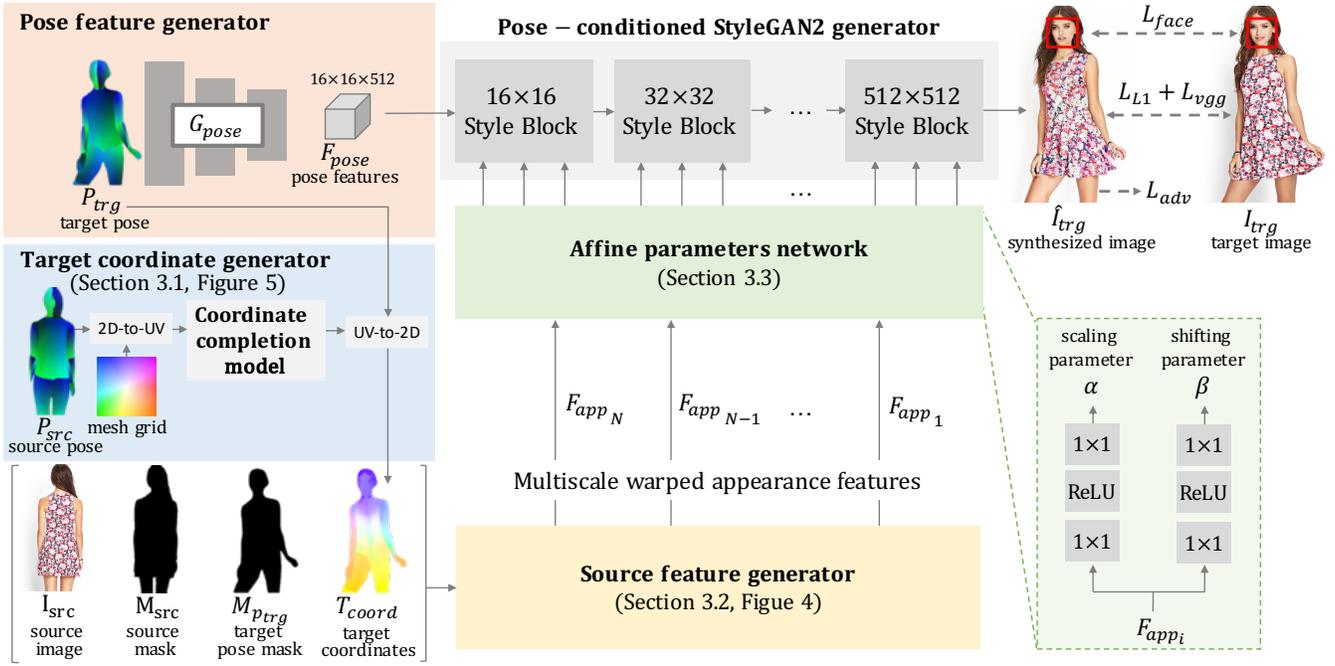


Fig. 3. **Method overview.** Our human reposing model builds upon a pose-conditioned StyleGAN2 generator [Karras et al. 2020]. We extract the DensePose [Güler et al. 2018] representation P_{trg} and use a pose encoder G_{pose} to encode P_{trg} into $16 \times 16 \times 512$ pose features F_{pose} which is used as input to the StyleGAN2 generator [Karras et al. 2020]. To preserve the source image appearance, we encode the input source image I_{src} into multiscale warped appearance features F_{app_i} using the source feature generator (Figure 4). To warp the feature from the source pose to the target pose we use the target coordinates T_{coord} . We compute these target coordinates T_{coord} using 1) the target dense pose P_{trg} and 2) the completed coordinates in the UV-space inpainted using the coordinate completion model (Figure 5). We pass the multi-scale warped appearance features F_{app_i} through the affine parameters network to generate scaling and shifting parameters α and β that are used to modulate the StyleGAN2 generator features in a *spatially varying* manner (Figure 7). Our training losses include adversarial loss, reconstruction losses, and a face identity loss.

3 METHOD

Given an image of a person I_{src} and a desired target pose P_{trg} represented by Image-space UV coordinate map per body part (shortly IUUV) extracted from DensePose [Güler et al. 2018], our goal is to generate an image preserving the appearance of the person in I_{src} in the desired pose P_{trg} . Note that this IUUV representation of dense pose entangles both the pose and shape representation.

We show an overview of our proposed approach in Figure 3. We use a pose-guided StyleGAN2 generator [Karras et al. 2020] that takes $16 \times 16 \times 512$ pose features F_{pose} as input. The pose features F_{pose} are encoded from the DensePose representation [Güler et al. 2018] using a pose feature generator G_{pose} that is composed of several residual blocks. Using the source and target pose, we use coordinate completion model to produce target coordinate that establishes the correspondence between target and source image. To encode the appearance information, we use a feature pyramid network [Lin et al. 2017] G_{app} to encode the source image into multiscale features and warp them according to the target pose. We then use the warped appearance features to generate scaling and shifting parameters to spatially modulate the latent space of the StyleGAN generator.

3.1 Coordinate Completion Model

The IUUV map of the source pose P_{src} allows us to represent the *pose-independent* appearance of the person in the UV-space. However, only the appearance of *visible* body surface can be extracted. This leads to incomplete UV-space appearance representation and thus may not handle the dis-occluded appearance for the target pose P_{trg} . Previous work [Sarkar et al. 2021] encodes the partial UV-space appearance to a *global latent vector* for modulating the generator. This works well for clothing with uniform colors or homogeneous textures, but inevitably loses the spatially-distributed appearance details. We propose to inpaint the UV-space appearance by a neural network guided by the human body mirror-symmetry prior. Instead of directly inpainting pixel values in UV-space, we choose to complete the mapping from image-space to UV-space established by P_{src} and represented by UV-space source image coordinates, in order to avoid generating unwanted appearance artifacts while best preserving the source appearance. We refer to this network as *coordinate completion model*. We show an overview of our coordinate completion model in Figure 5.

