

3D-Aware Semantic-Guided Generative Model for Human Synthesis

Jichao Zhang¹, Enver Sanginetto², Hao Tang³, Aliaksandr Siarohin^{1,4},
Zhun Zhong¹, Nicu Sebe¹, and Wei Wang¹

¹University of Trento ²University of Modena and Reggio Emilia ³ETH Zurich ⁴Snap Research

Abstract. Generative Neural Radiance Field (GNeRF) models, which extract implicit 3D representations from 2D images, have recently been shown to produce realistic images representing rigid/semi-rigid objects, such as human faces or cars. However, they usually struggle to generate high-quality images representing non-rigid objects, such as the human body, which is of a great interest for many computer graphics applications. This paper proposes a 3D-aware Semantic-Guided Generative Model (3D-SGAN) for human image synthesis, which combines a GNeRF with a texture generator. The former learns an implicit 3D representation of the human body and outputs a set of 2D semantic segmentation masks. The latter transforms these semantic masks into a real image, adding a realistic texture to the human appearance. Without requiring additional 3D information, our model can learn 3D human representations with a photo-realistic, controllable generation. Our experiments on the DeepFashion dataset show that 3D-SGAN significantly outperforms the most recent baselines. The code is available at <https://github.com/zhangqianhui/3DSGAN>.

Keywords: Generative Neural Radiance Fields, Human image generation

1 Introduction

Recent deep generative models can generate and manipulate high-quality images. For instance, Generative Adversarial Networks (GANs) [14], have been applied to different tasks, such as image-to-image translation [87, 10, 25], portrait editing [2, 69, 77, 70], and semantic image synthesis [57], to mention a few. However, most state-of-the-art GAN models [19, 28, 32, 33, 29, 20, 30] are trained using 2D images only, operate in the 2D domain, and ignore the 3D nature of the world. Thus, they often struggle to disentangle the underlying 3D factors of the represented objects.

Recently, different 3D-aware generative models [50, 51, 80] have been proposed to solve this problem. Since most of these methods do not need 3D annotations, they can create 3D content while reducing the hardware costs of common computer graphics alternatives. Differently from generating 3D untextured shapes [80, 13], some of these methods [88, 9, 50, 39, 51] focus on 3D-aware realistic image generation and controllability. Generally speaking, these models mimic the traditional computer graphics rendering pipeline: they first model the 3D structure, then they use a (differentiable) projection module to project the 3D structure into 2D images. The latter may be a depth map [9], a sketch [88] or a feature map [50] which is finally mapped into the real image



Fig. 1: A qualitative comparison between different generation methods: GRAF [68], pi-GAN [7], GIRAFFE [53], ShadeGAN [56] CIPS-3D [85], and 3D-SGAN (Ours).

by a rendering module. During training, some methods require 3D data [88, 9], and some [50, 39, 51] can learn a 3D representation directly from raw images.

An important class of *implicit* 3D representations are the Neural Radiance Fields (NeRFs), which can generate high-quality unseen views of complex scenes [48, 26, 12, 59, 58, 60, 7]. Generative NeRFs (GNeRFs) combine NeRFs with GANs in order to condition the generation process with a latent code governing the object’s appearance or shape [68, 7, 53]. However, these methods [68, 7, 53, 56, 85] focus on relatively simple and “rigid” objects, such as cars and faces, and they usually struggle to generate highly non-rigid objects such as the human body (e.g., see Fig. 1). This is likely due to the fact that the human body appearance is highly variable because of both its articulated poses and the variability of the clothes texture, being these two factors entangled with each other. Thus, adversarially learning the data distribution modeling all these factors, is a hard task, especially when the training set is relatively small.

To mitigate this problem, we propose to *split* the human generation process in two separate steps and use intermediate segmentation masks as the bridge of these two stages. Specifically, our 3D-aware Semantic-Guided Generative model (3D-SGAN) is composed of two generators: a GNeRF model and a texture generator. The GNeRF model learns the 3D structure of the human body and generates a semantic segmentation of the main body components, which is largely invariant to the surface texture. The texture generator translates the previous segmentation output into a photo-realistic image. To control the texture style, a Variational AutoEncoder (VAE [35]) approach with a StyleGAN-like [33] decoder is used to modulate the final decoding process. The similar idea has been used in [38], but their semantic generator is 2D, and it cannot perform 3D manipulations. We empirically show that splitting the human generation process into these two stages brings the following three advantages. First, the GNeRF model is able to learn the intrinsic 3D geometry of the human body, even when trained with a small dataset. Second, the texture generator can successfully translate semantic information into a textured object. Third, both generators can be controlled by explicitly varying their respective conditioning latent codes. Moreover, we propose two consistency losses to further disentangle the latent codes representing the garment type (which we call the “semantic” code) and the human pose. Finally, since there is no general metric which can be used to evaluate the 3D consistency of image generation with

multiple viewpoints, we propose a point matching-based metric which we name *average Matched Points (aMP)*. Experiments conducted on the DeepFashion dataset [41] show that 3D-SGAN can generate high-quality person images significantly outperforming state-of-the-art approaches. In summary, the main contributions of this work are:

- 1) We propose 3D-SGAN, which combines a GNeRF with a VAE-conditioned texture generator for human-image synthesis.
- 2) We propose two consistency losses to increase the disentanglement between semantic information (e.g., garment type) and the human pose.
- 3) We show that 3D-SGAN generates high-quality human images, significantly outperforming the previous controllable state-of-the-art methods.
- 4) We propose a new metric (aMP) to evaluate the 3D-view consistency.

