

Dense Pose Transfer

Natalia Neverova¹, Rıza Alp Güler², and Iasonas Kokkinos¹

¹ Facebook AI Research, Paris, France, {nneverova, iasonask}@fb.com

² INRIA-CentraleSupélec, Paris, France, riza.guler@inria.fr

Abstract. In this work we integrate ideas from surface-based modeling with neural synthesis: we propose a combination of surface-based pose estimation and deep generative models that allows us to perform accurate pose transfer, i.e. synthesize a new image of a person based on a single image of that person and the image of a pose donor. We use a dense pose estimation system that maps pixels from both images to a common surface-based coordinate system, allowing the two images to be brought in correspondence with each other. We inpaint and refine the source image intensities in the surface coordinate system, prior to warping them onto the target pose. These predictions are fused with those of a convolutional predictive module through a neural synthesis module allowing for training the whole pipeline jointly end-to-end, optimizing a combination of adversarial and perceptual losses. We show that dense pose estimation is a substantially more powerful conditioning input than landmark-, or mask-based alternatives, and report systematic improvements over state of the art generators on DeepFashion and MVC datasets.




Fig. 1. Overview of our pose transfer pipeline: given an input image and a target pose we use DensePose [1] to drive the generation process. This is achieved through the complementary streams of (a) a data-driven predictive module, and (b) a surface-based module that warps the texture to UV-coordinates, interpolates on the surface, and warps back to the target image. A blending module combines the complementary merits of these two streams in a single end-to-end trainable framework.

1 Introduction

Deep models have recently shown remarkable success in tasks such as face [2], human [3–5], or scene generation [6, 7], collectively known as “neural synthesis”. This opens countless possibilities in computer graphics applications, including cinematography, gaming and virtual reality settings. At the same time, the potential malevolent use of this technology raises new research problems, including the detection of forged images or videos [8], which in turn requires training forgery detection algorithms with multiple realistic samples. In addition, synthetically generated images have been successfully exploited for data augmentation and training deep learning frameworks for relevant recognition tasks [9].

In most applications, the relevance of a generative model to the task directly relates to the amount of control that one can exert on the generation process. Recent works have shown the possibility of adjusting image synthesis by controlling categorical attributes [10, 3], low-dimensional parameters [11], or layout constraints indicated by a conditioning input [12, 6, 7, 3–5]. In this work we aspire to obtain a stronger hold of the image synthesis process by relying on surface-based object representations, similar to the ones used in graphics engines.

Our work is focused on the human body, where surface-based image understanding has been most recently unlocked [13–18, 1]. We build on the recently introduced SMPL model [13] and the DensePose system [1], which taken together allow us to interpret an image of a person in terms of a full-fledged surface model, getting us closer to the goal of performing *inverse graphics*.

In this work we close the loop and perform image generation by rendering the same person in a new pose through surface-based neural synthesis. The target pose is indicated by the image of a ‘pose donor’, i.e. another person that guides the image synthesis. The DensePose system is used to associate the new photo with the common surface coordinates and copy the appearance predicted there.

The purely geometry-based synthesis process is on its own insufficient for realistic image generation: its performance can be compromised by inaccuracies of the DensePose system as well as by self-occlusions of the body surface in at least one of the two images. We account for occlusions by introducing an inpainting network that operates in the surface coordinate system and combine its predictions with the outputs of a more traditional feedforward conditional synthesis module. These predictions are obtained independently and compounded by a refinement module that is trained so as to optimize a combination of reconstruction, perceptual and adversarial losses.

We experiment on the DeepFashion [19] and MVC [20] datasets and show that we can obtain results that are quantitatively better than the latest state-of-the-art. Apart from the specific problem of pose transfer, the proposed combination of neural synthesis with surface-based representations can also be promising for the broader problems of virtual and augmented reality: the generation process is more transparent and easy to connect with the physical world, thanks to the underlying surface-based representation. In the more immediate future, the task of pose transfer can be useful for dataset augmentation, training forgery

detectors, as well as texture transfer applications like those showcased by [1], without however requiring the acquisition of a surface-level texture map.

2 Previous works

Deep generative models have originally been studied as a means of unsupervised feature learning [21–24]; however, based on the increasing realism of neural synthesis models [25, 2, 6, 7] such models can now be considered as components in computer graphics applications such as content manipulation [7].