Given a mesh grid and the dense pose of the input source image P_{src} , we use a pre-computed image-space to/from UV-space mapping to map coordinates from the mesh grid to appropriate locations in the UV-space (using bilinear sampling for handling fractional

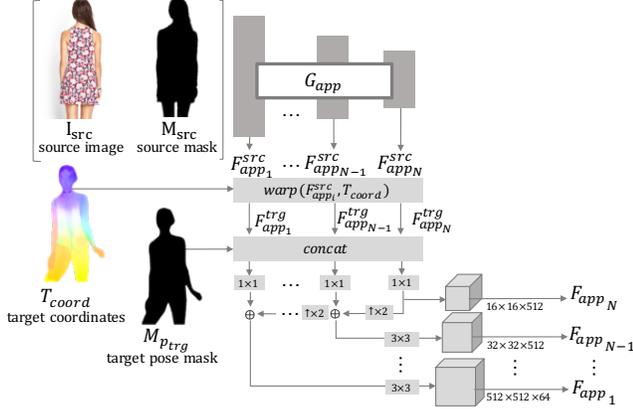


Fig. 4. **Source feature generator.** To preserve the source image appearance, we encode the input source image I_{src} into multiscale features $F_{app_i}^{src}$ and warp them from the source pose to the target pose $F_{app_i}^{trg}$ using the target coordinates T_{coord} computed from the target coordinate generator (Figure 3). We further process the warped features with a feature pyramid network [Lin et al. 2017] to obtain the multi-scale warped appearance features F_{app_i} which go through affine parameters network to generate scaling and shifting parameters α and β that are used to modulate the StyleGAN2 generator features in a *spatially varying* manner (Figure 7).

coordinates). We denote these base mapped coordinates as C_{base} and the mask indicating where these coordinates are as M_{base} .

Since human appearances are often left-right symmetrical, in addition to these base coordinates, we also map the left-right mirrored coordinates to the UV-space $C_{mirrored}$ and denote their respective mask as $M_{mirrored}$. Visualization of mapped base and mirrored coordinates are shown in Figure 6.

We combine the incomplete UV-space base and mirrored coordinates and their respective masks, such that:

$$M_{mirrored} = M_{mirrored} - (M_{base} \cdot M_{mirrored}) \quad (1)$$

$$M_{in} = M_{base} + M_{mirrored} \quad (2)$$

$$C_{in} = C_{base} \cdot M_{base} + C_{mirrored} \cdot M_{mirrored} \quad (3)$$

We concatenate the combined coordinates C_{in} and their mask M_{in} and pass them as input to the coordinate completion model. To implement our coordinate completion model, we follow a similar architecture to the coordinate inpainting architecture proposed by [Grigorev et al. 2019] with gated convolutions [Yu et al. 2019].

We train our model to minimize the ℓ_1 loss between the generated coordinates C_{out} and the input coordinates, such that:

$$L_{coord} = \|C_{out} \cdot M_{base} - C_{base} \cdot M_{base}\|_1 + \lambda_{mirrored} \cdot \|C_{out} \cdot M_{mirrored} - C_{mirrored} \cdot M_{mirrored}\|_1, \quad (4)$$

where $\lambda_{mirrored}$ is set to 0.5.

We also utilize the source-target pairs to train the coordinate completion model. Specifically, we use the source dense pose P_{src} to map the generated complete coordinates from the UV-space to the source image-space S_{coord} . Similarly, we also use the target dense pose P_{trg} to map the generated complete coordinates from the UV-space to the target image-space T_{coord} using the pre-computed mapping table. We then use these target and source coordinates to

warp pixels from the input source image I_{src} and minimize the ℓ_1 loss between the foreground of the warped images and the foreground of the ground truth images, such that:

$$L_{rgb} = \|warp(I_{src}, S_{coord}) \cdot M_{p_{src}} - I_{src} \cdot M_{p_{src}}\|_1 + \|warp(I_{src}, T_{coord}) \cdot M_{p_{trg}} - I_{trg} \cdot M_{p_{trg}}\|_1, \quad (5)$$

where $M_{p_{src}}$ and $M_{p_{trg}}$ are the source pose mask and target pose mask, respectively.

The total loss to train the coordinate completion model is:

$$L = L_{coord} + L_{rgb} \quad (6)$$

3.2 Source Feature Generator

To preserve the appearance in source image I_{src} , we encode it using several residual blocks into multi-scale features $F_{app_i}^{src}$. We utilize the pretrained coordinate completion model to obtain the target image-space coordinates T_{coord} such that it could warp the source features $F_{app_i}^{src}$ from the source pose to the target pose $F_{app_i}^{trg}$. We then concatenate these warped features with the target dense pose mask $M_{p_{trg}}$ and pass them into a feature pyramid network [Lin et al. 2017] to get our multi-scale warped appearance features F_{app_i} . We show our source feature generator in Figure 4.

3.3 Affine Parameters Network and Spatial Modulation

Prior to every convolution layer in each style block of StyleGAN2, we pass the warped source features F_{app_i} into an affine parameters network to generate scaling α , and shifting β parameters. Each convolution layer has its own independent affine parameters network which is composed of two 1×1 convolutions separated by a ReLU activation function for each parameter. To preserve spatial details, we modify every convolution layer in each style block of StyleGAN2 [Karras et al. 2020]. Instead of performing spatially invariant weight modulation and demodulation, we use the generated scaling and shifting tensor parameters, α and β , to perform spatially varying modulation of the features F_i as follows:

$$\hat{F}_i = \alpha_i \otimes F_i + \beta_i, \quad (7)$$

where \hat{F}_i is now the modulated features that will be passed as input to the 3×3 convolution of styleGAN2 generator. The output features of the convolution \hat{F}_i^{out} is then normalized to zero mean and unit standard deviation. Such that:

$$\bar{F}_i = \frac{\hat{F}_i^{out} - \text{mean}(\hat{F}_i^{out})}{\text{std}(\hat{F}_i^{out})} \quad (8)$$

Similar to StyleGAN2, we then add the noise broadcast operation B and the bias to \bar{F}_i to get F_{i+1} which will be fed into the next convolution layer. In Figure 7, we illustrate our detailed spatial modulation approach as well as the non-spatial weight modulation and demodulation of StyleGAN2.

