

# 360-Degree Textures of People in Clothing from a Single Image

Verica Lazova

Eldar Insafutdinov

Gerard Pons-Moll

Max Planck Institute for Informatics  
Saarland Informatics Campus, Germany

{vlazova, eldar, gpons}@mpi-inf.mpg.de

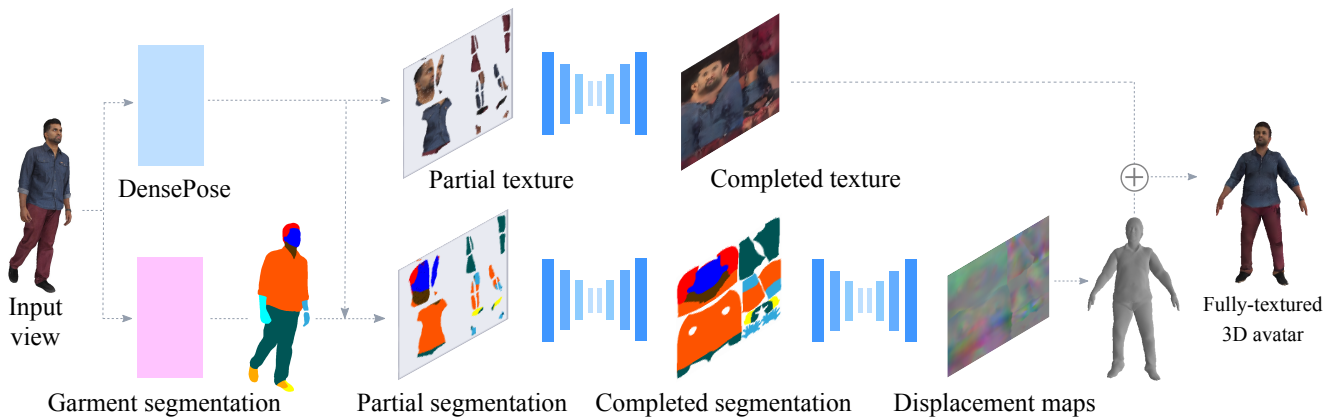


Figure 1: Given a single view of a person we predict a complete texture map in the UV space, complete clothing segmentation as well as a displacement map for the SMPL model [41], which we then combine to obtain a fully-textured 3D avatar.

## Abstract

*In this paper we predict a full 3D avatar of a person from a single image. We infer texture and geometry in the UV-space of the SMPL model using an image-to-image translation method. Given partial texture and segmentation layout maps derived from the input view, our model predicts the complete segmentation map, the complete texture map, and a displacement map. The predicted maps can be applied to the SMPL model in order to naturally generalize to novel poses, shapes, and even new clothing. In order to learn our model in a common UV-space, we non-rigidly register the SMPL model to thousands of 3D scans, effectively encoding textures and geometries as images in correspondence. This turns a difficult 3D inference task into a simpler image-to-image translation one. Results on rendered scans of people and images from the DeepFashion dataset demonstrate that our method can reconstruct plausible 3D avatars from a single image. We further use our model to digitally change pose, shape, swap garments between people and edit clothing. To encourage research in this direction we will make the source code available for research purpose [5].*

## 1. Introduction

3D models of humans, comprising personalized surface geometry, and full *texture* are required for numerous applications including VR/AR, gaming, entertainment or human tracking for surveillance. Methods capable of recovering such 3D models from a single image would democratize the acquisition process, and allow people to easily recover avatars of themselves.

While there is extensive work on 3D human pose and shape recovery (surface geometry) from a single image, almost no work addresses the problem of predicting a complete texture of a person from a single image. We, humans, can do it to some degree, because we can guess what the person might look like from another view. We make these guesses effortlessly because we have seen people from many angles, and we have built a mental model of typical correlations. For example, the texture patterns within a garment are often repetitive; skin color and hair is roughly homogeneous; and the appearance of the left and right shoe is typically the same. In this work, we introduce a model that automatically learns such correlations from real appearances of people. Specifically, our model predicts a full (360°) texture map of the person in the UV space of the

SMPL model [41] given a partial texture map derived from the visible part in an image, see Fig. 1. Instead of learning an image-based model that has to generalize to every possible pose, shape and viewpoint [30, 37, 43], our idea is to learn to complete the full 3D texture and generalize to new poses, shapes and viewpoints with SMPL and traditional rendering (Figure 9). Learning in a common UV-space requires texture maps of people in clothing in good correspondence, which is a highly non-trivial task. To that end, we non-rigidly deform the SMPL body model to 4541 static 3D scans of people in different poses and clothing. This effectively brings the 3D appearances of people into a common 2D UV-map representation, see Fig. 2. Learning from this data has several advantages. First, since every pixel in a UV-map corresponds to one location on the body surface, it is easier for a neural network to implicitly learn localized and specialized models for particular parts of the body surface. For example, our model learns to complete different textures for the different garments, the face region and the skin, with clearly defined boundaries. Second, at test time, once the texture map is recovered, we can apply it to SMPL and animate it directly in 3D with any pose and shape and viewpoint to generate a coherent movement and appearance—without the risk of failing to generalize to these factors.