Loss functions used to train such networks largely determine the realism of the resulting outputs. Standard reconstruction losses, such as ℓ_1 or ℓ_2 norms typically result in blurry results, but at the same time enhance stability [12]. Realism can be enforced by using an adapted discriminator loss trained in tandem with a generator in Generative Adversarial Network (GAN) architectures [23] to ensure that the generated and observed samples are indistinguishable. However, this training can often be unstable, calling for more robust variants such as the squared loss of [26], WGAN and its variants [27] or multi-scale discriminators as in [7]. An alternative solution is the perceptual loss used in [28, 29] replacing the optimization-based style transfer of [25] with feedforward processing. This was recently shown in [6] to deliver substantially more accurate scene synthesis results than [12], while compelling results were obtained more recently by combining this loss with a GAN-style discriminator [7].

Person and clothing synthesis has been addressed in a growing body of recently works [3, 5, 4, 30]. All of these works assist the image generation task through domain, person-specific knowledge, which gives both better quality results, and a more controllable image generation pipeline.

Conditional neural synthesis of humans has been shown in [5, 4] to provide a strong handle on the output of the generative process. A controllable surface-based model of the human body is used in [3] to drive the generation of persons wearing clothes with controllable color combinations. The generated images are demonstrably realistic, but the pose is determined by controlling a surface based model, which can be limiting if one wants e.g. to render a source human based on a target video. A different approach is taken in the pose transfer work of [4], where a sparse set of landmarks detected in the target image are used as a conditioning input to a generative model. The authors show that pose can be generated with increased accuracy, but often losing texture properties of the source images, such as cloth color or texture properties. In the work of [31] multi-view supervision is used to train a two-stage system that can generate images from multiple views. In more recent work [5] the authors show that introducing a correspondence component in a GAN architecture allows for substantially more accurate pose transfer.

Image inpainting helps estimate the body appearance on occluded body regions. Generative models are able to fill-in information that is labelled as occluded, either by accounting for the occlusion pattern during training [32], or

by optimizing a score function that indicates the quality of an image, such as the negative of a GAN discriminator loss [33]. The work of [33] inpaints arbitrarily occluded faces by minimizing the discriminator loss of a GAN trained with fully-observed face patches. In the realm of face analysis impressive results have been generated recently by works that operate in the UV coordinate system of the face surface, aiming at photorealistic face inpainting [34], and pose-invariant identification [35]. Even though we address a similar problem, the lack of access to full UV recordings (as in [35, 34]) poses an additional challenge.

3 Dense Pose Transfer

We develop our approach to pose transfer around the *DensePose* estimation system [1] to associate every human pixel with its coordinates on a surface-based parameterization of the human body in an efficient, bottom-up manner. We exploit the DensePose outputs in two complementary ways, corresponding to the *predictive module* and the *warping module*, as shown in Fig. 1. The warping module uses DensePose surface correspondence and inpainting to generate a new view of the person, while the predictive module is a generic, *black-box*, generative model conditioned on the DensePose outputs for the input and the target.

These modules corresponding to two parallel streams have complementary merits: the *predictive module* successfully exploits the dense conditioning output to generate plausible images for familiar poses, delivering superior results to those obtained from sparse, landmark-based conditioning; at the same time, it cannot generalize to new poses, or transfer texture details. By contrast, the *warping module* can preserve high-quality details and textures, allows us to perform inpainting in a uniform, canonical coordinate system, and generalizes for free for a broad variety of body movements. However, its body-, rather than clothing-centered construction does not take into account hair, hanging clothes, and accessories. The best of both worlds is obtained by feeding the outputs of these two blocks into a *blending module* trained to fuse and refine their predictions using a combination of reconstruction, adversarial, and perceptual losses in an end-to-end trainable framework.

The DensePose module is common to both streams and delivers dense correspondences between an image and a surface-based model of the human body. It does so by firstly assigning every pixel to one of 24 predetermined surface parts, and then regressing the part-specific surface coordinates of every pixel. The results of this system are encoded in three output channels, comprising the part label and part-specific UV surface coordinates. This system is trained discriminatively and provides a simple, feed-forward module for dense correspondence from an image to the human body surface. We omit further details, since we rely on the system of [1] with minor implementation differences described in Sec. 4.