3.4 Training Losses

In addition to StyleGAN2 adversarial loss L_{adv} , we train our model to minimize the following reconstruction losses:

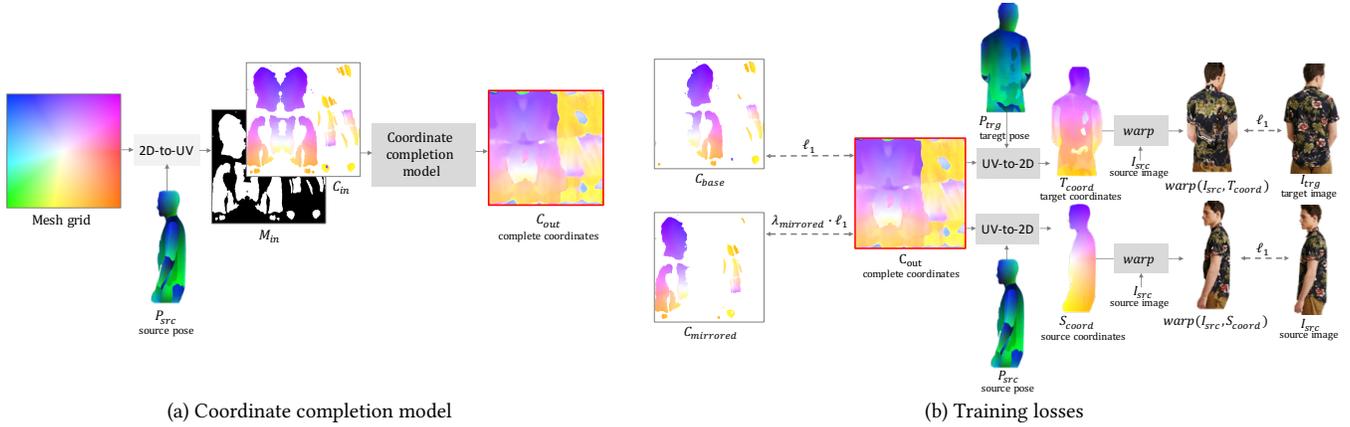


Fig. 5. **Coordinate completion model.** The goal of the coordinate completion model is to learn how to *reuse* the local features of the visible parts of the human in the source image for the invisible parts (unseen in the source pose) in the target pose. (a) Given a mesh grid and the dense pose of the input source image P_{src} , we map the base coordinates C_{base} and their symmetric counterpart $C_{mirrored}$ from the 2D mesh grid to the UV-space using a pre-computed mapping table. We then concatenate the combined coordinates C_{in} and their corresponding visibility mask M_{in} as input to the coordinate completion model. (b) We train the model to minimize the L1 loss between the predicted coordinates C_{out} and the input coordinates C_{in} as shown in Eqn. 4. We also minimize the L1 loss between the warped source image and the warped target image as shown in Eqn. 5.

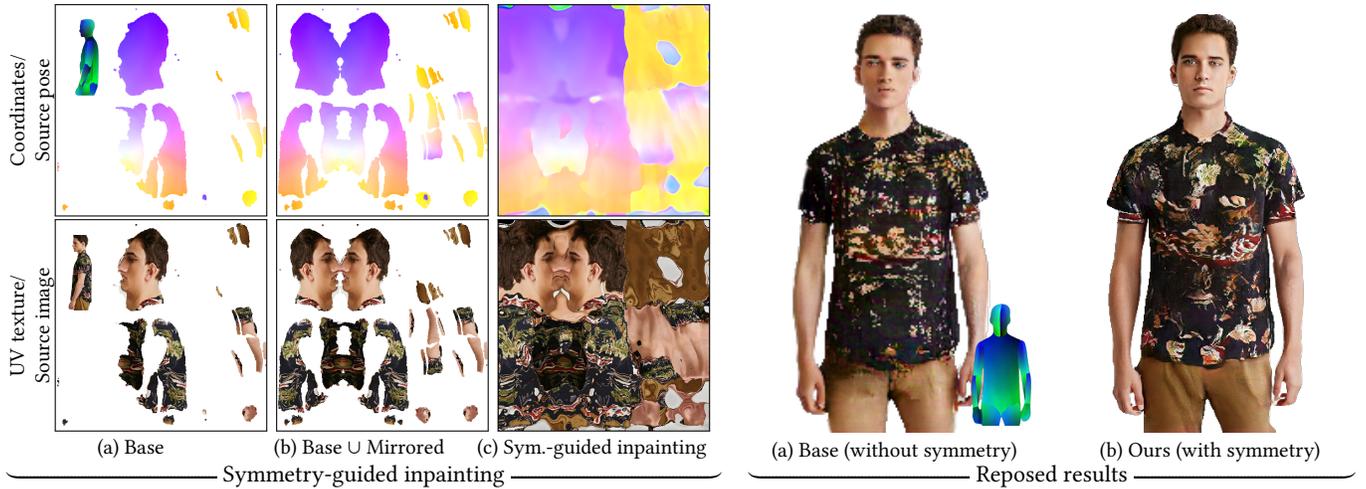


Fig. 6. **Symmetry-guided inpainting.** (Left) Given a mesh grid and the source image dense pose P_{src} , we first map the coordinates from the 2D mesh grid to appropriate locations in the UV-space using a pre-computed mapping table. We can then use these mapped base coordinates C_{base} to warp RGB pixels from the input source image I_{src} . We show the base coordinates and their warped RGB pixels in (a). (b) In addition to these base coordinates, we can also map the left-right *mirrored* coordinates $C_{mirrored}$ from the 2D mesh grid to the UV space. To train our coordinate completion model, we combine the incomplete base and mirrored coordinates in the UV-space. We then concatenate these combined coordinates with their respective mask and pass them as input to our coordinate completion model. We show our completed coordinates and the UV texture map in (c). (Right) We compare the reposing results *without* and *with* the proposed symmetry prior.