Additionally, our model can predict, on the SMPL UV map, a full clothing segmentation layout, from which we predict a *plausible* clothing geometry of the garments worn by the subject. The latter is a highly multi-modal problem where many clothing geometries can explain the same segmentation layout. However, as observed for other tasks [33] an image translation network is capable of producing a plausible output. To allow for further control, we can additionally easily edit the predicted segmentation layout in order to control the shape and appearance of the clothing. In particular, we can modify the length of the sleeves in shirts and t-shirts, and interpolate from shorts to pants and the other way around.

We train our method on our newly created dataset (*360° People Textures*) consisting of 4541 scans registered to a common SMPL body template. Our results generalize on real images and show that our model is capable of completing textures of people, often reproducing clothing boundaries and completing garment textures and skin. In summary, our model can reconstruct a full 3D animatable avatar from a single image, including texture, geometry and segmentation layout for further control.

## 2. Related work

**Pose guided image and video generation** Given a source image and a target 2D pose, image-based methods [43, 44, 22, 59, 49] produce an image with the source appearance in the target pose. To deal with pixel miss-

alignments, it is helpful to transform the pixels in the original image to match the target pose within the network [60, 11, 42]. Following similar ideas to growing GANs [36], high resolution anime characters can be generated [29]. Dis-occlusions and out of plane rotations are a problem for such methods.

Sharing our goal of incorporating more information about the human body during the generation process, some works [37, 49] also leverage the SMPL body model. By conditioning on a posed SMPL rendering, realistic images of people in the same pose can be generated [37]—the model however cannot maintain the appearance of the person when the pose changes. Recent approaches [49, 26] leverage DensePose [27] to map a source appearance to the SMPL surface, and condition the generated image on the target DensePose mask. Although their goal is generating an image, the model generates a texture map on the SMPL surface as an internal representation. During training, multiple views of the same subject are mapped to a custom designed UV-map, and the network is forced to complete the texture on the visible parts. While training from images is practical, DensePose was not designed to accommodate for hair and clothing deviating from the body, and significant misalignments might occur in the UV-map when DensePose fails or is inaccurate. Consequently, texture map completion results are limited by the ability of DensePose to parse images during training. We use DensePose solely to create the partial texture map from the input view, but train from high quality aligned and complete texture maps, which are obtained by registering SMPL to 3D scans. This has several advantages. By virtue of the registration, the textures are very well aligned in the SMPL-UV space, and cover the full extent of the appearance including clothing and hair. This allows us to generalize beyond poses seen in a particular dataset such as DeepFashion [40]. Furthermore, we go beyond [49, 26] and predict a full 3D textured avatar including geometry and 3D segmentation layout for further control.

While image-based methods are able to generate plausible images, many applications (VR/AR, gaming) require a full 3D representation of the person. Traditional graphics rendering could potentially be replaced by querying such models for every desired novel viewpoint and pose. Unfortunately, image-based models do not produce temporally coherent results for sequences of target poses. Recent work focus on video generation [18, 11, 73, 58], but these models are often trained specifically for a particular subject. Furthermore, the formation of more complex scenes with objects or multiple people would be a challenge for image-based methods.

**Try-on and conditional clothing synthesis** A growing number of recent works focus on conditioning the person image generation on a particular clothing item [30, 66], or



Figure 2: *360° People Textures*: Example registrations and texture maps used to train our model. After registering SMPL to the scans their appearances are encoded as texture maps in a common UV-space.

a description [78]. Other works demonstrated models capable of swapping the appearances between two different subjects [54, 75].

**Multi-view texture generation.** Texture generation is challenging even in the case of multi-view images. The difficulty lies in combining the partial textures created from different views by using blending [13, 20, 51, 61], mosaicing [12, 39, 50, 55], graph cuts [38], or flow based registration [21, 14, 65, 23, 76, 7] in order to reduce ghosting and stitching artifacts. A learning based model that leverages multi-view images for training [49, 26] will suffer similar problems. Hence, in this work, we learn from complete texture maps obtained from 3D registrations.

**3D person reconstruction from images** While promising, recent methods for 3D person reconstruction either require video as input [6, 7, 8], scans [74], do not allow control over pose, shape and clothing [48, 56], focus only on faces [72, 32, 63, 57, 47, 62], or only on garments [68].

### 3. Method

We introduce an image-to-image translation method for inferring texture and geometry of a person from one image. Our method predicts a full texture map, displacement map and full clothing segmentation for additional control. Our models are trained on large number of high quality, diverse 3D scans of people in clothing.

#### 3.1. Synthetic data generation

To generate our training set, we take 4541 static scans, and register non-rigidly the SMPL template to them to obtain complete textures aligned in the SMPL UV space. We will briefly review SMPL and explain our non-rigid registration procedure in 3.1.1.