Having outlined the overall architecture of our system, in Sec. 3.1 and Sec. 3.3 we present in some more detail our components, and then turn in Sec. 3.4 to the loss functions used in their training. A thorough description of architecture

details is left to the supplemental material. We start by presenting the architecture of the predictive stream, and then turn to the surface-based stream, corresponding to the upper and lower rows of Fig. 1, respectively.

3.1 Predictive stream

The predictive module is a conditional generative model that exploits the DensePose system results for pose transfer. Existing conditional models indicate the target pose in the form of heat-maps from keypoint detectors [4], or part segmentations [3]. Here we condition on the concatenation of the input image and DensePose results for the input and target images, resulting in an input of dimension $256 \times 256 \times 9$. This provides conditioning that is both global (part-classification), and point-level (continuous coordinates), allowing the remaining network to exploit a richer source of information.

The remaining architecture includes an encoder followed by a stack of residual blocks and a decoder at the end, along the lines of [28]. In more detail, this network comprises (a) a cascade of three convolutional layers that encode the $256 \times 256 \times 9$ input into $64 \times 64 \times 256$ activations, (b) a set of six residual blocks with $3 \times 3 \times 256 \times 256$ kernels, (c) a cascade of two deconvolutional and one convolutional layer that deliver an output of the same spatial resolution as the input. All intermediate convolutional layers have 3×3 filters and are followed by instance normalization [36] and ReLU activation. The last layer has tanh non-linearity and no normalization.

3.2 Warping stream

Our warping module performs pose transfer by performing explicit texture mapping between the input and the target image on the common surface UV-system. The core of this component is a Spatial Transformer Network (STN) [37] that warps according to DensePose the image observations to the UV-coordinate system of each surface part; we use a grid with 256×256 UV points for each of the 24 surface parts, and perform scattered interpolation to handle the continuous values of the regressed UV coordinates. The inverse mapping from UV to the output image space is performed by a second STN with a bilinear kernel. As shown in Fig. 3, a direct implementation of this module would often deliver poor results: the part of the surface that is visible on the source image is typically small, and can often be entirely non-overlapping with the part of the body that is visible on the target image. This is only exacerbated by DensePose failures or systematic errors around the part seams. These problems motivate the use of an inpainting network within the warping module, as detailed below.

Inpainting autoencoder. This model allows us to extrapolate the body appearance from the surface nodes populated by the STN to the remainder of the surface. Our setup requires a different approach to the one of other deep inpainting methods [33], because we never observe the full surface texture during training. We handle the partially-observed nature of our training signal by using

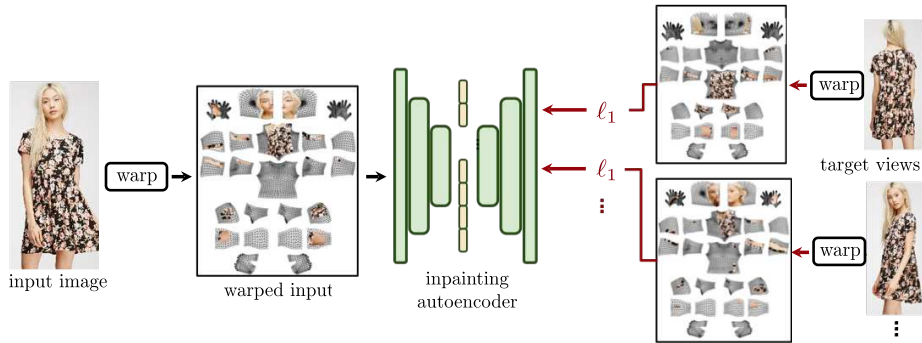


Fig. 2. Supervision signals for pose transfer on the warping stream: the input image on the left is warped to intrinsic surface coordinates through a spatial transformer network driven by DensePose. From this input, the inpainting autoencoder has to predict the appearance of the same person from different viewpoints, when also warped to intrinsic coordinates. The loss functions on the right penalize the reconstruction only on the observed parts of the texture map. This form of multi-view supervision acts like a surrogate for the (unavailable) appearance of the person on the full body surface.

a reconstruction loss that only penalizes the observed part of the UV map, and lets the network freely *guess* the remaining domain of the signal. In particular, we use a masked ℓ_1 loss on the difference between the autoencoder predictions and the target signals, where the masks indicate the visibility of the target signal.