2 Related work

3D-aware image synthesis is based on generative models which incorporate a 3D scene representation. This allows rendering photo-realistic images from different viewpoints. Early methods use GAN-based architectures for building 3D voxel [80, 13, 43, 22] or mesh [62, 21] representations. However, they mostly focus on learning untextured 3D structures. More recently, different methods learn textured representations directly from 2D images [79, 88, 50, 68, 12, 52, 53]. The resulting controllable 3D scene representation can be used for image synthesis. Some of these methods [88, 9] require extra 3D data for disentangling shape from texture. The main idea is to generate an internal 3D shape and then project this shape into 2D sketches [88] or depth maps [9], which are finally rendered in a realistic image. Other methods are directly trained on 2D images without using 3D data [50, 68, 51, 53, 81, 74]. For instance, inspired by StyleGANv2 [33], Thu et al. [50] propose HoloGAN, which predicts 3D abstract features using 3D convolutions, and then projects these features into a 2D representation which is finally decoded into an image. However, the learnable projection function, e.g., the decoder, results in an entangled representation, thus the view-consistency of the generated images is degraded. Katja et al. [68] use a NeRF to represent the 3D scene and a volume rendering technique to render the final image. However, this model works at relatively low image resolutions and it is restricted to single-object scenes. To tackle these issues, some works propose object-aware scene representations. For example, Liao et al. [39] combine a 2D generator and a projection module with a 3D generator which outputs multiple abstract 3D primitives. Every stage in this model outputs multiple terms to separately represent each object. Instead of abstract 3D primitives, Phuoc et al. [51] use a voxel feature grid as the 3D representation, but their method fails to generate consistent images at high-resolution. Michael et al. [53] recently introduced GIRAFFE, a multiple-object scene representation based on NeRFs, jointly with an object composition operator. GIRAFFE is the state-of-the-art 3D-aware approach for both single and multiple object generation tasks. A few very recent papers [56, 75] propose to learn an accurate object geometry by introducing a relighting module into the rendering process. Xu et al. [82], explicitly learn a structural and a textural representation (a feature volume), which is used jointly with the implicit NeRF mechanism. Chan et al. [8], propose to replace the 3D volume with three projection feature planes. Finally, 3D-consistency

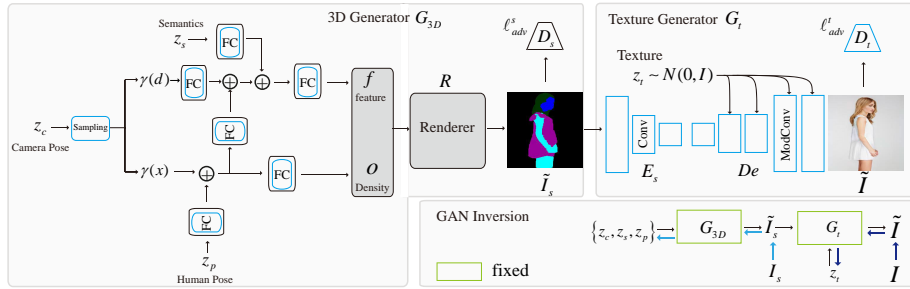


Fig. 2: An overview of the proposed 3D-SGAN architecture, composed of two main generators. G_{3D} (on the left) follows a GNeRF structure, with a NeRF kernel used to represent implicit 3D information, latent codes governing different appearance variations and a discriminator (D_s) which is used for adversarial training. The output of G_{3D} is the semantic masks \tilde{I}_s (middle). The second generator (G_t , right) translates the semantic masks into a photo-realistic image \tilde{I} . Also G_t is trained adversarially (see top right, the second discriminator D_t). The human generation process can be controlled by interpolating different latent codes: the semantics z_s , the pose z_p , the camera z_c , and the texture code z_t . The bottom of the figure shows the GAN inversion scheme.

is addressed in [17, 85, 54], where, e.g., StyleGAN-based networks are used for neural rendering. However, most of these works fail to disentangle the semantic attributes.

While previous methods can achieve impressive rigid-object generation and manipulation results, they usually struggle to deal with non-rigid objects with complex pose and texture variations. For instance, the human body is a non-rigid object which is very important in many generative applications.

GANs for human generation. GANs [14] have been widely used for different object categories, such as, for instance, faces [31, 28, 32, 6], cars [31, 32, 68, 51], and churches [53]. However, GANs still struggle to produce high-quality full-human body images, because of the complex pose variations. Very recently, Sarkar et al. [66] proposed a VAE-GAN model for the pose transfer and the part sampling tasks. In more detail, this model extracts an UV texture map from the input image using DensePose [3], and then encodes the texture into a Gaussian distribution. Then, it samples from this distribution and warps the sample into the target pose space. Finally, the warped latent code is used as input to the decoder. Compared to Sarkar et al. [66], our method does not use an SMPL nor DensePose to extract the point correspondences as additional supervised information. Despite that, 3D-SGAN can learn 3D representations of the human body and control the generation process (e.g., by changing the input camera parameters).

StylePeople [16] is based on a full-body human avatar, which combines StyleGANv2 [33] with *neural dressing*. The StyleGANv2 module samples neural textures, and these textures are superimposed on the meshes of an SMPL. The textured meshes are finally rendered into an image. In contrast, our 3D-SGAN can perform semantic disentanglement and manipulation using semantic codes.

Pose transfer aims to synthesize person images in a novel view or in a new pose. This is a very challenging task, since it requires very complicated spatial transformations to

account for different poses. Most works in this field can be categorized by the way in which the human pose is represented. Early works are based on keypoints [45, 71, 89, 61, 76, 46, 47, 24, 5, 44, 27, 72, 86, 63, 73, 83, 64]. More recent methods [15, 49, 67, 65, 40] use correspondences between pixel location in 2D images and points in SMPL [42] (usually estimated using DensePose [18]). However, these approaches usually struggle to simultaneously provide a realistic and a 3D controllable person generation.

3 Preliminaries

NeRF [48] is an implicit model which represents a 3D scene using the weights of a multilayer perceptron (MLP). This MLP (h) takes as input a 3D coordinate $\mathbf{x} \in \mathbb{R}^3$ and a view direction $\mathbf{d} \in \mathbb{R}^2$, and outputs the density (or “opacity”, o) and the view-dependent RGB color value \mathbf{c} :

$$(\mathbf{c}, o) = h(\gamma(\mathbf{x}), \gamma(\mathbf{d})), \quad (1)$$

where γ is a positional encoding function [78]. On the other hand, Generative NeRF (GNeRF) [68] is a NeRF conditioned on the latent codes \mathbf{z}_g and \mathbf{z}_a , respectively representing the geometric shape and the object appearance, and drawn from a priori distributions. GNeRFs [68, 53] are trained using an adversarial approach. In GIRAFFE [53], the color value (\mathbf{c} in Eq. 1) is replaced by an intermediate feature vector \mathbf{f} :

$$(\mathbf{f}, o) = h(\gamma(\mathbf{x}), \gamma(\mathbf{d}), \mathbf{z}_g, \mathbf{z}_a). \quad (2)$$

\mathbf{f} is mapped into a photo-realistic image using a volume and neural rendering module R and fed to a discriminator (more details in [68, 53]).

Our 3D generator (see Sec. 4) is inspired by GIRAFFE [53]. However, it learns to produce a segmentation image, a simpler task with respect to directly generating a photo-realistic image (see Sec. 1).