- ℓ_1 loss. We minimize the ℓ_1 loss between the foreground human regions of the synthesized image \hat{I}_{trg} and of the ground truth target I_{trg} .

$$L_{\ell_1} = \|\hat{I}_{trg} \cdot M_{trg} - I_{trg} \cdot M_{trg}\|_1, \quad (9)$$

where M_{trg} is the human foreground mask estimated using a human parsing method [Gong et al. 2018].

- Perceptual loss. We minimize the weighted sum of the ℓ_1 loss between the pretrained VGG features of the synthesized \hat{I}_{trg}

foreground and the ground truth I_{trg} foreground such that:

$$L_{vgg} = \sum_{i=1}^5 w_i \cdot \|VGG_{l_i}(\hat{I}_{trg} \cdot M_{trg}) - VGG_{l_i}(I_{trg} \cdot M_{trg})\|_1 \quad (10)$$

We use $w = [\frac{1}{32}, \frac{1}{16}, \frac{1}{8}, \frac{1}{4}, 1.0]$ and VGG ReLU output layers $l = [1, 6, 11, 20, 29]$ following [Park et al. 2019; Wang et al. 2018b].

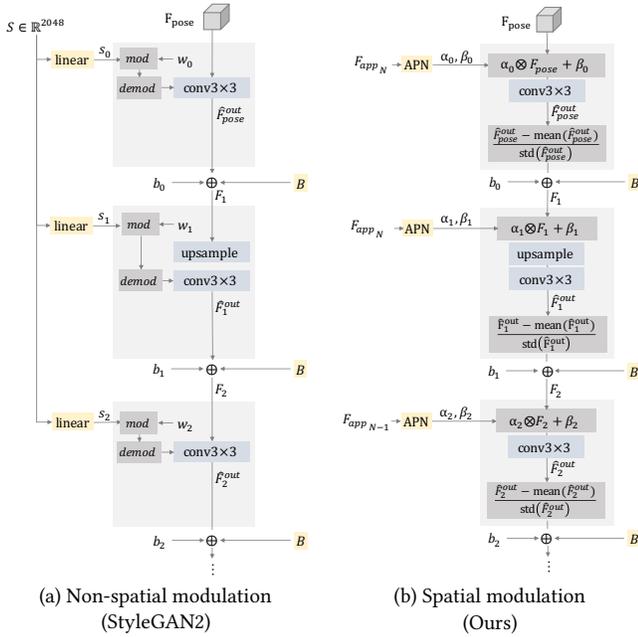


Fig. 7. **Spatial vs. non-spatial modulation of StyleGAN2 features.** Our input to the StyleGAN2 generator is the encoded target pose features F_{pose} (a) StyleGAN2 [Karras et al. 2020] performs *non-spatial modulation* of features by modulating and demodulating the weights of the convolutions using the learned style vector S . After the convolution the bias is added as well as StyleGAN2 noise broadcast operation B . (b) To better leverage the spatial features for preserving appearance details, we propose *spatial modulation* of styleGAN2 features. Instead of modulating and demodulating the weights of the convolutions, we modulate the mean and standard deviation of the features. We perform this modulation before the convolution using the shifting and scaling parameters, α and β , generated by the affine parameters network (APN). We then normalize the output of the convolution to zero mean and unit standard deviation before adding the bias and StyleGAN2 noise broadcast operation B .

- **Face identity loss.** We use MTCNN [Zhang et al. 2016] to detect, crop, and align faces from the generated image \hat{I}_{trg} and ground truth target I_{trg} . When a face is detected, we maximize the cosine similarity between the pretrained SphereFace [Liu et al. 2017] features of the generated face and the ground truth target face. Such that:

$$L_{face} = 1 - \left(\frac{SF(\hat{I}_{trg})^T SF(I_{trg})}{\max(\|SF(\hat{I}_{trg})\|_2 \cdot \|SF(I_{trg})\|_2, \epsilon)} \right) \quad (11)$$

where SF is the pretrained SphereFace feature extractor, and $SF(\hat{I}_{trg})$ and $SF(I_{trg})$ are features of the aligned faces of the generated and ground truth image respectively. $\epsilon = e^{-8}$ is a very small value to avoid zero-dividing.

Therefore, our final loss is: $L = L_{adv} + L_{\ell_1} + L_{vgg} + L_{face}$.

4 EXPERIMENTAL RESULTS

4.1 Experimental setup

4.1.1 Implementation details. We implement our model with PyTorch. We use ADAM optimizer with a learning rate of $ratio \cdot 0.002$

Table 1. Quantitative comparison with the state-of-the-art methods on the DeepFashion dataset [Liu et al. 2016].

	PSNR \uparrow	SSIM \uparrow	FID \downarrow	LPIPS \downarrow
<i>Resolution = 174 × 256</i>				
PATN [Zhu et al. 2019]	17.7041	0.7543	21.8568	0.195
ADGAN [Men et al. 2020]	17.7223	0.7544	16.2686	0.175
GFLA [Ren et al. 2020]	18.0424	0.7625	15.1722	0.167
Ours	18.5062	0.7784	8.7449	0.134
<i>Resolution = 348 × 512</i>				
GFLA [Ren et al. 2020]	17.9718	0.7540	18.8519	0.170
Ours	18.3567	0.7640	9.0002	0.143

Table 2. Quantitative comparison on 348 × 512 resolution with StylePoseGAN [Sarkar et al. 2021] on their DeepFashion dataset train/test split.

	PSNR \uparrow	SSIM \uparrow	FID \downarrow	LPIPS \downarrow
StylePoseGAN [Sarkar et al. 2021]	17.7568	0.7508	7.4804	0.167
Ours	18.5029	0.7711	6.0557	0.144

and beta parameters $(0, 0.99^{ratio})$. We set the generator ratio to $\frac{4}{5}$ and discriminator ratio to $\frac{16}{17}$.

