

StylePeople: A Generative Model of Fullbody Human Avatars

Artur Grigorev^{1,2*}, Karim Iskakov^{1*}, Anastasia Ianina¹, Renat Bashirov¹,
Ilya Zakharkin^{1,2}, Alexander Vakhitov¹, Victor Lempitsky^{1,2}

¹ Samsung AI Center, Moscow

² Skolkovo Institute of Science and Technology, Moscow

Abstract

We propose a new type of full-body human avatars, which combines parametric mesh-based body model with a neural texture. We show that with the help of neural textures, such avatars can successfully model clothing and hair, which usually poses a problem for mesh-based approaches. We also show how these avatars can be created from multiple frames of a video using backpropagation. We then propose a generative model for such avatars that can be trained from datasets of images and videos of people. The generative model allows us to sample random avatars as well as to create dressed avatars of people from one or few images. The code for the project is available at saic-violet.github.io/style-people.

1. Introduction

Creating realistically-looking and articulated human avatars is a challenging task with many applications in telepresence, gaming, augmented and virtual reality. In recent years, sophisticated and powerful models of “naked” people that model shape of the body including facial and hand deformations have been developed [33, 50]. These models are based on mesh geometry and are learned from several existing datasets of body scans. Clothing and hair literally add an extra layer of complexity on top of body modeling, they are even more challenging for appearance modeling, and accurate 3D data for these elements are scarce and hard to obtain. And yet creating realistic avatars is not possible without modeling of these elements.

Here, we propose a new approach that we call *neural dressing* that allows to create 3D realistic full body avatars from videos and in a few-shot mode (from one or several images). Similarly to previous works, this approach uses deformable meshes (specifically, SMPL-X model [33]) to model and animate body geometry in 3D. On top of the body mesh, the approach superimposes a multi-channel neural texture [44] that is processed by a rendering network

in order to generate images of a full-body avatar with clothing and hair. Our first contribution is thus to show that the combination of deformable mesh models and neural textures (*neural dressing*) can model appearance of full-body avatars with loose clothing and hair well and to account for the geometry missing in parametric body models.

As our second contribution, we build a generative model of full-body avatars. The key component of the new model is a generative network for neural body texture. The generative network is derived from StyleGANv2 [21] generator. To build the complete model, we thus incorporate neural texture synthesis, mesh rendering, and neural rendering into the joint generation process, which is trained in an adversarial fashion on a large-scale dataset of full-body images. We also address the need to ensure that avatars have consistent appearance across variety of poses and camera positions. This is ensured by adding an additional discriminator network and by modifying the training process accordingly.

Using the resulting generative model, we can sample new realistic 3D “artificial humans” (*StylePeople* – Figure 1). Furthermore, the availability of a generative model allows us to build avatars for existing people by fitting the model to a single image or few images of a given person. As the generative model is heavily over-parameterized, we investigate how the fitting process can be regularized. In the experiments, we then compare our model to previous approaches to few-shot avatar modeling.

2. Related work

An early generative model for people in clothing [24] translates renderings of a parametric body model to human parsing (semantic segmentation) maps, which are subsequently translated into human images. The pose-guided human image generation [29] methods synthesize human images in new poses or with new clothes from a single image by either relying on feature warping in pose-to-image translation [41, 11, 10, 14] or on predicting surface coordinates in the source and target frames [5] and sampling RGB texture maps [31, 13]. In [30] the textures are transferred from clothing images to 3D garments worn on top of SMPL [28]

*equal contribution



Figure 1. **Style people**, i.e. random samples from our generative models of human avatars (truncation factor 0.8) . Each avatar is shown from two different viewpoints. The samples show diversity in terms of clothing and demographics. Loose clothing and hair are present.

using only shape information and ignoring texture.

In a video-to-video setting, [48] trains a system that transforms a sequence of body poses into corresponding temporally-consistent monocular videos. Likewise, [39] proposes to translate renderings of body joints to clothed body surface coordinates and use the latter to sample the learnable RGB texture stack achieving better visual quality. The geometry-texture decomposition helps occlusion reasoning and generalization to unseen poses, however surface coordinate regression mistakes result in visual artifacts, especially for the face and hands. Our approach goes one step further by switching from body coordinate regression to fitting a deformable model to images, improving alignment and geometric consistency.

The works [3, 2, 4, 7, 25, 16] take a single or multiple images of a person and produce a texture and displacement map for a body model, which later can be rendered from any viewpoint and in arbitrary pose. However the quality of rendering the human avatar in this case is limited by the classic graphical pipeline, while in our case a rendering network allows to increase photorealism.