4 The proposed 3D-SGAN

Fig. 2 shows the proposed 3D-SGAN architecture, composed of two main modules: a 3D-based segmentation mask generator and a texture generator. The former (G_{3D}) generates semantic segmentation masks of the human body which correspond to the main body parts and depend on the type of clothes, the camera viewpoint and the human pose. On the other hand, the texture generator (G_t) takes as input these segmentation masks and translates them into a photo-realistic image, adding a texture style randomly drawn from a pre-learned marginal distribution. The two modules are trained separately.

4.1 3D generator for semantic mask rendering

Given a set of 2D human image samples $\{I^i\}_{i=1}^N$, we first use an off-the-shelf human parsing tool [4] to obtain the corresponding ground-truth semantic segmentation masks

$\{I_s^i\}_{i=1}^N$. Using $T = \{(I^i, I_s^i)\}_{i=1}^N$ as our training set, the goal is to train a two-step generative model:

$$\tilde{I} = G(\mathbf{z}_c, \mathbf{z}_s, \mathbf{z}_p, \mathbf{z}_t) = G_t(G_{3D}(\mathbf{z}_c, \mathbf{z}_s, \mathbf{z}_p), \mathbf{z}_t), \quad (3)$$

where \tilde{I} is the final generated image. The latent codes $\mathbf{z}_c \sim P_c$ (see Sec.5), $\mathbf{z}_s \sim \mathcal{N}(0, \mathbf{I})$, $\mathbf{z}_p \sim \mathcal{N}(0, \mathbf{I})$, and $\mathbf{z}_t \sim \mathcal{N}(0, \mathbf{I})$ represent, respectively: the camera view-point, the semantics (i.e., the garment type), the body pose and the human texture.

The structure of our 3D Generator G_{3D} is inspired by GIRAFFE [53] (Sec. 3). However, differently from [68, 53], which learn to generate a textured object, in our case, h learns to generate a semantically segmented image. Specifically, we use a latent semantic code (\mathbf{z}_s) to condition the final segmentation output on the type of garment. As shown in Fig. 2, \mathbf{z}_s does not influence the opacity generation branch, and it is injected into the direction-dependent branch, which finally outputs a feature vector \mathbf{f} , representing a point-wise semantic content. Formally, we have:

$$(\mathbf{f}, o) = h(\gamma(\mathbf{x}), \gamma(\mathbf{d}), \mathbf{z}_c, \mathbf{z}_s, \mathbf{z}_p). \quad (4)$$

Following [68, 53], we generate a set of pairs $\{(\mathbf{f}, o)\}$ which are finally projected into the 2D plane using a rendering module R [48, 53] (see Sec. 3), and represented by the segmentation masks \tilde{I}_s . Specifically, \tilde{I}_s is a tensor composed of n_s channels, where each channel represents a segmentation mask of the same spatial resolution of the real images in T (Fig. 2).

G_{3D} is trained jointly with a discriminator D_s , which learns to discriminate between real (I_s) and fake (\tilde{I}_s) segmentation masks (more details in Sec. 4.4).

4.2 VAE-conditioned texture generator

The goal of our texture generator G_t is twofold: (1) mapping the segmentation masks \tilde{I}_s generated by G_{3D} into a textured human image and (2) learning a marginal distribution of the human texture using the dataset T . The latter is obtained using a Variational AutoEncoder (VAE [35]) framework, which we use to learn how to *modulate* the texture style of the decoder. Specifically, as shown in Fig. 3, G_t is composed of a semantic encoder E_s , a texture encoder E_t , and a decoder De . De is based on a StyleGANv2 architecture [33], in which a style code is used to “demodulate” the weights of each convolutional layer. We modify this architecture using a variational approach, in which the style code, *at inference time*, is extracted from a learned marginal distribution. In more detail, given a segmentation tensor I_s , we use E_s to extract the semantic content which is decoded into the final image using De and a texture code \mathbf{z}_t . The latter is sampled using the VAE encoder E_t , which converts a real input image I into a latent-space normal distribution ($\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma})$), from which \mathbf{z}_t is randomly chosen:

$$(\boldsymbol{\mu}, \boldsymbol{\sigma}) = E_t(I), \mathbf{z}_t \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}), \tilde{I} = De(E_s(I_s), \mathbf{z}_t). \quad (5)$$

G_t is trained using the pairs in T . Specifically, given a pair of samples (I^i, I_s^i) , we use an adversarial loss ℓ_{adv}^t (and a dedicated discriminator D_t), jointly with a reconstruction loss ℓ_r , and a standard Kullback-Leibler divergence (\mathcal{D}_{kl}) loss (ℓ_{kl}) [35]:

$$\ell_r = \|G_t(I_s^i, \mathbf{z}_t) - I^i\|_1, \ell_{kl} = \mathcal{D}_{kl}(\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}) || \mathcal{N}(\mathbf{0}, \mathbf{I})). \quad (6)$$




Fig. 3: The proposed VAE-conditioned texture generator. ModConv stands for “Modulated Convolution” [33].




Fig. 4: A schematic representation of our consistency losses ℓ_{ss} (a) and ℓ_{ps} (b).

Note that, in the reconstruction process, G_t cannot ignore the segmentation tensor (I_s^i) and its corresponding encoder E_s . In fact, the information extracted from the real image I , and encoded using E_t , is not enough for the decoder to represent the image content, since z_t is used only as a style modulator in De .

4.3 Consistency losses for semantics and pose disentanglement

In G_{3D} , the opacity value (o), computed by h , does not depend on the latent code z_s . Despite that, we have empirically observed that the semantics (z_s) and the pose (z_p) representations are highly entangled. We presume this is due to the convolutional filters in R (Sec. 3), where the two latent factors are implicitly merged. In order to increase the disentanglement of these factors, we propose two self-supervised consistency losses.

Silhouette-based geometric consistency. This loss is based on the idea that two different body segmentations (e.g., long-sleeve vs. short-sleeve, etc.), produced using two different semantic codes z_{s1} and z_{s2} , *but keeping fixed the pose and the camera codes*, once they are binarized, should correspond to roughly the same silhouette (see Fig. 4 (a)). Formally, the proposed geometric consistency loss ℓ_{ss} is defined as:

$$\ell_{ss} = \|B(G_{3D}(z_c, z_{s1}, z_p)) - B(G_{3D}(z_c, z_{s2}, z_p))\|_1, \quad (7)$$

where $B(\tilde{I}_s)$ maps the segmentation masks \tilde{I}_s into a binary silhouette image.