4.1.2 Training. We first train our model by focusing on generating the foreground. We apply the reconstruction loss and the adversarial loss only on the foreground. We set the batch size to 1 and train for 50 epochs. This training process takes around 7 days on 8 NVIDIA 2080 Ti GPUs. We then finetune the model by applying the adversarial loss globally on the entire image. We set the batch size to 8 and train for 10 epochs. This training process takes less than 2 days on 2 A100 GPUs. At test time, generating a reposing results with 384 × 512 resolution takes 0.4 seconds using 1 NVIDIA 2080 Ti GPU.

4.1.3 Dataset. We use the DeepFashion dataset [Liu et al. 2016] for training and evaluation. We follow the train/test splits (101,967 training and 8,570 testing pairs) of recent methods [Men et al. 2020; Ren et al. 2020; Zhu et al. 2019].

4.2 Evaluations

4.2.1 Quantitative evaluation. is reported in Table 1. We report the human foreground peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), learned perceptual image patch similarity (LPIPS) [Zhang et al. 2018], and Frchet Inception Distance (FID) [Heusel et al. 2017]. PSNR/SSIM often do not correlate well with perceived quality, particularly for synthesis tasks. For example, PSNR may favor blurry results over sharp ones. We report these metrics only for completeness.

Our method compares favorably against existing works such as PATN [Zhu et al. 2019] ADGAN [Men et al. 2020], and GFLA [Ren et al. 2020]. Our method also compares favorably against the concurrent work StylePoseGAN [Sarkar et al. 2021]. We report the quantitative evaluation in Table 2. We train and test our method using their DeepFashion dataset train/test split.

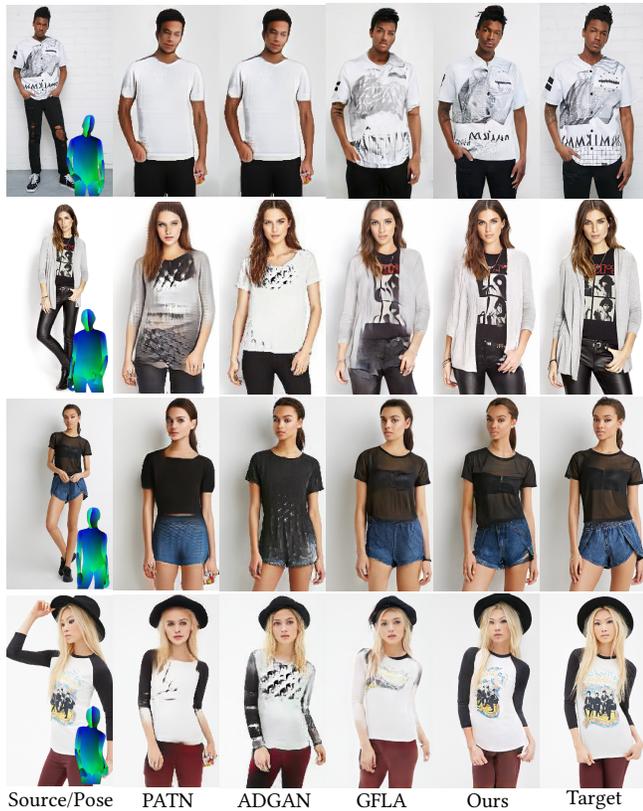


Fig. 8. **Visual comparison for human reposing.** We show visual comparison of our method with PATN [Zhu et al. 2019], ADGAN [Men et al. 2020], and GFLA [Ren et al. 2020] on DeepFashion dataset [Liu et al. 2016]. Our approach successfully captures the local details from the source image.

Table 3. The ability to preserve the identity of the reposed person.

	Arcface	Spherenet
GFLA [Ren et al. 2020]	0.117	0.197
Ours	<u>0.373</u>	<u>0.438</u>
Ground truth	0.555	0.444

4.2.2 *Visual comparison.* in Figure 8 and Figure 9 show that our proposed approach captures finer-grained appearance details from the input source images.

4.2.3 *Face identity.* We evaluate our model’s ability to preserve the reposed person’s identity. For test images with visible faces (7,164 from 8,570), we report in Table 3 the averaged cosine similarity between the face features (Arcface [Deng et al. 2019] and Spherenet [Coors et al. 2018]) extracted from the aligned faces in the source/target images. Our method compares favorably against GFLA [Ren et al. 2020].

4.3 Ablation study

For the ablation study, we report the results of the foreground-focused trained model.

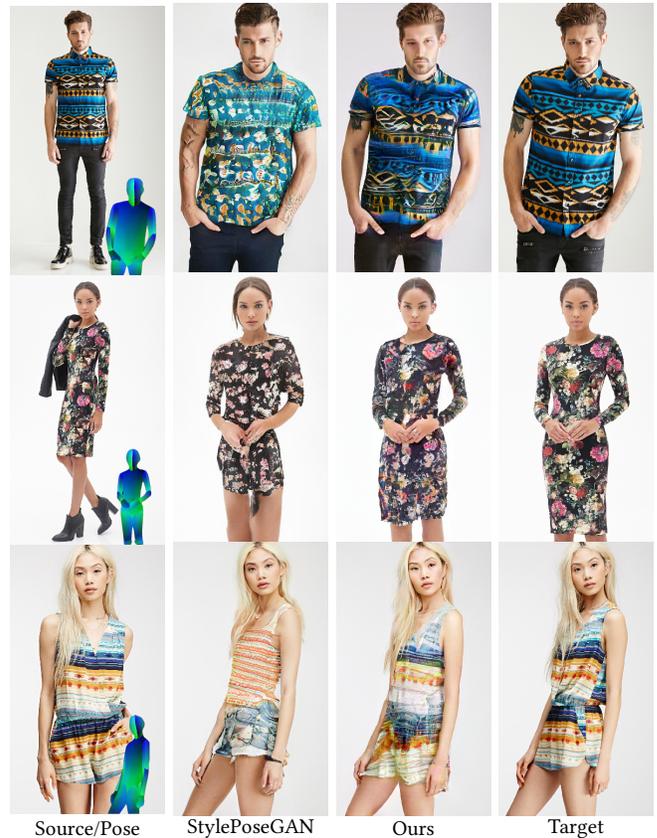


Fig. 9. **Visual comparison for human reposing.** We compare our method with StylePoseGAN [Sarkar et al. 2021] on their train/test split of DeepFashion dataset [Liu et al. 2016]. Our approach preserves the appearance and captures the fine-grained details of the source image.