##### 3.1.1 SMPL Body Model with Clothing

SMPL is a function  $M(\theta, \beta)$  that takes as input pose  $\theta \in \mathbb{R}^{72}$  and shape parameters  $\beta \in \mathbb{R}^{10}$  and outputs the  $N = 6890$  vertices of a mesh. SMPL applies deformations to the mean shape  $\mathbf{T}_\mu$  using a learned skinning function:

$$M(\beta, \theta) = W(T(\beta, \theta), J(\beta), \theta, \mathbf{W}) \quad (1)$$

$$T(\beta, \theta) = \mathbf{T}_\mu + B_s(\beta) + B_p(\theta), \quad (2)$$

where  $W$  is a linear blend-skinning function applied to the morphed shape  $T(\beta, \theta)$  based on the skeleton joints  $J(\beta)$ ; the morphed shape is obtained by applying pose-dependent deformations  $B_p(\theta)$  and shape-dependent deformations  $B_s(\beta)$ . The shape space  $B_s(\beta)$  was learned from undressed scans of people and therefore cannot accommodate for clothing. Since our goal is registering the SMPL surface to scans of people in clothing, we modify SMPL by adding a set of offsets  $\mathbf{D} \in \mathbb{R}^{3N}$  to the template:

$$M(\theta, \beta, \mathbf{D}) = W(T(\theta, \beta, \mathbf{D}), J(\beta), \theta, \mathbf{W}) \quad (3)$$

$$T(\theta, \beta, \mathbf{D}) = \mathbf{T}_\mu + B_s(\beta) + B_p(\theta) + \mathbf{D}, \quad (4)$$

which are responsible to explain clothing, hair and details beyond the shape space of SMPL.

**Non-rigid registration** of humans (Figure 3) is challenging even for scans without clothing [16, 53]. The scans in *360° People Textures* include subjects in clothing scanned in variety of poses. Registration without a good initialization of pose and shape would fail. To obtain a good initialization, we render the scans in 8 camera viewpoints around the vertical axis and run a 2D pose joint detector [17] in each of the views. Then, we minimize the re-projection error between SMPL joint positions  $J(\theta, \beta)$  and the detected joints  $\mathbf{j}_{2D}^i$  with respect to pose and shape:

$$\arg \min_{\theta, \beta} \sum_i \|\pi_{C_i}(J(\theta, \beta)) - \mathbf{j}_{2D}^i\|_F^2 + \quad (5)$$

$$\lambda_\theta D_M(\theta; \mu_\theta, \Sigma_\theta) + \lambda_\beta D_M(\beta; \mu_\beta, \Sigma_\beta), \quad (6)$$

where  $\pi_{C_i}$  projects the 3D joints  $J(\theta, \beta)$  onto the view of camera  $C_i$ , and  $D_M(\cdot)$  is a Mahalanobis distance prior that we use to regularize pose and shape. This brings the model close to the scan, but it is not registered at this point because SMPL cannot explain details such as clothing, hair or jewellery. Let  $\mathcal{S}$  denote a scan we want to register to, and  $\mathcal{A} = \{\mathbf{A}, \mathbf{F}\}$  be the registration mesh with free-form vertices  $\mathbf{A} \in \mathbb{R}^{3N}$  and faces  $\mathbf{F}$  defining the same topology as SMPL. We obtain a registration by minimizing the following objective function:

$$\arg \min_{\theta, \beta, \mathbf{A}} \sum_{\mathbf{x}_s \in \mathcal{S}} \rho(\text{dist}(\mathcal{A}, \mathbf{x}_s)) + \sum_{i=0}^N w_i \|\mathbf{A}_i - M_i(\theta, \beta)\|_F^2 + \lambda_\theta D_M(\theta; \mu_\theta, \Sigma_\theta) + \lambda_\beta D_M(\beta; \mu_\beta, \Sigma_\beta), \quad (7)$$

where the first term measures the distance from every point in the scan  $\mathbf{x}_s \in \mathcal{S}$  to the surface of the registration  $\mathcal{A}$ , the second term forces the registration vertices  $\mathbf{A}_i$  to remain close to the model vertices  $M(\theta, \beta)_i$ , and the other terms




Figure 3: Registration: we bring all scan appearances into correspondence by non-rigidly deforming the SMPL model. This allows us to learn in a pose invariant space.

are the aforementioned priors. Here,  $\rho(\cdot)$  denotes a Geman-McClure robust cost function, and  $w_i$  are the weights that penalize deviations from the model more heavily for the vertices on the hands and the feet. The first term in Eq. (7) pulls the registration to the scan, while the second penalizes deviations from the body model ensuring the registration looks like a human. That allows to accurately bring all scans into correspondence.

### 3.2. 3D Model Generation

Our full pipeline for 3D model generation from an image is shown in Figure 1 and consists of three stages: texture completion, segmentation completion and geometry prediction. Given an image of a person and its garment segmentation we extract partial texture map as well as partial segmentation map. We have put together inpainting models to complete the garment segmentation and the texture in UV-space. The texture completion part recovers the full appearance of the person in the image. The completed segmentation is further used to predict the geometry. To obtain geometry details we predict a displacement map from the segmentation, which gives us a set of offsets from the average SMPL [41] body model, that correspond to clothing geometry. All three parts are trained separately and allow us to go from a single image to a 3D model. Additionally, the complete segmentation map allows us to edit the final model by changing the garment properties such as length or texture, or to completely redress the person.