Pose-based semantic consistency. Analogously to ℓ_{ss} , the proposed pose-based semantic consistency loss is based on the idea that two different pose codes should produce a similar body segmentation. However, as shown in Fig. 4 (b), despite the body being partitioned in similar semantic segments (e.g., because the clothes have not changed), when the human pose changes, the overall spatial layout of these segments can also change (e.g., see the two different arm positions in Fig. 4 (b)). For this reason, we formulate a semantic consistency loss (ℓ_{ps}) which is spatial-invariant, and it is based on the channel-by-channel comparison of two segmentation masks. In more detail, given two different pose codes \mathbf{z}_{p1} and \mathbf{z}_{p2} , and fixing the semantics and the camera code, we first produce two corresponding segmentation tensors \tilde{I}_s^1 and \tilde{I}_s^2 . Then, for each tensor and each channel, we sum all the channel-specific mask values over the spatial dimension and we get two spatial invariant vectors $\mathbf{s}_{p1}, \mathbf{s}_{p2} \in \mathbb{R}^{n_s}$. Finally, ℓ_{ps} is given by:

$$\ell_{ps} = \sum_{i=1}^{n_s} [\max(\frac{|\mathbf{s}_{p1}[i] - \mathbf{s}_{p2}[i]|}{\mathbf{s}_{p1}[i] + \epsilon}, \rho) - \rho], \quad (8)$$

where $\mathbf{s}[i]$ is the i -th channel value of vector \mathbf{s} , ϵ is a small value used for numerical stability, and ρ is a margin representing the tolerable channel-wise difference.

4.4 Training and inference

G_{3D} is trained using an adversarial loss (ℓ_{adv}^s) jointly with ℓ_{ss} and ℓ_{ps} (Sec. 4.3):

$$\ell_{3D} = \ell_{adv}^s + \lambda_1 \ell_{ss} + \lambda_2 \ell_{ps}, \quad (9)$$

where λ_1 and λ_2 are hyper-parameters controlling the contribution of each loss term.

G_t is trained using a variational-adversarial approach (VAE-GAN [37]):

$$\ell_{tr} = \ell_{adv}^t + \lambda_3 \ell_r + \lambda_4 \ell_{kl}, \quad (10)$$

where λ_3 and λ_4 are hyper-parameters, and ℓ_{tr} is the overall objective function of G_t .

G_{3D} and G_t are trained separately. However, at inference time, the tensor \tilde{I}_s , generated by G_{3D} , is fed to G_t , along with a texture code \mathbf{z}_t , randomly drawn from a standard normal distribution:

$$\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tilde{I} = De(E_s(\tilde{I}_s), \mathbf{z}_t). \quad (11)$$

4.5 Real image editing using GAN inversion

The variational method proposed in Sec. 4.2 cannot completely reconstruct the input image. For real image editing, we use a GAN inversion technique [1] to optimize the values of the latent codes corresponding to a real input image I . Since we have two separate generators (G_{3D} and G_t), the optimization process is based on two steps (see Fig. 2, bottom). Specifically, given a pair of real image and its corresponding segmentation masks (extracted using [4], see Sec. 4.1) (I, I_s), we first generate $\tilde{I}_s = G_{3D}(\mathbf{z}_c, \mathbf{z}_s, \mathbf{z}_p)$ and we optimize $\|\tilde{I}_s - I_s\|_1$ with respect to $\mathbf{z}_c, \mathbf{z}_s$ and \mathbf{z}_p . Let $\mathbf{z}_c^*, \mathbf{z}_s^*$ and \mathbf{z}_p^* be the

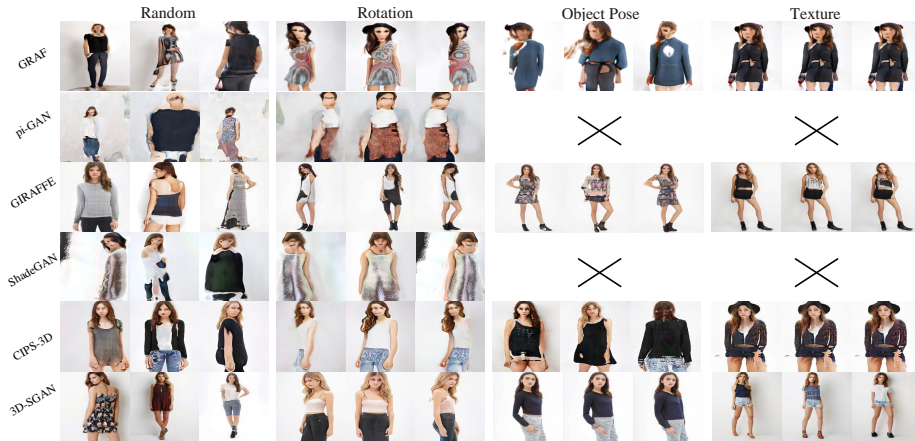


Fig. 5: A qualitative comparison. ‘Random’ means that the results are generated by randomly sampling the latent codes from the corresponding learned marginal distributions. The other 3 columns show controllable person generations with respect to the rotation, the human pose, and the texture attribute. The lack of the ‘Object Pose’ and the ‘Texture’ results for both pi-GAN and ShadeGAN is due to the fact that both methods use a single latent code to model both the texture and the geometry.

optimal values so found, and let $\tilde{I}_s^* = G_{3D}(z_c^*, z_s^*, z_p^*)$. Then, we use $\tilde{I} = G_t(\tilde{I}_s^*, z_t)$ and we optimize $\|\tilde{I} - I\|_1 + \tau LPIPS(\tilde{I}, I)$ with respect to z_t , where $LPIPS(I^1, I^2)$ is the $LPIPS$ distance between two images [84] and we use $\tau = 10$.

Once obtained the latent codes $(z_c^*, z_s^*, z_p^*, z_t^*)$ corresponding to a real image, editing can be easily done by changing these codes.

5 Experiments

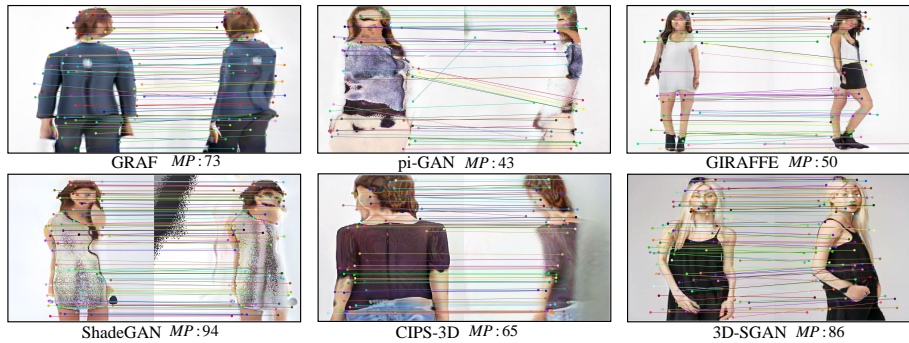
Datasets. We use the DeepFashion In-shop Clothes Retrieval benchmark [41], which consists of 52,712 high-resolution (1101×750 pixels) person images with various appearances and poses. This dataset has been widely used in pose transfer tasks. We use the following preprocessing. First, we remove overly cropped images, such as incomplete images of humans. Then, the remaining 42,977 images are resized into a 256×256 resolution, and are divided into 41,001 training and 1,976 testing images.