Table 4. The effect of symmetry-guided coordinate inpainting on the DeepFashion dataset [Liu et al. 2016].

	PSNR \uparrow	SSIM \uparrow	FID \downarrow	LPIPS \downarrow
Without symmetry	18.8810	0.7886	8.5240	0.129
With symmetry (Ours)	18.9657	0.7919	8.1434	0.124

4.3.1 *Symmetry prior.* We evaluate the effectiveness of adding the symmetry prior to the input of the coordinate completion model. We train our networks *with* and *without* this symmetry prior and report the quantitative results in Table 4. Results show that adding the symmetry prior indeed improves the quality of the synthesis. We note that the symmetry prior generally works well for repetitive/textured patterns, but may introduce artifacts for unique patterns (e.g., text).

4.3.2 *Modulation schemes.* We quantitatively and qualitatively demonstrate the effectiveness of our proposed spatial modulation. We show quantitative evaluation in Table 5. We show qualitative results in Figure 10. Results show that spatially varying modulation improves the quality of the synthesized human foreground and captures the spatial details of the source images regardless of the input source image type.



Fig. 10. **Ablation.** We compare our results with other variants, including the modulation types and source of appearance features. We show that the proposed spatial modulation captures finer-grained details from the source image. Transferring appearance features from the source image leads to fewer artifacts compared to features from the UV space.

Table 5. Ablation on source for appearance (Incomplete UV, Complete UV, and Image) and modulation types (Spatial and Non-spatial).

ID	Input Source	Spatial	PSNR \uparrow	SSIM \uparrow	FID \downarrow	LPIPS \downarrow
A	Incomplete UV	-	18.7385	0.7710	9.4150	0.151
B	Incomplete UV	✓	18.6005	0.7696	9.2321	0.146
C	Complete UV	-	18.9407	0.7770	9.7435	0.147
D	Complete UV	✓	18.7063	0.7720	<u>9.0236</u>	<u>0.143</u>
E	Source image	-	18.6027	0.7678	9.4367	0.154
F	Source image	✓	<u>18.7420</u>	<u>0.7739</u>	8.8060	0.139

4.3.3 *Sources of appearance.* We experiment with multiple variants for encoding source appearance. Specifically, the *incomplete UV*, *complete UV* (completed using the coordinate completion model), and *source image* (our approach shown in Figure 3). We report the quantitative results in Table 5 and visual results in Figure 10. The results show that extracting appearance features directly from the source image preserves more details than other variants.

4.4 Garment transfer Results

Using the UV-space pre-computed mapping table, we can segment the UV-space into human body parts (Figure 11). We can then use this UV-space segmentation map to generate the target pose segmentation map using the target dense pose P_{trg} . The target segmentation

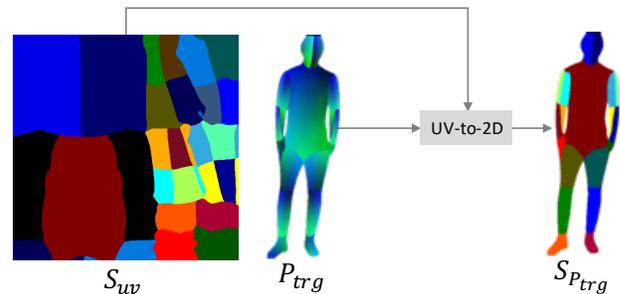


Fig. 11. **Human body segmentation.** We create a UV-space segmentation map S_{uv} of the human body using the UV-space pre-computed mapping table. We use the target dense pose P_t to map this UV-space segmentation map to 2D target pose S_{P_t} , which can then be used to combine features from multiple source images to perform garment transfer.

map allows us to combine partial features from multiple source images to perform garment transfer. Figure 12 shows examples of bottom and top garment transfer.

4.5 Limitations

4.5.1 *Failure cases.* Human reposing from a single image remains challenging. Figure 13 shows two failure cases where our approach fails to synthesize realistic hands and clothing textures. Hands are

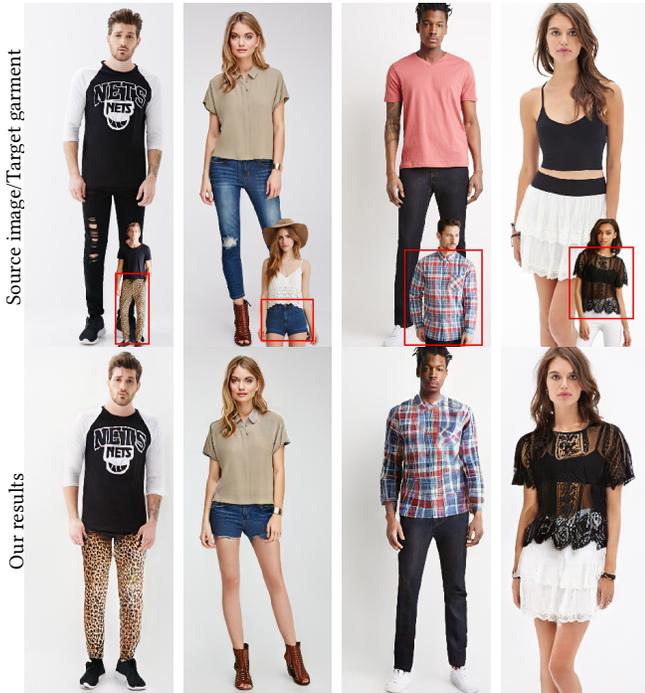


Fig. 12. **Garment transfer.** We show examples of garment transfer for bottom (left) and top (right) garment sources.