#### 3.2.1 Texture Completion

The texture completion part recovers the full appearance of the person in the image. We use DensePose [27] to find pixel correspondences between the image and the SMPL body model, which we use to remap a partial texture map of the person in the SMPL texture template (Figure 1). We use an inpainting network based on image-to-image translation [33, 19, 67] that learns to complete the partial texture map. The inaccuracies of texture extraction with DensePose introduce distortions in the partial texture map. As a result,

the input and the output image for the texture completion network are not in perfect alignment. This makes our task different from the classical image inpainting, where the network can learn to copy the visible parts of the image to the target. We use an architecture for the generator based on residual blocks similar to [19, 77, 67]. We refer the reader to the supplementary material for the detailed description of the network layers.

We train the inpainting network using the conditional Generative Adversarial Network (GAN) approach with a fully-convolutional patch discriminator to discriminate between real and fake texture maps. The partial texture map  $x$  is fed through the generator  $G$  to infer the full texture map  $G(x)$ . The conditional discriminator  $D$  takes pair of partial and complete texture map and classifies them as real or fake, i.e.  $D(x, y) = 1$  and  $D(x, G(x)) = 0$ . We train the generator  $G$  by optimizing the compound objective  $L_{\text{tex}}$  consisting of a GAN, reconstruction, perceptual and dissimilarity term.

$$L_{\text{tex}} = \lambda_1 L_{\text{GAN}} + \lambda_2 L_{\text{recon}} + \lambda_3 L_{\text{perc}} + \lambda_4 L_{\text{DSSIM}} \quad (8)$$

where  $\lambda_1, \dots, \lambda_4$  are the weights for each of the terms. We use the Wasserstein GAN loss [10, 28] as our adversarial loss, which produced visually sharper results:

$$L_{\text{GAN}} = \mathbb{E}_{x,y}[D(x, y)] - \mathbb{E}_x[D(x, G(x))] \quad (9)$$

The reconstruction loss is an  $L1$  distance between the target  $y$  and the reconstructed texture map  $G(x)$ .

$$L_{\text{recon}}(G) = \mathbb{E}_{x,y} \|G(x) - y\|_1 \quad (10)$$

Additionally we use perceptual loss which involves matching deep features between the target and the reconstructed image. The features are extracted from the convolutional layers of a pretrained VGG-19 neural network,  $VGG_l(y)$  and  $VGG_l(G(x))$  for each layer  $l$ . The perceptual loss is the  $L1$  distance between features extracted from the target and the generated texture image.

$$L_{\text{perc}}(G) = \mathbb{E}_{x,y} \sum_l \|VGG_l(y) - VGG_l(G(x))\|_1 \quad (11)$$

Finally, as another perceptual term we minimize the dissimilarity index between the target and generated texture map:

$$L_{\text{DSSIM}}(G) = \frac{1 - \text{MSSIM}(y, G(x))}{2}, \quad (12)$$

where  $\text{MSSIM}$  is a Multiscale Structural Similarity index [69, 70], a structural similarity computed at 5 different scales. Please check the supplementary material for more details on the effect of each of the loss functions.



### 3.2.2 Segmentation completion

Similar to the texture completion stage, the segmentation completion starts with a partial clothing segmentation obtained from a single image, and the task is to obtain a valid full body clothing segmentation in UV space. An image-to-image translation model is trained to inpaint the missing part of the segmentation. The architecture for generator and discriminator is the same as in the texture completion network. We train the segmentation completion network by minimizing the  $L_{\text{seg}}$  loss:

$$L_{\text{seg}} = \lambda_1 L_{\text{GAN}} + \lambda_2 L_{\text{recon}} \quad (13)$$

where  $L_{\text{GAN}}$  and  $L_{\text{recon}}$  are defined in Equations 9 and 10. Finally, the inpainted segmentation is discretized by assigning the label for each pixel in a nearest-neighbour fashion.

### 3.2.3 Displacement map prediction

The final step of our pipeline is geometry prediction where the purpose is to generate geometry that corresponds to the clothes the person is wearing. For this purpose, to capture the clothing shape, we learn to generate vertex offsets from the average SMPL body conditioned on garment type, i.e the segmentation. The offsets are stored as displacement map, where each pixel location corresponds to a point on the SMPL model, and each pixel stores the normalized offset at that body location. This allows us to treat the displacements as images and apply the image-to-image translation method we described previously. A similar approach is used in [9], where the authors learn full displacement and normal maps from a partial texture. Our intent is to learn a model which is able to sample plausible geometry that fits the clothing segmentation. The predicted geometry might not exactly coincide with the image (eg. it is impossible to know from a segmentation in UV-space, if a T-shirt is tight or loose). Nevertheless, relying on the segmentation has advantages: it provides additional flexibility, making editing and adjusting the final 3D model straightforward. We discuss some of the model editing possibilities in Section 4.