We observed that by its own this does not urge the network to inpaint successfully; results substantially improve when we accompany every input with multiple supervision signals, as shown in Fig. 2, corresponding to UV-wrapped shots of the same person at different poses. This fills up a larger portion of the UV-space and forces the inpainting network to predict over the whole texture domain. As shown in Fig. 3, the inpainting process allows us to obtain a uniformly observed surface, which captures the appearance of skin and tight clothes, but does not account for hair, skirts, or apparel, since these are not accommodated by DensePose’s surface model.

Our inpainting network is comprised of N autoencoders, corresponding to the decomposition of the body surface into N parts used in the original DensePose system [1]. This is based on the observation that appearance properties are non-stationary over the body surface. Propagation of the context-based information from visible to invisible parts that are entirely occluded at the present view is achieved through a fusion mechanism that operates at the level of latent representations delivered by the individual encoders, and injects global pose context in the individual encoding through a concatenation operation.

In particular, we denote by \mathcal{E}_i the individual encoding delivered by the encoder for the i -th part. The fusion layer concatenates these obtained encodings into a single vector which is then down-projected to a 256-dimensional global pose embedding through a linear layer. We pass the resulting embedding through



Fig. 3. Warping module results. For each sample, the top row shows interpolated textures obtained from DensePose predictions and projected on the surface of the 3D body model. The bottom row shows the same textures after inpainting in the UV space.

a cascade of ReLU and Instance-Norm units and transform it again into an embedding denoted by \mathcal{G} .

Then the i -th part decoder receives as an input the concatenation of \mathcal{G} with \mathcal{E}_i , which combines information particular to part i , and global context, delivered by \mathcal{G} . This is processed by a stack of deconvolution operations, which delivers in the end the prediction for the texture of part i .

3.3 Blending module

The blending module’s objective is to combine the complementary strengths of the two streams to deliver a single fused result, that will be ‘polished’ as measured by the losses used for training. As such it no longer involves an encoder or decoder unit, but rather only contains two convolutional and three residual blocks that aim at combining the predictions and refining their results.

In our framework, both predictive and warping modules are first pretrained separately and then finetuned jointly during blending. The final refined output is obtained by learning a residual term added to the output of the predictive stream. The blending module takes an input consisting of the outputs of the predictive and the warping modules combined with the target dense pose.

3.4 Loss Functions

As shown in Fig. 1, the training set for our network comes in the form of pairs of input and target images, \mathbf{x} , \mathbf{y} respectively, both of which are of the same person-clothing, but in different poses. Denoting by $\hat{\mathbf{y}} = G(\mathbf{x})$ the network’s prediction, the difference between $\hat{\mathbf{y}}$, \mathbf{y} can be measured through a multitude of loss terms, that penalize different forms of deviation. We present them below for completeness and ablate their impact in practice in Sec. 4.

Reconstruction loss. To penalize reconstruction errors we use the common ℓ_1 distance between the two signals: $\|\hat{\mathbf{y}} - \mathbf{y}\|_1$. On its own, this delivers blurry results but is important for retaining the overall intensity levels.

Perceptual loss. As in Chen and Koltun [6], we use a VGG19 network pretrained for classification [38] as a feature extractor for both $\hat{\mathbf{y}}$, \mathbf{y} and penalize the ℓ_2 distance of the respective intermediate feature activations Φ^v at 5 different network layers $v = 1, \dots, N$:

$$\mathcal{L}_p(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{v=1}^N \|\Phi^v(\mathbf{y}) - \Phi^v(\hat{\mathbf{y}})\|_2. \quad (1)$$

This loss penalizes differences in low- mid- and high-level feature statistics, captured by the respective network filters.

Style loss. As in [28], we use the Gram matrix criterion of [25] as an objective for training a feedforward network. We first compute the Gram matrix of neuron activations delivered by the VGG network Φ at layer v for an image \mathbf{x} :

$$\mathcal{G}^v(\mathbf{x})_{c,c'} = \sum_{h,w} \Phi_c^v(\mathbf{x})[h,w] \Phi_{c'}^v(\mathbf{x})[h,w] \quad (2)$$

where h and w are horizontal and vertical pixel coordinates and c and c' are feature maps of layer v . The style loss is given by the sum of Frobenius norms for the difference between the per-layer Gram matrices \mathcal{G}^v of the two inputs:

$$\mathcal{L}_{\text{style}}(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{v=1}^B \|\mathcal{G}^v(\mathbf{y}) - \mathcal{G}^v(\hat{\mathbf{y}})\|_F. \quad (3)$$