Training details. Following GIRAFFE [53], the camera distribution P_c can be implemented by first sampling the camera code from a uniform distribution over the dataset-dependent camera elevation angles, and then applying an object affine transformation to sample 3D points and rays. Both G_t and G_{3D} are trained using the RMSprop optimizer [36]. The learning rate for both the discriminator and the generator is set to 10^{-4} . For the loss weights, we use: $\lambda_1=0.01$, $\lambda_2=0.01$, $\lambda_3=1$, and $\lambda_4=1$. For GAN inversion, we use the Adam optimizer [34] with a learning rate of 10^{-2} .

Baselines. We compare 3D-SGAN with five state-of-the-art 3D-aware generative approaches, i.e., GRAF [68], pi-GAN [7], GIRAFFE [53], ShadeGAN [56] and CIPS-

Table 1: A quantitative comparison using the FID (\downarrow) and the aMP (\uparrow) scores.

Method	FID \downarrow			aMP \uparrow	
	Random Rotation	Object-Pose	Texture	Rotation	
GRAF [68]	52.68	176.9	57.76	220.9	64.0
pi-GAN [7]	137.6	213.7	-	-	58.0
GIRAFFE [53]	42.73	123.4	82.61	98.41	51.0
ShadeGAN [56]	134.7	232.4	-	-	89.0
CIPS-3D [85]	69.45	156.9	233.6	36.16	60.0
3D-SGAN	8.240	117.3	54.00	60.63	81.0

Fig. 6: Computing MP between pairs of generated images with 2 different viewpoints.

3D [85]. For each baseline, we use the corresponding publicly available code with a few minor adaptations for the DeepFashion dataset. Note that some concurrent methods, such as StyleNeRF [17], GRAM [11], Tri-plane [8], StyleSDF [55] achieve a performance very similar to CIPS-3D. Moreover, for some of them there is no released code yet, thus, a direct comparison is not possible. The comparison with the 2D-GAN model HumanGAN [66] can be found in the supplementary material.

Metrics. We adopt the widely used FID [23] scores to evaluate the quality of the generated human images, following common protocols (e.g., using 5,000 fake samples, etc.). And we propose the average Matched Points (aMP) to evaluate the 3D-view consistency of the generated images. The supplementary document provides the introduction of this metric.

5.1 Comparisons with state-of-the-art methods

Unconditioned human generation. Fig. 5 (“Random” column) shows a qualitative comparison between image samples generated by all the models. GRAF [68], pi-GAN [7] and ShadeGAN [56] fail to generate realistic human images. GIRAFFE [53] and CIPS-3D [85] generate reasonable human images, but they suffer from visual artifacts and texture blurs. In contrast, 3D-SGAN synthesizes much better and more photo-realistic



Fig. 7: Controllable person generation by interpolating latent codes (Rows 1-4). The fifth row shows texture generation results obtained randomly sampling z_t .

images. This qualitative analysis is confirmed in Tab. 1, where the corresponding FID scores show that 3D-SGAN significantly outperforms all the other baselines.

Controllable human generation. We analyse the representation controllability of all the models, which reflects the ability to disentangle different attributes from each other. We do this by manipulating a single latent code while fixing the others. Fig. 5 (columns “Rotation”, “Object Pose” and “Texture”) shows a qualitative comparison by varying only a single latent code. We observe that all the models can rotate the camera viewpoint. However, GRAF and CIPS-3D fail to disentangle the object pose and the texture. Moreover, pi-GAN and ShadeGAN also suffer from the same problem, since they use one single latent code to model both texture and geometry. On the other hand, both GIRAFFE [53] and 3D-SGAN can effectively disentangle the different variation factors, but GIRAFFE [53] suffers from multi-view inconsistencies and mode collapse for the texture generation. In Table 1, we use FID scores to evaluate the realistic degree of each attribute (e.g., “Rotation”, etc.). This is done computing FIDs using only the manipulated (e.g., rotated) fake images, which are compared with all the real images in the dataset. Note that this protocol cannot measure the attribute-based consistency. In most cases, 3D-SGAN has better FID scores than the other baselines.

In order to evaluate the 3D-view consistency, we use our proposed aMP metric (Sec. 5). Table 1 shows that 3D-SGAN gets the best aMP scores with respect to all the other methods except from ShadeGAN, which however generates much less realistic images, as testified by the very high FIDs (134.7 vs. our 8.24, Table 1, first column) and qualitatively shown in Fig. 6.

Fig. 7 shows additional controllable human image generation results obtained with 3D-SGAN. The generated images are realistic and, in most cases, the attributes are effectively disentangled. Specifically, Fig. 7 (1-st row) shows camera rotation results. The images generated by interpolating the camera pose parameter are consistent, and

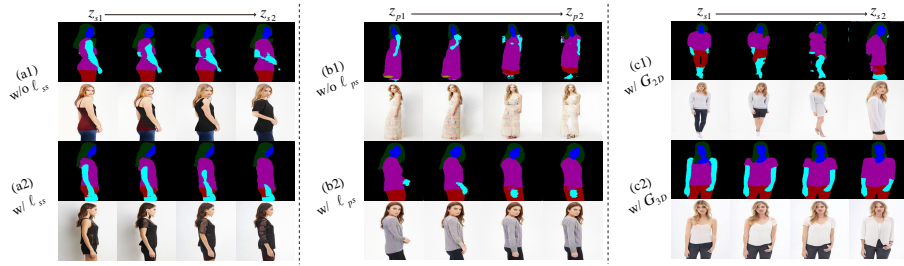


Fig. 8: A qualitative analysis of ℓ_{ss} (a), ℓ_{ps} (b), and G_{3D} (c). (a) and (c) show interpolation results between semantics codes z_{s1} and z_{s2} . (b) shows interpolation results between pose codes z_{p1} and z_{p2} .