As in the previous stages, the model for displacements prediction consists of a generator  $G$ , that takes a segmentation map as input and produces a displacement map at the output. The generator and the discriminator have the same architecture as in the texture completion scenario. We train the model by minimizing the  $L_{\text{disp}}$  objective:

$$L_{\text{disp}} = \lambda_1 L_{\text{GAN}} + \lambda_2 L_{\text{recon}} + \lambda_4 L_{\text{DSSIM}} \quad (14)$$

where  $L_{\text{GAN}}$ ,  $L_{\text{recon}}$  and  $L_{\text{DSSIM}}$  are defined in Eq. 9, 10 and 12.

The generated displacement map applied to the average SMPL body yields the full untextured 3D model. Applying the generated texture from the texture completion part on the model, gives us the full 3D avatar (Figure 1)

## 4. Experiments and Results

In the following section we give more details on our  $360^\circ$  *People Textures* dataset and we describe our experiments. For details regarding the training process please check the supplementary material. We have made an attempt for quantitative evaluations, however the typical metrics for evaluating GAN models known in the literature did not correspond to the human-perceived quality of the generated texture. Therefore, here we only present our qualitative results.

### 4.1. Dataset

We trained our models on a dataset of 4541 3D scans: 230 come from Renderpeople [1]; 116 from XYZ [2]; 45 are courtesy of Treedy’s [4] and the rest 4150 are courtesy of Twindom [3]. We used all of them to train the texture inpainting part, plus additional 242 texture maps from the SURREAL [64] dataset. Each of the scans is rendered from 10 roughly frontal views, with the random horizontal rotation of the body sampled from the range  $[-\frac{\pi}{2}, \frac{\pi}{2}]$ . Out of these we used 40 scans for validation and 10 for testing. Additionally we test on the People Snapshot Dataset [8] and on preselected 526 images from the DeepFashion [40] dataset that satisfy the criteria of fully visible person from a roughly frontal view without overly complex clothing or poses. We use DensePose to extract  $256 \times 256$  partial texture maps which we aim to complete.

In order to obtain a full segmentation map first we render the scan from 20 different views. For each of the rendered images we run the clothing segmentation method in [25]. Each of the segmented images is projected back to the scan and remapped into a partial segmentation UV-map. All the partial segmentation maps are then stitched together using a graph-cut based method to finally obtain the full segmentation map. The segmentation maps obtained in this way are used as ground-truth for the segmentation completion method. The partial segmentation maps used as input are extracted with DensePose.

For the geometry prediction part we have selected a subset of 2056 scans for which we could obtain good ground truth segmentation and displacement maps. Out of these 30 scans were taken out for validation. We obtain the segmentation as explained above and the ground truth displacement maps are obtained directly from the scan registrations.

### 4.2. Results

Figure 4 shows sample results of going from real images to 3D models on the People Snapshot and DeepFashion datasets. Each row shows the original image, it’s garment segmentation, the recovered full texture and segmentation, predicted geometry and the fully textured avatar. Predicting the body shape is not part of our contribution, but an off-the-shelf method such as [52, 35, 15] could easily be deployed to obtain a more faithful representation of the body



Figure 4: Generalization of our method to real images. Left to right: real image, segmented image, complete texture map, complete segmentation map, generated untextured avatar, generated textured avatar.



Figure 5: Comparison of our single view method (bottom row) to the monocular video method of [6] (top row)

shape. Additionally, in Figure 5 we show comparison of our single-view method to the monocular video method of [6]. While our results look comparable and visually pleasing, should be noted that the method in [6] produces higher resolution reconstructions. The advantage of our method is that it requires just a single view.

The full clothing segmentation map gives us the flexibility to edit and adjust the model in different ways. Using the segmentation, one can change the texture of a specific garment and have it completed in a different style. Figure 6 shows examples where the original T-shirt of the subject is edited with samples of different textures. The network hallucinates wrinkles to make the artificially added texture look more cloth-like.



Figure 6: Segmentation-based texture editing: the 3D model reconstructed from an image is shown in the top row. The subsequent rows show the edited partial texture and the reconstructed model with different texture for the t-shirt

Additionally, the segmentation could be used for clothes re-targeting. Having two partial texture maps from two images one can rely on the segmentation to exchange garments between the subjects. The edited partial texture and segmentation are then completed, new geometry is produced and the final result is a redressed avatar. An example of this application is shown in Figure 7.



Figure 7: Garment re-targeting results: Four subjects are shown in the top row with their 3D reconstruction. In each of the subsequent rows the subjects are reconstructed wearing the selected garments from the person in the first image.

Finally, one can use the segmentation to edit the length of a garment, transforming shorts to pants or blouse to a T-shirt. We train a separate network that only completes the arms and the legs of the person conditioned on the segmentation. Therefore, once we have reconstructed a model, we



can run it through the editing network to change the length of the sleeves or the pants. An example of this scenario is shown in Figure 8.