Adversarial loss. We use adversarial training to penalize any detectable differences between the generated and real samples. Since global structural properties are largely settled thanks to DensePose conditioning, we opt for the patch-GAN [12] discriminator, which operates locally and picks up differences between texture patterns. The discriminator [12, 7] takes as an input \mathbf{z} , a combination of the source image and the DensePose results on the target image, and either the target image \mathbf{y} (real) or the generated output (fake) $\hat{\mathbf{y}}$. We want fake samples to be indistinguishable from real ones – as such we optimize the following objective:

$$L_{\text{GAN}} = \underbrace{\frac{1}{2} \mathbb{E}_{\mathbf{z}} [l(D(\mathbf{z}, \mathbf{y}) - 1)] + \frac{1}{2} \mathbb{E}_{\mathbf{z}} [l(D(\mathbf{z}, \hat{\mathbf{y}}))]}_{\text{Discriminator}} + \underbrace{\frac{1}{2} \mathbb{E}_{\mathbf{z}} [l(D(G(\mathbf{z})) - 1)]}_{\text{Generator}}, \quad (4)$$

where we use $l(x) = x^2$ as in the Least Squares GAN (LSGAN) work of [26].

4 Experiments

We perform our experiments on the DeepFashion dataset (In-shop Clothes Retrieval Benchmark) [19] that contains 52,712 images of fashion models demonstrating 13,029 clothing items in different poses. All images are provided at a resolution of 256×256 and contain people captured over a uniform background. Following [5] we select 12,029 clothes for training and the remaining 1,000 for testing. For the sake of direct comparison with state-of-the-art keypoint-based methods, we also remove all images where the keypoint detector of [39] does not detect any body joints. This results in 140,110 training and 8,670 test pairs.

In the supplementary material we provide results on the large scale MVC dataset [20] that consists of 161,260 images of resolution 1920×2240 crawled from several online shopping websites and showing front, back, left, and right views for each clothing item.

4.1 Implementation details

DensePose estimator. We use a fully convolutional network (FCN) similar to the one used as a teacher network in [1]. The FCN is a ResNet-101 trained on cropped person instances from the COCO-DensePose dataset. The output consists of 2D fields representing body segments (I) and $\{U, V\}$ coordinates in coordinate spaces aligned with each of the semantic parts of the 3D model.

Training parameters. We train the network and its submodules with Adam optimizer with initial learning rate $2 \cdot 10^{-4}$ and $\beta_1=0.5$, $\beta_2=0.999$ (no weight decay). For speed, we pretrain the predictive module and the inpainting module separately and then train the blending network while finetuning the whole combined architecture end-to-end; DensePose network parameters remain fixed. In all experiments, the batch size is set to 8 and training proceeds for 40 epochs. The balancing weights λ between different losses in the blending step (described in Sec. 3.4) are set empirically to $\lambda_{\ell_1}=1$, $\lambda_p=0.5$, $\lambda_{\text{style}}=5 \cdot 10^5$, $\lambda_{\text{GAN}}=0.1$.

4.2 Evaluation metrics

As of today, there exists no general criterion allowing for adequate evaluation of the generated image quality from the perspective of both structural fidelity and photorealism. We therefore adopt a number of separate structural and perceptual metrics widely used in the community and report our joint performance on them.

Structure. The geometry of the generations is evaluated using the perception-correlated *Structural Similarity* metric (SSIM) [40]. We also exploit its multi-scale variant MS-SSIM [44] to estimate the geometry of our predictions at a number of levels, from body structure to fine clothing textures.

Image realism. Following previous works, we provide the values of *Inception scores (IS)* [41]. However, as repeatedly noted in the literature, this metric is of limited relevance to the problem of within-class object generation, and we do not wish to draw strong conclusions from it. We have empirically observed its instability and high variance with respect to the perceived quality of generations

Table 1. Quantitative comparison with the state-of-the-art methods on the DeepFashion dataset [19] according to the Structural Similarity (SSIM) [40], Inception Score (IS) [41] and detection score (DS) [5] metrics. Our *best structure* model corresponds to the ℓ_1 loss, the *highest realism* model corresponds to the style loss training (see the text and Table 4). Our *balanced* model is trained using the full combination of losses.