Recently, approaches that rely on implicit functions to create 3D model of human bodies [36, 37] emerged. They

allow to generate detailed shapes of clothed human body. Although, they are restricted to manipulate those shapes only in terms of viewpoint since they lack structural information and can not render fitted humans in new poses.

We build our approach on a powerful *deferred neural rendering* model of [44] proposing learnable neural textures for meshes, which encode visual appearance and can be rendered using a standard graphics pipeline and then translated into photorealistic images of objects and scenes. In our system, the neural textures also implicitly encode non-modelled geometry, e.g. clothing and hair. While the system of Thies et al. needs to be trained from scratch requiring at least several dozens of images per scene, we propose a generative model of neural textures for people rendering, trained on a new large dataset of annotated in-the-wild images, which allows few-shot synthesis of a photorealistic human avatar. The most recent approaches [47, 52, 26, 40] report pose-to-image video-to-video translation in the few-shot regime, when the target appearance is specified by several images. Our approach leads to superior visual quality in the same setting by relying on a better geometric proxy as well as on the geometry-texture decomposition.

Finally, our StylePeople model builds on top of the re-




Figure 2. The neural dressing model. The texture of an avatar (left column; first three PCA components are visualized) is superimposed on the body mesh modeled with SMPL-X [33] (columns 2,3). Processing the rasterizations with the rendering network generates the renders (columns 4,5). Note that the model can handle loose clothing and long hair as well as complex garment textures gracefully.

cent years of advances in modeling high-resolution 2D images ([19, 8, 20]), in particular StyleGANv2 model [21]. Tewari et al. [43] propose to use StyleGAN model for rig-like control over face features. However, they can not control parts of the scene that are not explained by the parametric face model (e.g. hair style). The pairwise discriminator that ensures consistency between poses of the same avatar in our work has been recently proposed in a similar extension for StyleGAN model in [27] for the task of video synthesis (while in general, the use of multi-frame discriminators for consistent video synthesis can be traced to [45]). The inference ideas we apply to invert our generative model in few-shot learning mode are based on extended generator space [20, 1], learning the auxiliary encoders [21, 27], using discriminator feature matching loss [32], as well as fine-tuning generators to achieve the best fit [53, 6, 32].

3. Methods

We first describe the neural dressing model that combines together deformable mesh modeling with neural rendering. After that we focus on the generative model for human avatars that builds on top of the neural dressing model.

3.1. Neural dressing model

The neural dressing model (Figure 2) builds on top of the deformable shape model (SMPL-X in our case) that generates fixed-topology mesh $M(\mathbf{p}, \mathbf{s})$ driven by sets of pose parameters \mathbf{p} and body shape parameters \mathbf{s} . We assume that the mesh comes with a pre-defined texture mapping function, and denote with $\mathcal{R}(M, T; C)$ the *rasterization* function that takes mesh M , L -channel texture T and camera parameters C , to generate L -channel rasterization of the textured mesh using z-buffer algorithm.

We use L -channel texture T with $L=16$ in our experiments to encode local photometric and geometric information (including the geometry missing in parametric mesh). We then use a rendering image-to-image network f_θ with learnable parameters θ to translate the L -channeled rasterized image R to a six-channeled image I of the same size, where the first three channels of I correspond to RGB color channels and the last three channels correspond to segmentation masks for the whole foreground, head and hands. Additional segmentation channels for head and hands are used as additional supervision to improve visual quality of these body parts.

In the neural dressing model, an avatar A is characterized by body shape parameters \mathbf{s}_A and neural texture T_A . Given a pretrained rendering network f_θ , it can be rendered for an arbitrary pose \mathbf{p} and arbitrary camera parameters C . At test time, the rendering process runs at interactive speed (~ 25 FPS at one megapixel resolution).

3.2. Learning neural dressing via backpropagation

Given a collection of videos of several people, we can create their avatars by fitting our model via backpropagation. For person i , we assume that a set of video frames I_i^j is given ($j \in 1..N_i$). We then obtain three-channel segmentation map using a pretrained network [12], and fit parameters of body shape \mathbf{s}_i , pose \mathbf{p}_i^j and camera C_i^j corresponding to each individual frame. Our fitting is based on a modification of the SMPLify-X algorithm [33] that constrains the body shape parameters to be shared between frames.

We then jointly optimize parameters of the rendering network θ and neural textures T_i of all individuals using backpropagation. The optimization process minimizes following loss values: perceptual [18], adversarial [17], and feature matching [46] losses between ground truth images

I_i^j and rendered images $f_\theta[\mathcal{R}(M(\mathbf{p}_i^j, \mathbf{s}_i), T_i, C_i^j)]$. We use above-mentioned losses for color channels and Dice loss for ground truth and predicted segmentation masks. Importantly, when predicting the mask, we assign all pixels covered by the parametric body model to foreground.