Figure 8: Model editing results: Models reconstructed from an image are shown in the top row. In each of the subsequent rows the models are edited such that the length of the sleeves and pants matches the segmentation on the left.

For additional results and better preview please check the supplementary material and video <sup>1</sup>.

## 5. Conclusions and Future Work

We presented a method for completion of full texture and geometry from single image. Our key idea is to turn a hard 3D inference problem into an image-to-image translation which is amenable to CNNs by encoding appearance, geometry and clothing layout on a common SMPL UV-space. By learning in such space we obtain visually pleasing and realistic 3D model reconstructions. To allow for control we decompose the predictions into factors: texture maps to control appearance, displacement maps for geometry, segmentation layout to control clothing shape and appearance. We have demonstrated how to swap clothing between 3D avatars just from the image and how to edit the garment geometry and texture.

There are many avenues for future work. First, our model cannot predict clothing with vastly different topol-



Figure 9: Image-based reposing methods cannot handle challenging target poses outside of the training set. First row shows the original scan reposed. The second row shows image based reposing using [60]. In the last row is the full model obtained with our method, reposed. The input image is the first one in each row.

ogy than the human body, eg. skirts or dresses. Implicit function based representations [71, 56, 46, 48, 31] might be beneficial to deal with different topologies, but they do not allow control. Although it is remarkable that our model can predict the occluded appearance of the person, the model struggles to predict high frequency detail and complex texture patterns. We have extensively experimented with style-based losses and models [24, 34], but the results are either not photo-realistic, lack control or need to be re-trained per texture. In the future, we would like to investigate models capable of producing photo-realistic textures, and complex clothing without losing control.

There is a strong trend in the community towards data-driven image-based models for novel viewpoint synthesis [44, 43], virtual try-on [30, 66], reposing [18] and reshaping with impressive results, suggesting 3D modelling can be by-passed. In this paper, we have demonstrated that full 3D completion and 3D reasoning has several advantages, particularly in terms of control and ability to generalize to novel poses and views. In the future, we plan to explore hybrid methods that combine powerful 3D representations with neural based rendering [45].

**Acknowledgements:** We thank Christoph Lassner, Bernt Schiele and Christian Theobald for provided insight and expertise at the early stages of the project; Thiemo Alldieck and Bharat Lal Bhatnagar for the in-depth discussions and assistance in our dataset preparation. Special thanks should be given to Mohamed Omran for contributing with an abundance of valuable ideas and most honest critiques.

This work is partly funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 409792180 (Emmy Noether Programme, project: Real Virtual Humans) and the Google Faculty Research Award.

<sup>1</sup><https://youtu.be/h-HyFq2rY00>



## References

- [1] <https://renderpeople.com/>. 5
- [2] <https://secure.axyz-design.com/>. 5
- [3] <https://web.twindom.com/>. 5
- [4] <https://www.treedys.com/>. 5
- [5] <http://virtualhumans.mpi-inf.mpg.de/360tex>. 1
- [6] T. Alldieck, M. Magnor, B. L. Bhatnagar, C. Theobalt, and G. Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, jun 2019. 3, 7
- [7] T. Alldieck, M. Magnor, W. Xu, C. Theobalt, and G. Pons-Moll. Detailed human avatars from monocular video. In *International Conference on 3D Vision (3DV)*, sep 2018. 3
- [8] T. Alldieck, M. Magnor, W. Xu, C. Theobalt, and G. Pons-Moll. Video based reconstruction of 3d people models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3, 5
- [9] T. Alldieck, G. Pons-Moll, C. Theobalt, and M. Magnor. Tex2shape: Detailed full human body geometry from a single image. *arXiv preprint arXiv:1904.08645*, 2019. 5
- [10] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017. 4
- [11] G. Balakrishnan, A. Zhao, A. V. Dalca, F. Durand, and J. Guttag. Synthesizing images of humans in unseen poses. *arXiv preprint arXiv:1804.07739*, 2018. 2
- [12] A. Baumberg. Blending images for texturing 3d models. In *British Machine Vision Conference*, volume 3, page 5. Citeseer, 2002. 3
- [13] F. Bernardini, I. M. Martin, and H. Rushmeier. High-quality texture reconstruction from multiple scans. *IEEE Transactions on Visualization and Computer Graphics*, 7(4):318–332, 2001. 3
- [14] S. Bi, N. K. Kalantari, and R. Ramamoorthi. Patch-based optimization for image-based texture mapping. *ACM Transactions on Graphics*, 36(4), 2017. 3
- [15] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *European Conference on Computer Vision*. Springer International Publishing, Oct. 2016. 5
- [16] F. Bogo, J. Romero, G. Pons-Moll, and M. J. Black. Dynamic FAUST: Registering human bodies in motion. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2017. 3
- [17] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2017. 3
- [18] C. Chan, S. Ginosar, T. Zhou, and A. A. Efros. Everybody dance now. *arXiv preprint arXiv:1808.07371*, 2018. 2, 8
- [19] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8789–8797, 2018. 4
- [20] P. E. Debevec, C. J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *Annual Conf. on Computer Graphics and Interactive Techniques*, pages 11–20. ACM, 1996. 3
- [21] M. Eisemann, B. De Decker, M. Magnor, P. Bekaert, E. De Aguiar, N. Ahmed, C. Theobalt, and A. Sellent. Floating textures. *Computer Graphics Forum*, 27(2):409–418, 2008. 3
- [22] P. Esser, E. Sutter, and B. Ommer. A variational u-net for conditional appearance and shape generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8857–8866, 2018. 2
- [23] Y. Fu, Q. Yan, L. Yang, J. Liao, and C. Xiao. Texture mapping for 3d reconstruction with rgb-d sensor. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2018. 3
- [24] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 8
- [25] K. Gong, X. Liang, Y. Li, Y. Chen, M. Yang, and L. Lin. Instance-level human parsing via part grouping network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 770–785, 2018. 5
- [26] A. Grigorev, A. Sevastopolsky, A. Vakhitov, and V. Lempitsky. Coordinate-based texture inpainting for pose-guided image generation. *arXiv preprint arXiv:1811.11459*, 2018. 2, 3
- [27] R. A. Güler, N. Neverova, and I. Kokkinos. Densepose: Dense human pose estimation in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 4
- [28] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017. 4
- [29] K. Hamada, K. Tachibana, T. Li, H. Honda, and Y. Uchida. Full-body high-resolution anime generation with progressive structure-conditional generative adversarial networks. *arXiv preprint arXiv:1809.01890*, 2018. 2
- [30] X. Han, Z. Wu, Z. Wu, R. Yu, and L. S. Davis. Viton: An image-based virtual try-on network. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 8
- [31] Z. Huang, T. Li, W. Chen, Y. Zhao, J. Xing, C. LeGendre, L. Luo, C. Ma, and H. Li. Deep volumetric video from very sparse multi-view performance capture. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 336–354, 2018. 8
- [32] L. Huynh, W. Chen, S. Saito, J. Xing, K. Nagano, A. Jones, P. Debevec, and H. Li. Mesoscopic facial geometry inference using deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8407–8416, 2018. 3
- [33] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 2, 4
- [34] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European*