Table 2: A quantitative analysis of ℓ_{ss} (left) and ℓ_{ps} (right). In the latter case, we use LPIPS to measure the diversity of sample pairs generated by interpolating z_p .

Metrics w/o ℓ_{ss}		w/ ℓ_{ss}		Metrics w/o ℓ_{ps}		w/ ℓ_{ps}	
L1 ↓	5.2489	3.9614		LPIPS ↓	0.1132	0.0393	

the transition from one image to the next is smooth, while simultaneously preserving the other attributes such as the texture and the pose. On the other hand, the second row shows images generated by interpolating the pose code. We again observe that human identity has been well preserved. Similarly, the other rows show that the non-target attributes have been well preserved. Finally, the third row shows that the head poses from left to right undergo only minor changes (“face frontalization”). This is likely due to both the limited training data and the data bias of the typical fashion images, where people have a frontal face. Additional results are shown in the Supplementary Material.

5.2 Ablation study

The consistency losses. Fig. 8 (a) shows a comparison between the results generated by 3D-SGAN with and without ℓ_{ss} . The effectiveness of ℓ_{ss} is shown by observing that, when removed, the generation process suffers from serious geometric inconsistencies. Specifically, the segmentation masks in Fig. 8 (a1) have undesirable pose variations, while Fig. 8 (a2) shows that ℓ_{ss} can largely alleviate this problem. To quantitatively evaluate this effect, we randomly sample two different semantic codes and we compute the $L1$ distance between the silhouettes of the corresponding generated segmentations. We average the scores over 500 different samples. The results reported in Table 2 (left) validate the effectiveness of this loss for improving the geometric consistency.

Analogously, Fig. 8 (b) qualitatively evaluates the impact of ℓ_{ps} with respect to the semantic consistency over different pose codes. For instance, in Fig. 8 (b1) there is no “red” region in the segmentation masks in the first and in the second column. However, this region is present in columns 3 and 4. Conversely, Fig. 8 (a2) shows that ℓ_{ps} can alleviate this phenomenon. To quantitatively evaluate ℓ_{ps} , we use LPIPS [84], and we



Fig. 9: An analysis of the impact of the semantic masks (a) and the VAE-conditioned texture generator (b).

Table 3: A quantitative analysis of the 3D generator G_{3D} , the semantic masks (SMs) and the VAE in our 3D-SGAN.

Metrics	w/ G_{2D}	w/o SMs	w/o VAE	full
FID ↓	13.24	66.79	14.35	8.240

measure the average *pairwise* diversity of the sample pairs generated by interpolating z_p (the lower the diversity, the higher the intra-pair consistency). Tab. 2 (right) shows that the full model achieves a lower diversity than the variant without ℓ_{ps} .

The 3D generator. To evaluate the benefit of using a GNeRF based generator, we replace it with a vanilla GAN (G_{2D}), which takes the pose and the semantics code as inputs. In this experiment, we keep all the other modules fixed. Note that G_{2D} cannot manipulate the camera parameters and, thus, it cannot generate images from multiple viewpoints. Moreover, G_{3D} can better disentangle the semantics and the pose factors with respect to G_{2D} , as demonstrated by Fig. 8 (c), where we show interpolation results between two different semantic codes. Tab. 3 shows the G_{3D} (the full model) achieves significantly better FID scores than G_{2D} .

The semantic masks and the texture generator. Existing methods such as GRAF [68] and GIRAFFE [53] do not use an additional texture generator which translates semantic masks into textured images. In contrast, the effectiveness of our semantic-based approach is shown in Fig. 5 and Table 1. However, to provide an apple-to-apple comparison and further verify the effectiveness of the semantic masks, we use an additional baseline. Specifically, in this baseline, we render the 3D representation of G_{3D} into features rather than semantic masks and we use the texture generator to map these features into the final image. Fig. 9 (a) shows the comparison of our full model with this baseline. We observe that the baseline (w/o semantic masks) fails to generate high-quality human images. Tab. 3 shows that the full model quantitatively outperforms this baseline in terms of FID scores.

The Variational Autoencoder. We evaluate the effect of conditioning G_t using a VAE (Sec. 4.2). This is done by removing the texture encoder E_t jointly with ℓ_{kl} and ℓ_r from Eq. (10). Fig. 9 (b) shows the comparison between the VAE-based approach and this variant. Both models generate human images with a high texture variability. However, the variant w/o VAE fails to preserve semantic information, i.e., the coherence between

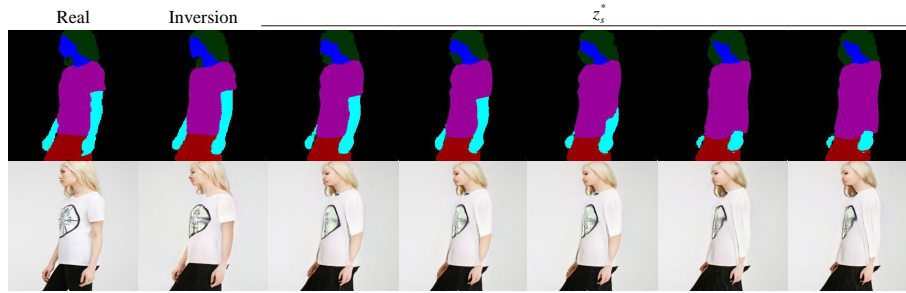


Fig. 10: Real data semantic editing results using GAN inversion.

the semantic masks, describing the clothes layout, and the final generated clothes. This shows that our VAE-based G_t learns to effectively map the semantic tensors to human images while modeling the texture distribution with the latent code z_t . Tab. 3 shows that this variant is significantly outperformed by the proposed VAE-based encoder.

5.3 Real human image editing

In this section, we use GAN inversion for real data editing tasks. The second column of Fig. 10 shows that the optimal code values (z_c^* , z_s^* , z_p^* , z_t^*), obtained using the procedure described in Sec. 4.5, lead to an effective reconstruction of the real input data (first column). In the other columns, we linearly manipulate the semantic code z_s^* while keeping fixed the other codes. Specifically, the second row of Fig. 10 shows the generated images corresponding to the semantic masks in the first row. These results demonstrate the effectiveness of the GAN inversion mechanism and the possibility to apply our model to a wide range of human image editing tasks. We computed the average LPIPS and MS-SSIM scores between real and inversion images, respectively obtaining 0.0301 and 0.912, which confirms the high reconstruction quality of our inversion.