In the first set of experiments, we pretrain the rendering network parameters θ on a dataset of 56 individuals, for which we have collected many thousands of frames. We can then create an avatar for a new person in three steps: (i) estimate pose, body shape, and camera parameters for a limited number of frames of the individual, (ii) generate three-channel segmentation masks for those frames and then (iii) use the fixed pretrained rendering network to optimize neural texture for the person by backpropagation. Note that in our current implementation, gradients are backpropagated through rasterization function only to texture parameters but not mesh or camera parameters, which means that non-differentiable renderers are suitable for our learning task.

3.3. Generative modeling

Learning with backpropagation allows to create avatars from relatively short videos or even few appropriately distributed frames (with each surface part of the person covered in at least one view). In some practical scenarios it may be beneficial to create fullbody avatars from impartial information, e.g. a single view in A-pose. To do that, system needs ability to reason about unobserved parts of the avatar. We take the generative modeling approach to handle this task.

We build a generator for neural textures that can be used in neural dressing. We start with the recent StyleGANv2 approach [21] for high-resolution image modeling. Following [21], we train a multi-layer perceptron $g_\phi(z) : \mathbb{R}^{512} \rightarrow \mathbb{R}^{512}$ and a convolutional network $h_\psi(\mathbf{w}^4, \dots, \mathbf{w}^{512}, \mathbf{N})$ that takes as an input a set of 512-dimensional *style vectors* controlling generation at different resolutions (from 4×4 to 512×512) via modulation-demodulation mechanism as well as the set \mathbf{N} of noise maps at these resolutions. Together, the two networks define the generative model $h_\psi(g_\phi(z), \mathbf{N})$, in which a random vector \mathbf{z} is passed through the MLP and then used as an input at multiple resolutions into the convolutional part.

Our architecture for the generator (Figure 3) closely follows [21] with two exceptions. First, in our case it outputs L -channeled neural texture instead of an RGB image. Second, we concatenate inputs of several later layers of the network with the 16-channel map of spectral coordinates of mesh vertices mapped into texture space to provide the generator with some information about the body mesh geometry. During training, we also sample body poses and shapes from the training dataset, and superimpose the generated textures over the deformed body meshes. The resulting textured meshes are then passed through the rendering network

that is trained as a part of the generator. Overall, the images produced by our generator are computed as follows:

$$\mathcal{G}_{\phi, \psi, \theta}(\mathbf{z}, \mathbf{N}, \mathbf{p}, \mathbf{s}, C) = \quad (1)$$

$$f_\theta \left[\mathcal{R} \left(M(\mathbf{p}, \mathbf{s}), h_\psi(g_\phi(\mathbf{z}), \mathbf{N}), C \right) \right], \quad (2)$$

where \mathbf{z} and \mathbf{N} variables are sampled from the unit Normal distributions, while the pose, the camera, and the body shape variables are sampled from the empirical distribution in the training set.

Training of the full model is performed adversarially, with a dataset of videos of people as a source for real images. For each video we obtain foreground segmentation as well as body shape, pose and camera parameters as discussed above. We utilize three different discriminator networks to train our system:

- **Unary discriminator** closely follows StyleGANv2 architecture, and simply considers individual images (generated by our model as well as real images). The unary discriminator thus focuses on image quality.
- **Binary discriminator** also follows StyleGANv2 architecture, however it takes pairs of images rather than individual images. Each real example is made up of two frames of a single video along with their segmentations. To obtain a fake example, we generate two examples $\mathcal{G}_{\phi, \psi, \theta}(\mathbf{z}, \mathbf{N}, \mathbf{p}_1, \mathbf{s}, C_1)$ and $\mathcal{G}_{\phi, \psi, \theta}(\mathbf{z}, \mathbf{N}, \mathbf{p}_2, \mathbf{s}, C_2)$, where \mathbf{z} , \mathbf{N} (and thus the avatar texture), as well as the body shape \mathbf{s} are shared between examples. Pose parameters $\mathbf{p}_1, \mathbf{p}_2$ and camera positions C_1, C_2 are different within each pair and correspond to two frames of the same video. The role of this discriminator is thus both to assess the visual realism of examples, as well as to ensure that the identity is preserved through camera change and pose change.
- **Face discriminator** considers individual images cropped around the face region (for both real and synthetic images) and is used to improve quality of models' face rendering.