Model	SSIM	IS	DS
Disentangled [42]	0.614	3.29	–
VariGAN [43]	0.620	3.03	–
G1+G2+D [4]	0.762	3.09	–
DSC [5]	0.761	3.39	0.966
Ours (best structure)	0.796	3.17	0.971
Ours (highest realism)	0.777	3.67	0.969
Ours (balanced)	0.785	3.61	0.971

and structural similarity. We also note that the ground truth images from the DeepFashion dataset have an average IS of 3.9, which indicates low degree of realism of this data according to the IS metric (for comparison, IS of CIFAR-10 is 11.2 [41] with best image generation methods achieving IS of 8.8 [2]).

In addition, we perform additional evaluation using *detection scores (DS)* [5] reflecting the similarity of the generations to the *person* class. Detection scores correspond to the maximum of confidence of the PASCAL-trained SSD detector [45] in the person class taken over all detected bounding boxes.

4.3 Comparison with the state-of-the-art

We first compare the performance of our framework to a number of recent methods proposed for the task of *keypoint guided* image generation or multi-view synthesis. Table 1 shows a significant advantage of our pipeline in terms of structural fidelity of obtained predictions. This holds for the whole range of tested network configurations and training setups (see Table 4). In terms of perceptual quality expressed through IS, the output generations of our models are of higher quality or directly comparable with the existing works. Some qualitative results of our method (corresponding to the *balanced* model in Table 1) and the best performing state-of-the-art approach [5] are shown in Fig. 4.

We have also performed a user study on Amazon Mechanical Turk, following the protocol of [5]: we show 55 real and 55 generated images in a random order to 30 users for one second. As the experiment of [5] was done with the help of fellow researchers and not AMT users, we perform an additional evaluation of images generated by [5] for consistency, using the official public implementation. We perform three evaluations, shown in Table ??: *Realism* asks users if the image is real or fake. *Anatomy* asks if a real, or generated image is anatomically plausible. *Pose* shows a pair of a target and a generated image and asks if they are in the same pose. The results (correctness, in %) indicate that generations




Fig. 4. Qualitative comparison with the state-of-the-art Deformable GAN (DSC) method of [5]. Each group shows the input, the target image, DSC predictions [5], predictions obtained with our full model. We observe that even though our cloth texture is occasionally not as sharp, we better retain face, gender, and skin color information.

of [5] have higher degree of perceived realism, but our generations show improved pose fidelity and higher probability of overall anatomical plausibility.

4.4 Effectiveness of different body representations

In order to clearly measure the effectiveness of the DensePose-based conditioning, we first compare to the performance of the ‘black box’, predictive module when used in combination with more traditional body representations, such as background/foreground masks, body part segmentation maps or body landmarks.



Fig. 5. Typical failures of keypoint-based pose transfer (top) in comparison with DensePose conditioning (bottom) indicate disappearance of limbs, discontinuities, collapse of 3D geometry of the body into a single plane and confusion in ordering in depth.

Table 2. On effectiveness of different body representations as a ground for pose transfer. The DensePose representation results in the highest structural quality.

Model	SSIM	MS-SSIM	IS
Foreground mask	0.747	0.710	3.04
Body part segmentation	0.776	0.791	3.22
Body keypoints	0.762	0.774	3.09
DensePose $\{I, U, V\}$	0.792	0.821	3.09
DensePose $\{\text{one-hot } I, U, V\}$	0.782	0.799	3.32

As a segmentation map we take the index component of DensePose and use it to form a one-hot encoding of each pixel into a set of class specific binary masks. Accordingly, as a background/foreground mask, we simply take all pixels with positive DensePose segmentation indices. Finally, following [5] we use [39] to obtain body keypoints and one-hot encode them.

In each case, we train the predictive module by concatenating the source image with a corresponding representation of the source and the target poses which results in 4 input planes for the mask, 27 for segmentation maps and 21 for the keypoints.

The corresponding results shown in Table 2 demonstrate a clear advantage of fine-grained dense conditioning over the sparse, keypoint-based, or coarse, segmentation-based, representations.

Complementing these quantitative results, typical failure cases of keypoint-based frameworks are demonstrated in Figure 5. We observe that these shortcomings are largely fixed by switching to the DensePose-based conditioning.