6 Conclusion

We proposed a 3D-aware Semantic-Guided Generative model (3D-SGAN) for human synthesis. We use a generative NeRF to implicitly represent the 3D human body and we render the 3D representation into 2D segmentation masks. Then, these masks are mapped into the final images using a VAE-conditioned texture generator. Moreover, we propose two consistency losses further disentangle the pose and the semantics factors. Our experiments show that the proposed approach generates human images which are significantly more realistic and more controllable than state-of-the-art methods.

Acknowledgements This work was supported by the EU H2020 projects AI4Media (No.951911) and SPRING (No. 871245).

References

1. Abdal, R., Qin, Y., Wonka, P.: Image2stylegan: How to embed images into the stylegan latent space? In: ICCV (2019) [8](#)
2. Abdal, R., Zhu, P., Mitra, N., Wonka, P.: Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. ACM TOG (2020) [1](#)
3. Alp Güler, R., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. In: CVPR (2018) [4](#)
4. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation. IEEE TPAMI **39** (2017) [5](#), [8](#)
5. Balakrishnan, G., Zhao, A., Dalca, A.V., Durand, F., Guttag, J.: Synthesizing images of humans in unseen poses. In: CVPR (2018) [5](#)
6. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: ICLR (2019) [4](#)
7. Chan, E., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In: CVPR (2021) [2](#), [9](#), [10](#)
8. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. arXiv preprint arXiv:2112.07945 (2021) [3](#), [10](#)
9. Chen, X., Cohen-Or, D., Chen, B., Mitra, N.J.: Towards a neural graphics pipeline for controllable image generation. CGF (2021) [1](#), [2](#), [3](#)
10. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In: CVPR (2018) [1](#)
11. Deng, Y., Yang, J., Xiang, J., Tong, X.: Gram: Generative radiance manifolds for 3d-aware image generation (2022) [10](#)
12. DeVries, T., Bautista, M.A., Srivastava, N., Taylor, G.W., Susskind, J.M.: Unconstrained scene generation with locally conditioned radiance fields. In: ICCV (2021) [2](#), [3](#)
13. Gadelha, M., Maji, S., Wang, R.: 3d shape induction from 2d views of multiple objects. In: 3DV (2017) [1](#), [3](#)
14. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NeurIPS (2014) [1](#), [4](#)
15. Grigorev, A., Sevastopolsky, A., Vakhitov, A., Lempitsky, V.: Coordinate-based texture inpainting for pose-guided image generation. In: CVPR (2019) [5](#)
16. Grigorev, A., Isakov, K., Ianina, A., Bashirov, R., Zakharkin, I., Vakhitov, A., Lempitsky, V.: Stylepeople: A generative model of fullbody human avatars. In: CVPR (2021) [4](#)
17. Gu, J., Liu, L., Wang, P., Theobalt, C.: Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. ICLR 2022 (2022) [4](#), [10](#)
18. Güler, R.A., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. In: CVPR (2018) [5](#)
19. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.: Improved training of wasserstein gans. In: NeurIPS (2017) [1](#)
20. He, Z., Kan, M., Shan, S.: Eigengan: Layer-wise eigen-learning for gans. In: ICCV (2021) [1](#)
21. Henderson, P., Ferrari, V.: Learning single-image 3d reconstruction by generative modelling of shape, pose and shading. IJCV (2019) [3](#)
22. Henzler, P., Mitra, N.J., Ritschel, T.: Escaping plato’s cave: 3d shape from adversarial rendering. In: ICCV (2019) [3](#)
23. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017) [10](#)

24. Huang, S., Xiong, H., Cheng, Z.Q., Wang, Q., Zhou, X., Wen, B., Huan, J., Dou, D.: Generating person images with appearance-aware pose stylizer. In: IJCAI (2020) [5](#)
25. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: ECCV (2018) [1](#)
26. Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: ICCV (2021) [2](#)
27. Jinsong, Z., Kun, L., Yu-Kun, L., Jingyu, Y.: PISE: Person image synthesis and editing with decoupled gan. In: CVPR (2021) [5](#)
28. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. In: ICLR (2018) [1, 4](#)
29. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. In: NeurIPS (2020) [1](#)
30. Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. In: NeurIPS (2021) [1](#)
31. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: CVPR (2020) [4](#)
32. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019) [1, 4](#)
33. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: CVPR (2020) [1, 2, 3, 4, 6, 7](#)
34. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015) [9](#)
35. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013) [2, 6](#)
36. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2013) [9](#)
37. Larsen, A.B.L., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: ICML (2016) [8](#)
38. Lassner, C., Pons-Moll, G., Gehler, P.V.: A generative model of people in clothing. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 853–862 (2017) [2](#)
39. Liao, Y., Schwarz, K., Mescheder, L., Geiger, A.: Towards unsupervised learning of generative models for 3d controllable image synthesis. In: CVPR (2020) [1, 2, 3](#)
40. Liu, W., Piao, Z., Tu, Z., Luo, W., Ma, L., Gao, S.: Liquid warping gan with attention: A unified framework for human image synthesis. IEEE TPAMI (2021) [5](#)
41. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: CVPR (2016) [3, 9](#)
42. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM TOG (2015) [5](#)
43. Lunz, S., Li, Y., Fitzgibbon, A.W., Kushman, N.: Inverse graphics GAN: learning to generate 3d shapes from unstructured 2d data. CoRR [abs/2002.12674](#) (2020), <https://arxiv.org/abs/2002.12674> [3](#)
44. Lv, Z., Li, X., Li, X., Li, F., Lin, T., He, D., Zuo, W.: Learning semantic person image generation by region-adaptive normalization. In: CVPR (2021) [5](#)
45. Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. In: NeurIPS (2017) [5](#)
46. Ma, L., Sun, Q., Georgoulis, S., Van Gool, L., Schiele, B., Fritz, M.: Disentangled person image generation. In: CVPR (2018) [5](#)
47. Men, Y., Mao, Y., Jiang, Y., Ma, W.Y., Lian, Z.: Controllable person image synthesis with attribute-decomposed gan. In: CVPR (2020) [5](#)
48. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020) [2, 5, 6](#)