Additional regularization. During training, we take additional steps to ensure identity preservation, as merely relying on binary discriminator turns out to be insufficient. We use three additional tricks. First, we train a predictor q_ξ that takes generated images $I = \mathcal{G}_{\phi, \psi, \theta}(\mathbf{z}, \mathbf{N}, \mathbf{p}, \mathbf{s}, C)$ and tries to recover vector $\mathbf{w} = g_\phi(\mathbf{z})$. The training loss $\|q_\xi(I) - \mathbf{w}\|^2$ is backpropagated through the entire generator, and thus ensures consistency of images created with same avatar texture.

The second regularizing trick is to ensure that rendering network is covariant to in-plane geometric transfor-

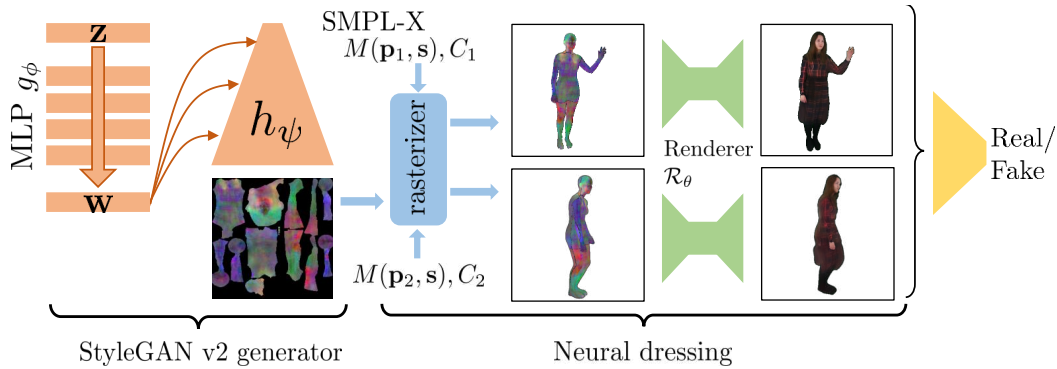


Figure 3. Our generative architecture is based on the combination of StyleGANv2 and Neural dressing. The StyleGAN part is used to generate neural textures, which are then superimposed on SMPL-X meshes and rendered with a neural renderer. During adversarial learning, the discriminator considers a pair of images of the same person. See text for more details.

mation of its input. We achieve this by applying a random in-plane rigid transformation Tr on rasterized images $J = \mathcal{R}(M(\mathbf{p}, \mathbf{s}), h_\psi(\mathbf{z}, \mathbf{N}), C)$. Transformed rasterizations $\text{Tr}(J)$ are then passed through the rendering network r_θ , penalizing difference between $\text{Tr}(f_\theta[J])$ and $f_\theta[\text{Tr}(J)]$.

Lastly, to mitigate the influence of poorly segmented images from our in-the-wild training dataset and those, where a part of a person’s body is occluded, we force foreground masks of generated samples to cover the whole rendered mesh given as an input. To do that, instead of a sampled segmentation mask we use its union with the binary mask of the mesh as a final foreground segmentation.

3.4. Encoders and inference

We now discuss how the generative model is used to fit a new avatar to one (or few) images of a person. Given an image of the person I (with corresponding segmentation mask), we estimate the parameters of its’ body shape \mathbf{s} , pose \mathbf{p} , and camera C . Our goal is to find a texture T such that the rendering $f_\theta[\mathcal{R}(M(\mathbf{p}, \mathbf{s}), T, C)]$ matches the observed image. The texture is parameterized by the convolutional generator $T = h_\psi(\mathbf{w}^4, \dots, \mathbf{w}^{512}, \mathbf{N})$, and thus depends on the style vectors $\mathbf{w}^4, \dots, \mathbf{w}^{512}$, the noise tensors \mathbf{N} and the parameters ψ of the generator h_ψ .

Encoders. As our texture parametrization is excessive (i.e. the number of texture elements is usually greater than the number of observations in I), very different sets of latent variables are able to fit the observed image, while leading to different degree of generalization to new poses and camera parameters. We have therefore found it important to learn *encoder networks* that would initialize latent vectors \mathbf{w} at the part of the latent space that leads to good generalization.

To learn the encoder network, we generate a dataset of synthetic samples from our texture generator. In particular, to get the k -th sample, we randomly and independently draw $\mathbf{w}_k^4, \dots, \mathbf{w}_k^{512}$ by sampling \mathbf{z} values and passing

them through the perceptron part of the generative model. The noise tensors are drawn from the Normal distribution, and the convolutional generator produces a random texture. We then superimpose the texture on a random body from our training set in an approximate A-pose and pick an approximately frontal camera. The rendering network then produces an image I_k of a random A-posed avatar. Our *A-Encoder* \mathcal{E}_A is then trained to recover the vectors $\mathbf{w}_k^4, \dots, \mathbf{w}_k^{512}$ from image I_k . The training uses the L1-loss on synthetic data from our generative model only.