4.5 Ablation study on architectural choices

Table 3 shows the contribution of each of the predictive module, warping module, and inpainting autoencoding blocks in the final model performance. For

Table 3. Contribution of each of the functional blocks of the framework

Model	SSIM	MS-SSIM	IS
predictive module only	0.792	0.821	3.09
predictive + blending (=self-refinement)	0.793	0.821	3.10
predictive + warping + blending	0.789	0.814	3.12
predictive + warping + inpainting + blending (full)	0.796	0.823	3.17

**Fig. 6.** Effects of training with different loss terms and their weighted combinations.

these experiments, we use only the reconstruction loss \mathcal{L}_{ℓ_1} , factoring out fluctuations in the performance due to instabilities of GAN training. As expected, including the warping branch in the generation pipeline results in better performance, which is further improved by including the inpainting in the UV space. Qualitatively, exploiting the inpainted representation has two advantages over the direct warping of the partially observed texture from the source pose to the target pose: first, it serves as an additional prior for the fusion pipeline, and, second, it also prevents the blending network from generating clearly visible sharp artifacts that otherwise appear on the borders of partially observed segments of textures.

4.6 Ablation study on supervision objectives

In Table 4 we analyze the role of each of the considered terms in the composite loss function used at the final stage of the training, while providing indicative results in Fig. 6.

The perceptual loss \mathcal{L}_p is most correlated with the image structure and least correlated with the perceived realism, probably due to introduced textural artifacts. At the same time, the style loss $\mathcal{L}_{\text{style}}$ produces sharp and correctly textured patterns while hallucinating edges over uniform regions. Finally, adversarial training with the loss \mathcal{L}_{GAN} tends to prioritize visual plausibility but often disregarding structural information in the input. This justifies the use of all these complementary supervision criteria in conjunction, as indicated in the last entry of Table 4.

Table 4. Comparison of different loss terms used at the final stage of the training. The perceptual loss is best correlated with the structure, and the style loss with IS. The combined model provides an optimal balance between the extreme solutions.

Model	SSIM	MS-SSIM	IS
$\{\mathcal{L}_{\ell_1}\}$	0.796	0.823	3.17
$\{\mathcal{L}_{\ell_1}, \mathcal{L}_p\}$	0.791	0.822	3.26
$\{\mathcal{L}_{\ell_1}, \mathcal{L}_{\text{style}}\}$	0.777	0.815	3.67
$\{\mathcal{L}_{\ell_1}, \mathcal{L}_p, \mathcal{L}_{\text{style}}\}$	0.784	0.820	3.41
$\{\mathcal{L}_{\ell_1}, \mathcal{L}_{\text{GAN}}\}$	0.771	0.807	3.39
$\{\mathcal{L}_{\ell_1}, \mathcal{L}_p, \mathcal{L}_{\text{GAN}}\}$	0.789	0.820	3.33
$\{\mathcal{L}_{\ell_1}, \mathcal{L}_{\text{style}}, \mathcal{L}_{\text{GAN}}\}$	0.787	0.820	3.32
$\{\mathcal{L}_{\ell_1}, \mathcal{L}_p, \mathcal{L}_{\text{style}}, \mathcal{L}_{\text{GAN}}\}$	0.785	0.807	3.61

5 Conclusion

In this work we have introduced a two-stream architecture for pose transfer that exploits the power of dense human pose estimation. We have shown that dense pose estimation is a clearly superior conditioning signal for data-driven human pose estimation, and also facilitates the formulation of the pose transfer problem in its natural, body-surface parameterization through inpainting. In future work we intend to further pursue the potential of this method for photorealistic image synthesis [2, 6] as well as the treatment of more categories.

References

1. Güler, R.A., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. In: CVPR. (2018)
2. Karras, T., Aila, T., Samuli, L., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. In: ICLR. (2018)
3. Lassner, C., Pons-Moll, G., Gehler, P.V.: A generative model of people in clothing. In: ICCV. (2017)
4. Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. In: NIPS. (2017)
5. Siarohin, A., Sangineto, E., Lathuiliere, S., Sebe, N.: Deformable gans for pose-based human image generation. In: CVPR. (2018)
6. Chen, Q., Koltun, V.: Photographic image synthesis with cascaded refinement networks. In: ICCV. (2017)
7. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Jan, K., Bryan, C.: High-resolution image synthesis and semantic manipulation with conditional gans. In: CVPR. (2018)
8. Rossler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Niener, M.: Face-forensics: A large-scale video dataset for forgery detection in human faces. In: arXiv:1803.09179v1. (2018)
9. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Weng, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: CVPR. (2017)