49. Neverova, N., Alp Guler, R., Kokkinos, I.: Dense pose transfer. In: ECCV (2018) 5
50. Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., Yang, Y.L.: Hologan: Unsupervised learning of 3d representations from natural images. In: ICCV (2019) 1, 2, 3
51. Nguyen-Phuoc, T., Richardt, C., Mai, L., Yang, Y.L., Mitra, N.: Blockgan: Learning 3d object-aware scene representations from unlabelled images. In: NeurIPS (2020) 1, 2, 3, 4
52. Niemeyer, M., Geiger, A.: CAMPARI: camera-aware decomposed generative neural radiance fields. In: 3DV (2021) 3
53. Niemeyer, M., Geiger, A.: GIRAFFE: Representing scenes as compositional generative neural feature fields. In: CVPR (2021) 2, 3, 4, 5, 6, 9, 10, 11, 13
54. Or-El, R., Luo, X., Shan, M., Shechtman, E., Park, J.J., Kemelmacher-Shlizerman, I.: StyleSDF: High-Resolution 3D-Consistent Image and Geometry Generation. CVPR2022 (2022) 4
55. Or-El, R., Luo, X., Shan, M., Shechtman, E., Park, J.J., Kemelmacher-Shlizerman, I.: Stylesdf: High-resolution 3d-consistent image and geometry generation (2022) 10
56. Pan, X., Xu, X., Loy, C.C., Theobalt, C., Dai, B.: A shading-guided generative implicit model for shape-accurate 3d-aware image synthesis. NeurIPS (2021) 2, 3, 9, 10
57. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: CVPR (2019) 1
58. Peng, S., Dong, J., Wang, Q., Zhang, S., Shuai, Q., Bao, H., Zhou, X.: Animatable neural radiance fields for human body modeling. In: ICCV (2021) 2
59. Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., Zhou, X.: Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In: CVPR (2021) 2
60. Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: ICCV (2021) 2
61. Ren, Y., Yu, X., Chen, J., Li, T.H., Li, G.: Deep image spatial transformation for person image generation. In: CVPR (2020) 5
62. Rezende, D.J., Eslami, S.M.A., Mohamed, S., Battaglia, P., Jaderberg, M., Heess, N.: Unsupervised learning of 3d structure from images. In: NeurIPS (2016) 3
63. Sanyal, S., Vorobiov, A., Bolkart, T., Loper, M., Mohler, B., Davis, L.S., Romero, J., Black, M.J.: Learning realistic human reposing using cyclic self-supervision with 3d shape, pose, and appearance consistency. In: ICCV (2021) 5
64. Sanyal, S., Vorobiov, A., Bolkart, T., Loper, M., Mohler, B., Davis, L.S., Romero, J., Black, M.J.: Learning realistic human reposing using cyclic self-supervision with 3d shape, pose, and appearance consistency. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11138–11147 (2021) 5
65. Sarkar, K., Golyanik, V., Liu, L., Theobalt, C.: Style and pose control for image synthesis of humans from a single monocular view. arXiv preprint arXiv:2102.11263 (2021) 5
66. Sarkar, K., Liu, L., Golyanik, V., Theobalt, C.: Humangan: A generative model of humans images (2021) 4, 10
67. Sarkar, K., Mehta, D., Xu, W., Golyanik, V., Theobalt, C.: Neural re-rendering of humans from a single image. In: ECCV (2020) 5
68. Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: Graf: Generative radiance fields for 3d-aware image synthesis. In: NeurIPS (2020) 2, 3, 4, 5, 6, 9, 10, 13
69. Shen, Y., Gu, J., Tang, X., Zhou, B.: Interpreting the latent space of gans for semantic face editing. In: CVPR (2020) 1
70. Shen, Y., Zhou, B.: Closed-form factorization of latent semantics in gans. In: CVPR (2021) 1
71. Siarohin, A., Lathuilière, S., Sangineto, E., Sebe, N.: Appearance and Pose-Conditioned Human Image Generation using Deformable GANs. IEEE TPAMI (2020) 5

72. Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., Sebe, N.: First order motion model for image animation. In: *NeurIPS (2019)* [5](#)
73. Song, S., Zhang, W., Liu, J., Mei, T.: Unsupervised person image generation with semantic parsing transformation. In: *CVPR (2019)* [5](#)
74. Sun, J., Wang, X., Zhang, Y., Li, X., Zhang, Q., Liu, Y., Wang, J.: Fenerf: Face editing in neural radiance fields. *arXiv preprint arXiv:2111.15490 (2021)* [3](#)
75. Tan, F., Fanello, S., Meka, A., Orts-Escolano, S., Tang, D., Pandey, R., Taylor, J., Tan, P., Zhang, Y.: Volux-gan: A generative model for 3d face synthesis with hdri relighting. *arXiv preprint arXiv:2201.04873 (2022)* [3](#)
76. Tang, H., Bai, S., Zhang, L., Torr, P.H., Sebe, N.: Xinggan for person image generation. In: *ECCV (2020)* [5](#)
77. Tov, O., Alaluf, Y., Nitzan, Y., Patashnik, O., Cohen-Or, D.: Designing an encoder for style-gan image manipulation. *ACM TOG (2021)* [1](#)
78. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *NeurIPS (2017)* [5](#)
79. Wang, X., Gupta, A.: Generative image modeling using style and structure adversarial networks. In: *ECCV (2016)* [3](#)
80. Wu, J., Zhang, C., Xue, T., Freeman, W.T., Tenenbaum, J.B.: Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In: *NeurIPS (2016)* [1, 3](#)
81. Xu, X., Pan, X., Lin, D., Dai, B.: Generative occupancy fields for 3d surface-aware image synthesis. In: *Advances in Neural Information Processing Systems(NeurIPS) (2021)* [3](#)
82. Xu, Y., Peng, S., Yang, C., Shen, Y., Zhou, B.: 3d-aware image synthesis via learning structural and textural representations. *CVPR (2022)* [3](#)
83. Yildirim, G., Jetchev, N., Vollgraf, R., Bergmann, U.: Generating high-resolution fashion model images wearing custom outfits. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*. pp. 0–0 (2019) [5](#)
84. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *CVPR (2018)* [9, 12](#)
85. Zhou, P., Xie, L., Ni, B., Tian, Q.: CIPS-3D: A 3D-Aware Generator of GANs Based on Conditionally-Independent Pixel Synthesis (2021) [2, 4, 10](#)
86. Zhou, X., Zhang, B., Zhang, T., Zhang, P., Bao, J., Chen, D., Zhang, Z., Wen, F.: Cocosnet v2: Full-resolution correspondence learning for image translation. In: *CVPR (2021)* [5](#)
87. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *ICCV (2017)* [1](#)
88. Zhu, J.Y., Zhang, Z., Zhang, C., Wu, J., Torralba, A., Tenenbaum, J.B., Freeman, W.T.: Visual object networks: Image generation with disentangled 3D representations. In: *NeurIPS (2018)* [1, 2, 3](#)
89. Zhu, Z., Huang, T., Shi, B., Yu, M., Wang, B., Bai, X.: Progressive pose attention transfer for person image generation. In: *CVPR (2019)* [5](#)