In addition to the A-encoder that is trained only on synthetic data, and is suitable for A-posed images, we have trained a generic encoder (*G-encoder*) on both synthetic data and pairs of video frames from real dataset. To train on real data the generic encoder \mathcal{E}_G takes a pair of real images (I_k, J_k) extracted from the same video of the same person, predicts latent variables $\mathbf{w}_k^4, \dots, \mathbf{w}_k^{512}$ from I_k , augments them with random noise tensors \mathbf{N} resulting in the texture $T_k = h_\psi(\mathcal{E}_G(I_k), \mathbf{N})$. The resulting texture is then superimposed on the mesh, posed and rendered according to the SMPL-X parameters $\mathbf{p}_k, \mathbf{s}_k$ and camera pose C_k observed in the image J_k . Learned perceptual similarity (LPIPS) [54] between J_k and $f_\theta[\mathcal{R}(M(\mathbf{p}_k, \mathbf{s}_k), T_k, C_k)]$ is then used as a loss function for real data. For synthetic data, similarly to \mathcal{E}_A , we use the L1-loss between synthetic and predicted latent vectors. These two loss values are taken with equal weights resulting in the loss for the generic encoder \mathcal{E}_G .

The encoder’s architecture is inspired by the pSp-architecture proposed in [34]. The main idea is that each level of the encoder predicts a latent vector corresponding to the generator’s resolution of this level. As a backbone we use EfficientNet-B7 [42] pretrained on ImageNet [35]. More details about architecture is provided in Supplementary Material.

Fitting. Avatar fitting to the given image(s) is done by optimization. The input to the optimization process is a set of images of a person $\mathbf{I} = \{I_i | i \in [1, N]\}$. Latent vectors \mathbf{w} are initialized by passing the images \mathbf{I} through one of the two pretrained encoders. If more than one input image are available ($N > 1$), the predicted latent vectors are averaged across all images. For the sake of simplicity we further assume that only one image I is given as input. The optimization is performed over (i) the latent vectors \mathbf{w} , (ii) the parameters ψ of generator h_ψ and (iii) the noise tensors \mathbf{N} to further minimize the difference between the recovered image $f_\theta[\mathcal{R}(M(\mathbf{p}, \mathbf{s}), T, C)]$ and I . As a last stage of the optimization, we perform direct optimization of the texture values for 100 iterations.

During the optimization we use multiple loss functions, namely LPIPS-loss [54] between recovered image and input image, Mean Squared Error (MSE) between recovered image and input image, Mean Absolute Error (MAE) on latent variables w deviation from the initialization predicted by the encoder, MAE on generator parameters h_ψ deviation from the initial ones, MAE on texture deviation from the texture values, optimized to the beginning of the last stage, LPIPS-loss [54] on face regions of recovered and input images, and feature matching loss based on the trained face discriminator [32]. In case of multiple input images, losses are averaged across all images. We provide additional analysis on the role of losses in the Supplementary Material. Optimization is performed via backpropagation with ADAM optimizer [22].

4. Experiments

Below, we present experimental validation of the neural dressing approach. We first evaluate the approach in the scenario, where an avatar is created from a short video of a person, using backpropagation without generative prior on the neural texture. We then proceed to evaluate samples from our generative model and its performance for one-shot and few-shot image-based full-body avatar creation. We provide additional results including qualitative comparisons in the Supplementary material.

Datasets. We use our recently collected *AzurePeople* dataset to pretrain rendering network in backpropagation-based approach. It was collected with five Kinect Azure sensors and contains multi-view RGBD-videos of 56 people (though we do not use depth information in this work).

Four other datasets were used to train generative model. The first is a multi-view *HUMBI* dataset [51]. At the moment of experiments it has 442 subjects available with total of 61,975 images. This dataset alone was used to train binary discriminator and ensure multi-view consistency of generated samples. The second dataset *TEDXPeople* was collected by us and contains 41,233 videos of TED and

TED-X talks. For our purposes we have extracted and processed 8 to 16 frames from each video, in which a person is visible in either full-body or close-up upper-body view. This dataset is very diverse in terms of demographics and clothing styles, though the resolution is limited and there are certain biases associated with peculiar lighting and camera viewpoints specific to TED(-X) talks. Because of these limitations, we only used this dataset for the unary discriminator.