10. Lample, G., Zeghidour, N., Usunier, N., Bordes, A., Denoyer, L., Ranzato, M.: Fader networks: Manipulating images by sliding attributes. In: NIPS. (2017)
11. Shu, Z., Yumer, E., Hadap, S., Sunkavalli, K., Shechtman, E., Samaras, D.: Neural face editing with intrinsic image disentangling. In: CVPR. (2017)
12. Isola, P., Zhu, J., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR. (2017)
13. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)* **34**(6) (October 2015) 248:1–248:16
14. Bogu, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In: ECCV. (2016)
15. Lassner, C., Romero, J., Kiefel, M., Bogu, F., Black, M.J., Gehler, P.V.: Unite the people: Closing the loop between 3d and 2d human representations. In: ICCV. (2017)
16. Varol, G., Romero, J., Martin, X., Mahmood, N., Black, M.J., Laptev, I., Schmid, C.: Learning from Synthetic Humans. In: CVPR. (2017)
17. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: CVPR. (2018)
18. Guler, R.A., Trigeorgis, G., Antonakos, E., Snape, P., Zafeiriou, S., Kokkinos, I.: Densereg: Fully convolutional dense shape regression in-the-wild. In: CVPR. (2017)
19. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: CVPR. (2016)
20. Liu, K.H., Chen, T.Y., Chen, C.S.: A dataset for view-invariant clothing retrieval and attribute prediction. In: ICMR. (2016)
21. Hinton, G.E., Salakhutdinov, R.R.: Reducing the dimensionality of data with neural networks. *Science* **313**(5786) (2006) 504–507
22. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR. (2014)
23. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS. (2014)
24. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. In: ICLR. (2016)
25. Gatys, L.A., Ecker, A.S., Bethge, M.: A neural algorithm of artistic style. In: CVPR. (2016)
26. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Smolley, S.P.: Least squares generative adversarial networks. In: ICCV. (2017)
27. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: ICML. (2017)
28. Johnson, J., Alahi, A., Li, F.F.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV. (2016)
29. Ulyanov, D., Lebedev, V., Vedaldi, A., Lempitsky, V.: Texture networks: Feed-forward synthesis of textures and stylized images. In: ICML. (2016)
30. Zhu, S., Fidler, S., Urtasun, R., Lin, D., Loy, C.C.: Be your own prada: Fashion synthesis with structural coherence. In: ICCV. (2017)
31. Zhao, B., Wu, X., Cheng, Z.Q., Liu, H., Feng, J.: Multi-view image generation from a single-view. In: ACM on Multimedia Conference. (2018)
32. Pathak, D., Krähenbühl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR. (2016)
33. Yeh, R.A., Chen, C., Lim, T., Hasegawa-Johnson, M., Do, M.N.: Semantic image inpainting with perceptual and contextual losses. In: CVPR. (2017)

34. Saito, S., Wei, L., Hu, L., Nagano, K., Li, H.: Photorealistic facial texture inference using deep neural networks. In: CVPR. (2017)
35. Deng, J., Cheng, S., Xue, N., Zhou, Y., Zafeiriou, S.: UV-GAN: adversarial facial UV map completion for pose-invariant face recognition. In: CVPR. (2018)
36. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In: CVPR. (2017)
37. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: NIPS. (2015)
38. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR. (2015)
39. Cao, Z., Simon, T., Wei, S., Sheikh, Y.: Realtime multiperson 2d pose estimation using part affinity fields. In: CVPR. (2017)
40. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. In: TIP. (2004)
41. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. In: NIPS. (2016)
42. Ma, L., Sun, Q., Georgoulis, S., Van Gool, L., Schiele, B., Fritz, M.: Disentangled person image generation. In: CVPR. (2018)
43. Zhao, B., Wu, X., Cheng, Z.Q., Liu, H., Feng, J.: Multi-view image generation from a single-view. In: ACM on Multimedia Conference. (2018)
44. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multi-scale structural similarity for image quality assessment. In: ACSSC. (2003)
45. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: ECCV. (2016)