Fig. 13. **Failure cases.** Our method produces artifacts on the hands (left) and the skirt (right).

difficult to capture due to the coarse granularity of DensePose. Explicitly parsing hands could help establish more accurate correspondence between source/target pose (e.g., using Monocular total capture [Xiang et al. 2019]). Long-hairs and loose-fit clothes (e.g., skirts) are challenging because they are not captured by DensePose. We believe that incorporating human/garment parsing in our framework may help mitigate the artifacts.

4.5.2 Diversity. DeepFashion dataset consists of mostly young fit models and very few dark-skinned individuals. Our trained model thus inherit the biases and perform worse on unrepresented individuals as shown in Figure 14. We believe that training and evaluating on diverse populations and appearance variations are important future directions.



Fig. 14. **Diverse in the wild cases.** Our model inherits the biases of DeepFashion dataset and thus performs worse on unrepresented individuals. Our method cannot accurately synthesize curly hair (left) and fails in reposing dark-skinned individuals (right). Images from Unplash.

5 CONCLUSIONS

We have presented a simple yet effective approach for pose-guided image synthesis. Our core technical novelties lie in 1) spatial modulation of a pose-conditioned StyleGAN generator and 2) a symmetry-guided inpainting network for completing correspondence field. We demonstrate that our approach is capable of synthesizing *photo-realistic* images in the desired target pose and *preserving details* from the source image. We validate various design choices through an ablation study and show improved results when compared with the state-of-the-art human reposing algorithms. Our controllable human image synthesis approach enables high-quality human pose transfer and garment transfer, providing a promising direction for rendering human images.

REFERENCES

- Rameen Abdal, Peihao Zhu, Niloy J Mitra, and Peter Wonka. 2021. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *TOG* 40, 3 (2021), 1–21.
- Kfir Aberman, Mingyi Shi, Jing Liao, Dani Lischinski, Baoquan Chen, and Daniel Cohen-Or. 2019. Deep video-based performance cloning. *CGF* 38, 2 (2019), 219–233.
- Badour AlBahar and Jia-Bin Huang. 2019. Guided image-to-image translation with bi-directional feature transformation. In *ICCV*.
- Yazeed Alharbi and Peter Wonka. 2020. Disentangled image generation through structured noise injection. In *CVPR*.
- Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. 2019. Tex2shape: Detailed full human body geometry from a single image. In *ICCV*.
- Guha Balakrishnan, Amy Zhao, Adrian V Dalca, Fredo Durand, and John Guttag. 2018. Synthesizing images of humans in unseen poses. In *CVPR*.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. 2019. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*.
- Lucy Chai, Jonas Wulff, and Phillip Isola. 2021. Using latent space regression to analyze and leverage compositionality in GANs. In *ICLR*.
- Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. 2019. Everybody dance now. In *ICCV*.
- Edo Collins, Raja Bala, Bob Price, and Sabine Susstrunk. 2020. Editing in style: Uncovering the local semantics of gans. In *CVPR*.
- Benjamin Coors, Alexandru Paul Condrache, and Andreas Geiger. 2018. Spherenet: Learning spherical representations for detection and classification in omnidirectional images. In *ECCV*.
- Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*.
- Patrick Esser, Ekaterina Sutter, and Björn Ommer. 2018. A variational u-net for conditional appearance and shape generation. In *CVPR*.
- Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2021. Dynamic Neural Radiance Fields for Monocular 4D Facial Avatar Reconstruction. In *CVPR*.
- Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. 2021. Dynamic View Synthesis from Dynamic Monocular Video. In *ICCV*.
- Chen Gao, Yichang Shih, Wei-Sheng Lai, Chia-Kai Liang, and Jia-Bin Huang. 2020. Portrait Neural Radiance Fields from a Single Image. *arXiv preprint arXiv:2012.05903* (2020).
- Ke Gong, Xiaodan Liang, Yicheng Li, Yimin Chen, Ming Yang, and Liang Lin. 2018. Instance-level human parsing via part grouping network. In *ECCV*.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In