Additionally, we used the dataset containing *curated SMPL-X fits* released by the authors of ExPose [9] and the *Google OpenImages v6* dataset [23] (from which we extracted human bounding boxes). After filtering out low-resolution and highly occluded samples from these two datasets, we added remaining 10869 *ExPose* samples and 37660 *OpenImages* samples to our training procedure.

For evaluation only, we also use two videos from the PeopleSnapshot dataset [3] ('female-1-casual' and 'male-2-casual'). Each video has a person rotating in front of the camera in A-pose. We choose these particular sequences because we have access to their avatars created with several previous methods.

4.1. Video-based avatars

In the first set of experiments, we assess the ability of our models to create avatars from videos, in which one person is seen in different poses. Here, we create avatars using backpropagation as described in Section 3.2. We show qualitative results of this approach on people from *AzurePeople* dataset in Figure 4 and in supplementary video.

For evaluation, we also compare our results to textured neural avatar system [39], and mesh-based approach [3]. Since we do not have any hold-out video sequences for the models from [3], we perform a user study using an online crowd-sourcing platform showing short videos of two people from PeopleSnapshot dataset in pose sequences sampled from *TEDXPeople* dataset. For visual quality, participants preferred our method in 69.8% of the cases when compared to [39] and in 70.6% of the cases when compared to [3]. For identity preservation, preferences were 67.7% and 69.8% respectively.

We have additionally compared our method to [39] on in-the-wild YouTube video considered in [39]. It contains much more challenging poses than *TEDXPeople* samples. Here participants preferred our method 63.9% of the cases when asked about visual realism and in 52.6% of the cases when asked about identity preservation.

In both studies participants were presented with two-second segments of 15 fps videos (10 segments for PeopleSnapshot sequences and 18 for the YouTube sequence) and asked to choose from pairs of side-by-side positioned samples. In identity preservation study, an image from ground truth video was presented along with the compared



Figure 4. Results for video-based avatars. Sample source frames are shown on the left. The remaining images show the video-based avatars in previously unseen poses shown within a simple augmented reality application. The pose sequence is transferred from a different person.

results. The order of compared results for each pair was randomized. Each pair was assessed by 30 people.

4.2. Few-shot avatars

We then focus our evaluation on the few-shot scenario. We first evaluate our avatar against the mesh-based Octopus approach [2] in the eight-shot mode (i.e. each avatar is created from eight photographs). For one-shot mode, we compare against the mesh-based method [25], as well as three neural rendering/warping one-shot methods: the first-order motion model [40], the liquidGAN model [26], and the coordinate-based inpainting model [13]. The first-order motion model was retrained on the TEDXPeople dataset (for completeness we also report the results for the author’s model trained on the TaiChi dataset). For other methods, author-provided pretrained models/results were used (see the supplementary material for details). For our method, we consider two variants initialized with two different encoders. Figure 5 demonstrates process of fitting our models to the data (using the A-encoder).

We use a set of standard metrics to evaluate image quality of the avatars, including learned perceptual similarity (LPIPS) [54], structured self-similarity (SSIM) [49], Fréchet Inception Distance (FID) [15], and Inception Score (IS) [38]. We use two-person People Snapshot dataset to compare with [2], [25], [26] and [40]. (Table 1) as those were available to us for all methods. We also compare to the state-of-the-art one-shot methods on three hold-out

video sequences from test split of our TEDXPeople dataset (Table 2, Figure 5). Both comparisons show advantage of our approach in all metrics, except the Inception Score, which we subjectively find least correlated with visual quality. Qualitative results and additional comparisons are presented in the supplementary material and video.

It has to be mentioned, that while all one-shot models create avatars from the same data (a single image), our model requires a lengthy optimization process to come up with an avatar, which is a limitation of our model.

	IS \uparrow	FID \downarrow	LPIPS \downarrow	SSIM \uparrow
One-shot				
360Degree [25]	1.8643	1383.1	0.2123	0.8079
LWGAN [26]	1.7159	1771.9	0.2727	0.2876
FOMM-TaiChi [40]	1.7643	1059.4	0.1767	0.7585
FOMM-TEDX [40]	1.5404	674.3	0.1477	0.8743
Neural Dressing (A-encoder)	1.8025	349.2	0.1957	0.8264
Neural Dressing (G-encoder)	1.7469	272.1	0.0836	0.9012
8-shot				
Octopus [2]	1.8123	403.6	0.1379	0.8324
Neural Dressing (A-encoder)	1.7335	183.3	0.1169	0.8764
Neural Dressing (G-encoder)	1.6608	158.2	0.0759	0.9079

Table 1. Comparison against state-of-the-art few-shot methods on two sequences from People Snapshot dataset [3]. We use the first frame as a source image for one-shot methods, and the frames used in [2] for eight-shot methods, and evaluate on the remaining frames of the sequences. Metrics are calculated with respect to the ground truth sequence.