- conference on computer vision, pages 694–711. Springer, 2016. 8
- [35] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [36] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2017. 2
- [37] C. Lassner, G. Pons-Moll, and P. V. Gehler. A generative model of people in clothing. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, Piscataway, NJ, USA, oct 2017. IEEE. 2
- [38] V. Lempitsky and D. Ivanov. Seamless mosaicing of image-based texture maps. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 1–6. IEEE, 2007. 3
- [39] H. P. Lensch, W. Heidrich, and H.-P. Seidel. A silhouette-based algorithm for texture registration and stitching. *Graphical Models*, 63(4):245–262, 2001. 3
- [40] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1096–1104, 2016. 2, 5
- [41] M. M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics*, 34(6):248:1–248:16, Oct. 2015. 1, 2, 4
- [42] D. Lorenz, L. Bereska, T. Milbich, and B. Ommer. Unsupervised part-based disentangling of object shape and appearance. *arXiv preprint arXiv:1903.06946*, 2019. 2
- [43] L. Ma, X. Jia, Q. Sun, B. Schiele, T. Tuytelaars, and L. Van Gool. Pose guided person image generation. In *Advances in Neural Information Processing Systems*, pages 406–416, 2017. 2, 8
- [44] L. Ma, Q. Sun, S. Georgoulis, L. Van Gool, B. Schiele, and M. Fritz. Disentangled person image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 99–108, 2018. 2, 8
- [45] R. Martin-Brualla, R. Pandey, S. Yang, P. Pidlypenskyi, J. Taylor, J. Valentin, S. Khamsi, P. Davidson, A. Tkach, P. Lincoln, et al. Lookingood: enhancing performance capture with real-time neural re-rendering. *arXiv preprint arXiv:1811.05029*, 2018. 8
- [46] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3d reconstruction in function space. *arXiv preprint arXiv:1812.03828*, 2018. 8
- [47] K. Nagano, J. Seo, J. Xing, L. Wei, Z. Li, S. Saito, A. Agarwal, J. Fursund, H. Li, R. Roberts, et al. pagan: real-time avatars using dynamic textures. In *SIGGRAPH Asia 2018 Technical Papers*, page 258. ACM, 2018. 3
- [48] R. Natsume, S. Saito, Z. Huang, W. Chen, C. Ma, H. Li, and S. Morishima. Siclope: Silhouette-based clothed people. *arXiv preprint arXiv:1901.00049*, 2018. 3, 8
- [49] N. Neverova, R. A. Güler, and I. Kokkinos. Dense pose transfer. *arXiv preprint arXiv:1809.01995*, 2018. 2, 3
- [50] W. Niem and J. Wingbermuhle. Automatic reconstruction of 3d objects using a mobile monoscopic camera. In *Proc. Inter. Conf. on Recent Advances in 3-D Digital Imaging and Modeling*, pages 173–180. IEEE, 1997. 3
- [51] E. Ofek, E. Shilat, A. Rappoport, and M. Werman. Multiresolution textures from image sequences. *IEEE Computer Graphics and Applications*, 17(2):18–29, Mar. 1997. 3
- [52] M. Omran, C. Lassner, G. Pons-Moll, P. Gehler, and B. Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *2018 International Conference on 3D Vision (3DV)*, pages 484–494. IEEE, 2018. 5
- [53] G. Pons-Moll, J. Romero, N. Mahmood, and M. J. Black. Dyna: a model of dynamic human shape in motion. *ACM Transactions on Graphics*, 34:120, 2015. 3
- [54] A. Raj, P. Sangkloy, H. Chang, J. Hays, D. Ceylan, and J. Lu. Swapnet: Image based garment transfer. In *European Conference on Computer Vision*, pages 679–695. Springer, Cham, 2018. 3
- [55] C. Rocchini, P. Cignoni, C. Montani, and R. Scopigno. Multiple textures stitching and blending on 3d objects. In *Rendering Techniques 99*, pages 119–130. Springer, 1999. 3
- [56] S. Saito, Z. Huang, R. Natsume, S. Morishima, A. Kanazawa, and H. Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. *arXiv preprint arXiv:1905.05172*, 2019. 3, 8
- [57] M. Sela, E. Richardson, and R. Kimmel. Unrestricted facial geometry reconstruction using image-to-image translation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1576–1585, 2017. 3
- [58] A. Shysheya, E. Zakharov, K.-A. Aliev, R. Bashirov, E. Burkov, K. Isakov, A. Ivakhnenko, Y. Malkov, I. Pasechnik, D. Ulyanov, et al. Textured neural avatars. *arXiv preprint arXiv:1905.08776*, 2019. 2
- [59] C. Si, W. Wang, L. Wang, and T. Tan. Multistage adversarial losses for pose-based human image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 118–126, 2018. 2
- [60] A. Siarohin, E. Sangineto, S. Lathuilière, and N. Sebe. Deformable gans for pose-based human image generation. In *CVPR 2018-Computer Vision and Pattern Recognition*, 2018. 2, 8
- [61] J. Starck and A. Hilton. Model-based multiple view reconstruction of people. In *IEEE international conference on computer vision*, pages 915–922, 2003. 3
- [62] A. Tewari, F. Bernard, P. Garrido, G. Bharaj, M. Elgharib, H.-P. Seidel, P. Pérez, M. Zollhöfer, and C. Theobalt. Fml: Face model learning from videos. *arXiv preprint arXiv:1812.07603*, 2018. 3
- [63] A. Tewari, M. Zollhöfer, P. Garrido, F. Bernard, H. Kim, P. Pérez, and C. Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 hz. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2549–2559, 2018. 3
- [64] G. Varol, J. Romero, X. Martin, N. Mahmood, M. J. Black, I. Laptev, and C. Schmid. Learning from synthetic humans. In *CVPR*, 2017. 5
- [65] M. Waechter, N. Moehrl, and M. Goesele. Let there be color! large-scale texturing of 3d reconstructions. In *Euro-*