- NeurIPS*.
- Artur Grigorev, Artem Sevastopolsky, Alexander Vakhitov, and Victor Lempitsky. 2019. Coordinate-based texture inpainting for pose-guided human image generation. In *CVPR*.
- Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. 2018. Densepose: Dense human pose estimation in the wild. In *CVPR*.
- Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. 2020. Ganspace: Discovering interpretable gan controls. In *NeurIPS*.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*.
- Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. 2018. Multimodal unsupervised image-to-image translation. In *ECCV*.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *CVPR*.
- Ali Jahani, Lucy Chai, and Phillip Isola. 2020. On the "steerability" of generative adversarial networks. In *ICLR*.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*.
- Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *CVPR*.
- Hyunso Kim, Yunje Choi, Junho Kim, Sungjoo Yoo, and Youngjung Uh. 2021. StyleMapGAN: Exploiting Spatial Dimensions of Latent in GAN for Real-time Image Editing. In *CVPR*.
- Hyeongwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. 2018. Deep video portraits. *TOG* 37, 4 (2018), 1–14.
- Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 2019. 360-degree textures of people in clothing from a single image. In *International Conference on 3D Vision*.
- Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. 2018. Diverse image-to-image translation via disentangled representations. In *ECCV*.
- Kathleen M Lewis, Srivatsan Varadharajan, and Ira Kemelmacher-Shlizerman. 2021. VOGUE: Try-On by StyleGAN Interpolation Optimization. *arXiv preprint arXiv:2101.02285* (2021).
- Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. 2021. Neural Scene Flow Fields for Space-Time View Synthesis of Dynamic Scenes. In *CVPR*.
- Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature pyramid networks for object detection. In *CVPR*.
- Lingjie Liu, Weipeng Xu, Marc Habermann, Michael Zollhofer, Florian Bernard, Hyeongwoo Kim, Wenping Wang, and Christian Theobalt. 2020. Neural human video rendering by learning dynamic textures and rendering-to-video translation. *TVCG* (2020).
- Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. 2017. Sphereface: Deep hypersphere embedding for face recognition. In *CVPR*.
- Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations. In *CVPR*.
- Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehmann, and Yaser Sheikh. 2019. Neural Volumes: Learning Dynamic Renderable Volumes from Images. *TOG* 38, 4, Article 65 (July 2019), 14 pages.
- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. 2015. SMPL: A skinned multi-person linear model. *TOG* 34, 6 (2015), 1–16.
- Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. 2017. Pose guided person image generation. In *NeurIPS*.
- Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. 2018. Disentangled person image generation. In *CVPR*.
- Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. 2020. Controllable person image synthesis with attribute-decomposed gan. In *CVPR*.
- Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. 2019. Neural re-rendering in the wild. In *CVPR*.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2020. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*.
- Natalia Neverova, Riza Alp Güler, and Iasonas Kokkinos. 2018. Dense pose transfer. In *ECCV*.
- Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. 2021. Neural Articulated Radiance Field. In *ICCV*.
- Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo-Martin Brualla. 2021. Nerfies: Deformable Neural Radiance Fields. In *ICCV*.
- Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. 2019. Semantic Image Synthesis with Spatially-Adaptive Normalization. In *CVPR*.
- William Peebles, John Peebles, Jun-Yan Zhu, Alexei Efros, and Antonio Torralba. 2020. The hessian penalty: A weak prior for unsupervised disentanglement. In *ECCV*.
- Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. 2021. Animatable Neural Radiance Fields for Human Body Modeling. In *ICCV*.
- Albert Pumarola, Antonio Agudo, Alberto Sanfeliu, and Francesc Moreno-Noguer. 2018. Unsupervised person image synthesis in arbitrary poses. In *CVPR*.
- Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. 2021. ANR: Articulated Neural Rendering for Virtual Avatars. In *CVPR*.
- Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H Li, and Ge Li. 2020. Deep image spatial transformation for person image generation. In *CVPR*.
- Kripasindhu Sarkar, Vladislav Golyanik, Lingjie Liu, and Christian Theobalt. 2021. Style and Pose Control for Image Synthesis of Humans from a Single Monocular View. *arXiv preprint arXiv:2102.11263* (2021).
- Kripasindhu Sarkar, Dushyant Mehta, Weipeng Xu, Vladislav Golyanik, and Christian Theobalt. 2020. Neural re-rendering of humans from a single image. In *ECCV*.
- Yujun Shen and Bolei Zhou. 2021. Closed-form factorization of latent semantics in gans. In *CVPR*.
- Assaf Shocher, Yossi Gandelsman, Inbar Mosseri, Michal Yarom, Michal Irani, William T Freeman, and Tali Dekel. 2020. Semantic pyramid for image generation. In *CVPR*.
- Alon Shoshan, Nadav Bihonker, Igor Kviatkovsky, and Gerard Medioni. 2021. GAN-Control: Explicitly Controllable GANs. In *ICCV*.
- Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, and Nicu Sebe. 2018. Deformable gans for pose-based human image generation. In *CVPR*.
- Sudipta N Sinha, Krishnan Ramnath, and Richard Szeliski. 2012. Detecting and reconstructing 3d mirror symmetric objects. In *ECCV*.
- Ayush Tewari, Mohamed Elgharib, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. 2020a. Pie: Portrait image embedding for semantic control. *TOG* 39, 6 (2020), 1–14.
- Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. 2020b. Stylerig: Rigging stylegan for 3d control over portrait images. In *CVPR*.
- Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. 2020c. State of the art on neural rendering. In *CGF*, Vol. 39. 701–727.
- Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2019. Deferred neural rendering: Image synthesis using neural textures. *TOG* 38, 4 (2019), 1–12.
- Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. 2021. Non-Rigid Neural Radiance Fields: Reconstruction and Novel View Synthesis of a Deforming Scene from Monocular Video. In *ICCV*.
- Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. 2019. Few-shot video-to-video synthesis. In *NeurIPS*.
- Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018a. Video-to-video synthesis. In *NeurIPS*.
- Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018b. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *CVPR*.
- Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021. One-Shot Free-View Neural Talking-Head Synthesis for Video Conferencing. In *CVPR*.
- Shangzhe Wu, Christian Ruppert, and Andrea Vedaldi. 2020. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *CVPR*.
- Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. 2021. Space-time Neural Irradiance Fields for Free-Viewpoint Video. In *CVPR*.
- Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. 2019. Monocular Total Capture: Posing Face, Body, and Hands in the Wild. In *CVPR*.
- Raymond A Yeh, Yuan-Ting Hu, and Alexander G Schwing. 2019. Chirality nets for human pose regression. In *NeurIPS*.
- Jaeh Shin Yoon, Lingjie Liu, Vladislav Golyanik, Kripasindhu Sarkar, Hyun Soo Park, and Christian Theobalt. 2021. Pose-Guided Human Animation from a Single Image in the Wild. In *CVPR*.
- Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. 2019. Free-Form Image Inpainting with Gated Convolution. In *ICCV*.
- Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. 2019. Few-shot adversarial learning of realistic neural talking head models. In *ICCV*.
- Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. 2019. Self-attention generative adversarial networks. In *ICML*.
- Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. 2016. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters* 23, 10 (2016), 1499–1503.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.
- Xiuming Zhang, Sean Fanello, Yun-Ta Tsai, Tiancheng Sun, Tianfan Xue, Rohit Pandey, Sergio Orts-Escolano, Philip Davidson, Christoph Rhemann, Paul Debevec, et al. 2021. Neural light transport for relighting and view synthesis. *TOG* 40, 1 (2021), 1–17.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*.
- Zhen Zhu, Tengpeng Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. 2019. Progressive Pose Attention Transfer for Person Image Generation. In *CVPR*.