Figure 5. One-shot inference results for test subjects from People Snapshot dataset [3]. Single *Input image* is fed into the encoder to produce initial texture estimation (*Encoder prediction*). Next, optimization step draws output renderings closer to the original person (*After optimization*). Note, how non-visible details in the source image propagate to the output from the initial encoder prediction (e.g. the back of the male’s shirt).

	IS \uparrow	FID \downarrow	LPIPS \downarrow	SSIM \uparrow
One-shot				
CBI [13]	1.4738	6761.3	0.1814	0.8293
LWGAN [26]	1.7467	10671.9	0.3284	0.6068
FOMM-TaiChi [41]	1.2981	3253.1	0.2186	0.7725
FOMM-TEDX [41]	1.2663	3976.1	0.2312	0.7944
Neural Dressing (A-encoder)	1.4448	601.7	0.0531	0.9369
Neural Dressing (G-encoder)	1.5903	765.7	0.0542	0.9361
8-shot				
Neural Dressing (A-encoder)	1.4372	505.2	0.0521	0.9389
Neural Dressing (G-encoder)	1.5548	559.8	0.0537	0.9381

Table 2. Comparison against state-of-the-art one-shot methods on three test sequences of the TEDXPeople dataset. Our method performs better than competitors in most metrics.

4.3. Generative model ablation

Additionally, we evaluate contributions of regularization mechanisms and additional discriminators in the training process of generative model. We compare ablations in terms of two metrics. First, we use Fréchet Inception Distance (FID) between the set of real validation images and the set of sampled images of people with the same meshes as in the validation set. The second metric is the Euclidean distance between vectors extracted with a re-identification network [55] from two images rendered with same neural texture, but positioned with pose difference of 180° (denoted as ReID). While the first metric focuses on the image quality, the second measures consistency of rendering.

With each subsequent ablation we turn off one of the modifications of the training procedure. Namely, *no-augmentations* loses in-plane rotation regularization *no-meshmask* uses plain sampled segmentation as foreground mask instead of its’ union with the mask of SMPL-X mesh

	FID \downarrow	ReID \downarrow
<i>Ours-full</i>	48.3	15.0
<i>no-augmentations</i>	48.8	16.3
<i>no-meshmask</i>	54.2	15.8
<i>no-speccoords</i>	47.8	15.5
<i>no-latent-predictor</i>	44.2	15.9
<i>no-face-discr</i>	43.4	16.0
<i>no-bin-discr</i>	42.9	17.3

Table 3. Ablation study for generative model of neural textures

projection, *no-speccoords* omits spectral coordinates passed into the texture generator, *no-latent-predictor* does not use latent predictor q_ξ , *no-face-discr* and *no-bin-discr* omit facial and binary discriminator respectively.

As seen from Table 3, although the most basic variant provides lowest FID for single images, the modifications improve consistency across views.