- pean Conf. on Computer Vision, pages 836–850. Springer, 2014. 3
- [66] B. Wang, H. Zheng, X. Liang, Y. Chen, L. Lin, and M. Yang. Toward characteristic-preserving image-based virtual try-on network. In *ECCV*, 2018. 2, 8
- [67] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8798–8807, 2018. 4
- [68] T. Y. Wang, D. Cylan, J. Popovic, and N. J. Mitra. Learning a shared shape space for multimodal garment design. *arXiv preprint arXiv:1806.11335*, 2018. 3
- [69] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 4
- [70] Z. Wang, E. P. Simoncelli, and A. C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003. 4
- [71] Q. Xu, W. Wang, D. Ceylan, R. Mech, and U. Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *arXiv preprint arXiv:1905.10711*, 2019. 8
- [72] S. Yamaguchi, S. Saito, K. Nagano, Y. Zhao, W. Chen, K. Olszewski, S. Morishima, and H. Li. High-fidelity facial reflectance and geometry inference from an unconstrained image. *ACM Transactions on Graphics (TOG)*, 37(4):162, 2018. 3
- [73] C. Yang, Z. Wang, X. Zhu, C. Huang, J. Shi, and D. Lin. Pose guided human video generation. In *ECCV*, 2018. 2
- [74] J. Yang, J.-S. Franco, F. Hétroy-Wheeler, and S. Wuhrer. Analyzing clothing layer deformation statistics of 3d human motions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 237–253, 2018. 3
- [75] M. Zanfir, A.-I. Popa, A. Zanfir, and C. Sminchisescu. Human appearance transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3
- [76] Q.-Y. Zhou and V. Koltun. Color map optimization for 3d reconstruction with consumer depth cameras. *ACM Transactions on Graphics*, 33(4):155, 2014. 3
- [77] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 4
- [78] S. Zhu, S. Fidler, R. Urtasun, D. Lin, and C. Change Loy. Be your own prada: Fashion synthesis with structural coherence. In *International Conference on Computer Vision (ICCV)*, 2017. 3