5. Conclusion

We have presented a new approach for modeling the human appearance. This approach integrates polygonal body mesh modeling with neural rendering. Thus, polygonal mesh controls and models coarse body geometry as well as the pose, while neural rendering component adds clothing and hair. On top of that, we train a generative model for neural textures, and show that this generative model allows to obtain avatars in few-shot mode. We show that proposed approach improves state-of-the-art for video-based and few image-based avatar acquisition. Admittedly, results of our method are still limited by the amount and quality of data, especially diversity the biases of video sequences. Improving data efficiency of our generative modeling is thus one of the important directions of our future work.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proc. ICCV*, pages 4432–4441, 2019.
- [2] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single rgb camera. In *Proc. CVPR*, pages 1175–1186, 2019.
- [3] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *Proc. CVPR*, pages 8387–8397, 2018.
- [4] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2293–2303, 2019.
- [5] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proc. CVPR*, pages 7297–7306, 2018.
- [6] David Bau, Hendrik Strobelt, William S. Peebles, Jonas Wulff, Bolei Zhou, Jun-Yan Zhu, and Antonio Torralba. Semantic photo manipulation with a generative image prior. *ACM Trans. Graph.*, 38(4):59:1–59:11, 2019.
- [7] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *Proc. ICCV*, pages 5420–5430, 2019.
- [8] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2018.
- [9] Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J. Black. Monocular expressive body regression through body-driven attention. In *Proc. ECCV*, 2020.
- [10] Haoye Dong, Xiaodan Liang, Ke Gong, Hanjiang Lai, Jia Zhu, and Jian Yin. Soft-gated warping-gan for pose-guided person image synthesis. In *Proc. NeurIPS*, pages 474–484, 2018.
- [11] Patrick Esser, Ekaterina Sutter, and Björn Ommer. A variational u-net for conditional appearance and shape generation. In *Proc. CVPR*, pages 8857–8866, 2018.
- [12] Ke Gong, Yiming Gao, Xiaodan Liang, Xiaohui Shen, Meng Wang, and Liang Lin. Graphonomy: Universal human parsing via graph transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7450–7459, 2019.
- [13] Artur Grigorev, Artem Sevastopolsky, Alexander Vakhitov, and Victor Lempitsky. Coordinate-based texture inpainting for pose-guided image generation. In *Proc. CVPR*, 2019.
- [14] Xintong Han, Xiaojun Hu, Weilin Huang, and Matthew R Scott. Clothflow: A flow-based model for clothed person generation. In *Proc. ICCV*, pages 10471–10480, 2019.
- [15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6626–6637. 2017.
- [16] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. *arXiv preprint arXiv:2004.04572*, 2020.
- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *Proc. CVPR*, 2017.
- [18] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Proc. ECCV*, 2016.
- [19] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018.
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proc. CVPR*, pages 4401–4410, 2019.
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. *CoRR*, abs/1912.04958, 2019.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.
- [24] Christoph Lassner, Gerard Pons-Moll, and Peter V Gehler. A generative model of people in clothing. In *Proc. ICCV*, pages 853–862, 2017.
- [25] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *Proc. 3DV*, pages 643–653. IEEE, 2019.
- [26] Wen Liu, Zhixin Piao, Jie Min, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *Proc. ICCV*, pages 5904–5913, 2019.
- [27] Elizaveta Logacheva, Roman Suvorov, Oleg Khomenko, Anton Mashikhin, and Victor Lempitsky. Deeplandscape: Adversarial modeling of landscape videos. In *Proc. ECCV*, pages 256–272. Springer, 2020.
- [28] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015.
- [29] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Proc. NeurIPS*, pages 406–416, 2017.

- [30] Aymen Mir, Thimeo Alldieck, and Gerard Pons-Moll. Learning to transfer texture from clothing images to 3d humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7023–7034, 2020.
- [31] Natalia Neverova, Riza Alp Guler, and Iasonas Kokkinos. Dense pose transfer. In *Proc. ECCV*, pages 123–138, 2018.
- [32] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. *Proc. ECCV*, 2020.
- [33] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [34] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. *arXiv preprint arXiv:2008.00951*, 2020.
- [35] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [36] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2304–2314, 2019.
- [37] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020.
- [38] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2234–2242. 2016.
- [39] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliev, Renat Bashirov, Egor Burkov, Karim Isakov, Aleksei Ivakhnenko, Yury Malkov, Igor Pasechnik, Dmitry Ulyanov, et al. Textured neural avatars. In *Proc. CVPR*, pages 2387–2397, 2019.
- [40] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *Proc. NeurIPS*, pages 7135–7145, 2019.
- [41] Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, and Nicu Sebe. Deformable gans for pose-based human image generation. In *Proc. CVPR*, pages 3408–3416, 2018.
- [42] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
- [43] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6142–6151, 2020.
- [44] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019.
- [45] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proc. CVPR*, pages 1526–1535, 2018.
- [46] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proc. CVPR*, pages 8798–8807. IEEE Computer Society, 2018.
- [47] Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Bryan Catanzaro, and Jan Kautz. Few-shot video-to-video synthesis. In *Proc. NeurIPS*, pages 5014–5025, 2019.
- [48] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Proc. NeurIPS*, pages 1144–1156, 2018.
- [49] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [50] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *Proc. CVPR*, 2019.
- [51] Zhixuan Yu, Jae Shin Yoon, In Kyu Lee, Prashanth Venkatesh, Jaesik Park, Jihun Yu, and Hyun Soo Park. Humbi: A large multiview dataset of human body expressions, 2020.
- [52] Polina Zablotskaia, Aliaksandr Siarohin, Bo Zhao, and Leonid Sigal. Dwnet: Dense warp-based network for pose-guided human video generation. In *Proc. BMVC*, 2019.
- [53] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *Proc. ICCV*, pages 9459–9468, 2019.
- [54] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. CVPR*, pages 586–595, 2018.
- [55] Kaiyang Zhou, Xiatian Zhu, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Learning generalisable omni-scale representations for person re-identification. *arXiv preprint arXiv:1910.06827*, 2019.