

3DHumanGAN: 3D-Aware Human Image Generation with 3D Pose Mapping

Zhuoqian Yang^{1,2†} Shikai Li¹ Wayne Wu^{1✉} Bo Dai¹
¹ Shanghai AI Laboratory ² School of Computer and Communication Sciences, EPFL
 zhuoqian.yang@epfl.ch lishikai@pjlab.org.cn
 wuwenyan0503@gmail.com daibo@pjlab.org.cn



Figure 1: **3DHumanGAN**. View-consistent full-body human images generated by our 3D-aware generative adversarial network (GAN). We show random sampled eight humans, each of them is in three view-angles.

Abstract

We present **3DHumanGAN**, a 3D-aware generative adversarial network that synthesizes photo-like images of full-body humans with consistent appearances under different view-angles and body-poses. To tackle the representational and computational challenges in synthesizing the articulated structure of human bodies, we propose a novel generator architecture in which a 2D convolutional backbone is modulated by a 3D pose mapping network. The 3D pose mapping network is formulated as a renderable implicit function conditioned on a posed 3D human mesh. This design has several merits: i) it leverages the strength of 2D GANs to produce high-quality images; ii) it generates consistent images under varying view-angles and poses; iii) the model can incorporate the 3D human prior and enable pose conditioning. Our model is adversarially learned from a collection of web images needless of manual annotation.

† Work done as research engineer at Shanghai AI Laboratory.
 Project page: <https://3dhumangan.github.io/>

1. Introduction

Human image generation is a long-standing topic in computer vision and graphics with applications across multiple areas of interest including movie production, social networking and e-commerce. Compared to physically-based methods, data-driven approaches are preferred due to the photorealism of their results, versatility and ease of use [61]. In this work, we are interested in synthesizing full-body human images with a 3D-aware generative adversarial network (GAN) that produces appearance-consistent images under different view-angles and body-poses.

Rapid developments have been seen in using 3D-aware GANs to generate view-consistent images of human faces [44, 56, 5, 4, 13, 68, 21]. However, these methods have limited capacity when dealt with complex and articulated objects such as human bodies. To begin with, methods based solely on neural volume rendering [5, 44, 56] are too memory inefficient. Rendering human bodies requires a volumetric representation that is much more dense than that of faces, which makes it computationally infeasible. A line of work improves the computational efficiency and rendering quality

of 3D-aware GANs by refining the rendered output with a convolutional neural network [68, 13, 4, 21]. However, we argue that this method is not optimal for generating full-body human images. This is because the 3D representation has to capture both the shape and the appearance of human bodies at the same time, which requires a high level of representational capacity. Meanwhile, the potential of the 2D-network is not fully exploited.

As depicted in Figure 2a, our work introduces a novel generator architecture in which a 2D convolutional backbone is modulated by a 3D pose mapping network. This design is motivated by the observation that in a StyleGAN2 [27] model trained on human images certain layers of styles correlate strongly with the pose of the generated human [9] while others correlate more apparently with the appearance. The 3D Pose Mapping Network is formulated as a renderable implicit function conditioned on a posed 3D human mesh derived with a parametric model [36]. In this way, the 3D representation handles the simplified task of parsing a geometric prior. As an additional benefit of explicitly conditioning on posed human mesh, the pose of the generated human can be specified. The output of the 3D pose mapping network is used to render a 2D low-resolution style map through ray integration [41]. The style map is passed into the first few layers of our backbone network. The appearance of the generated human is controlled by the Appearance Mapping Network. It is formulated as an MLP following a common practice in style-based generators [26, 27]. For the network to learn to parse the 3D geometric prior, we use a segmentation-based GAN loss [59] calculated using a U-Net [51] discriminator. This design enables the network to establish one-to-many mapping from 3D geometry to synthesized 2D textures using only collections of single-view 2D photographs without manual annotations.

Traditional CNN generator networks with 3×3 convolution and progressive upsampling are not equivariant and demonstrate inconsistency under geometric transformations [25]. In our case, the appearance of the generated human may change when pose and view-angle vary. To preserve consistency, we propose a network design with two key aspects: 1) Our backbone network is built entirely with 1×1 convolutions. This helps eliminate positional reference and promotes equalvariance. 2) Modulation from the pose mapping network is passed into the backbone by means of spatial adaptive batch normalization instead of the commonly used instance normalization [27, 46], so that underlying structure from the geometric information parsed by our 3D pose mapping network is preserved.

Our contributions can be summarized as follows: 1) We propose a 2D-3D hybrid generator which is both efficient and expressive. The model is supervised with segmentation-based GAN loss which helps establish a mapping between 3D coordinates and 2D human body semantics. 2) Our

generator is carefully designed to preserve the consistency of appearance when pose and view-angle vary. 3) Our work achieves state-of-the-art fidelity of 3D-aware generation of full-body human images.

2. Related Works

3D-Aware Image Generation. Generative Adversarial Networks [11] are able to generate images of human/animal faces, cars, indoor and natural scenes in eye-deceiving photorealisticness [26, 27, 25]. This motivates researchers to explore adversarial learning’s potential in achieving 3D-aware image synthesis. Earlier efforts devised different network structures to process meshes [32, 60], voxels [69, 10, 18] and block-based representations [16, 33]. These methods suffer from the common problem of insufficient 3D inductive bias [4]. Neural implicit representations rendered via ray integration have been found to be an effective representation for data-driven 3D synthesis [35]. Progress was gained in generating objects with simple structure such as human faces [5, 44, 56]. These approaches are burdened by heavy computational cost and limited to generating low-resolution images. Several works propose to remedy this by upsampling the rendered results with a 2D convolutional network and achieved impressive quality [68, 13, 4, 21], yet still not ready for the challenging 3d-aware human image generation. This is because the complexity of the articulated structure of human bodies challenges the 3D representations used by these works.

Human Image Generation. Previous works on person image generation focus on re-synthesizing the reference image, such as rendering humans in a novel pose [37, 70, 1, 6, 52], with different garments [15, 30] or with text [23]. Earlier methods condition sparse key-points [37, 38, 70, 49, 53] or semantic maps [7, 14, 64, 12] to manipulate the 2D images. To better preserve the appearance consistency of the source, surface-based methods [31, 34, 1] are introduced to establish the dense correspondences between pixel and human prior surface, *i.e.*, DensePose [50] and SMPL [36]. A limited amount of work focuses on human synthesis without a reference image, where [12] and [54] map the appearance into a gaussian distribution to enable free sampling of identities. Inspired by the success of StyleGAN [26, 27, 25] in unconditional generation, several works [9, 8] explore the capacity of the model for full-body synthesis from different aspects. Specifically, StyleGAN-Human [9] takes a data-centric perspective toward human generation and discusses the effect of manipulated style embedding on human images. InsetGAN [8] proposes a multi-GAN optimization framework to synthesize more plausible-looking humans. All methods mentioned above are 2D methods.

Three concurrent works [2, 65, 20] tackle 3D-aware human image generation based on different 3D representations. [2] and [65] render a low-resolution image with a 3D triplane

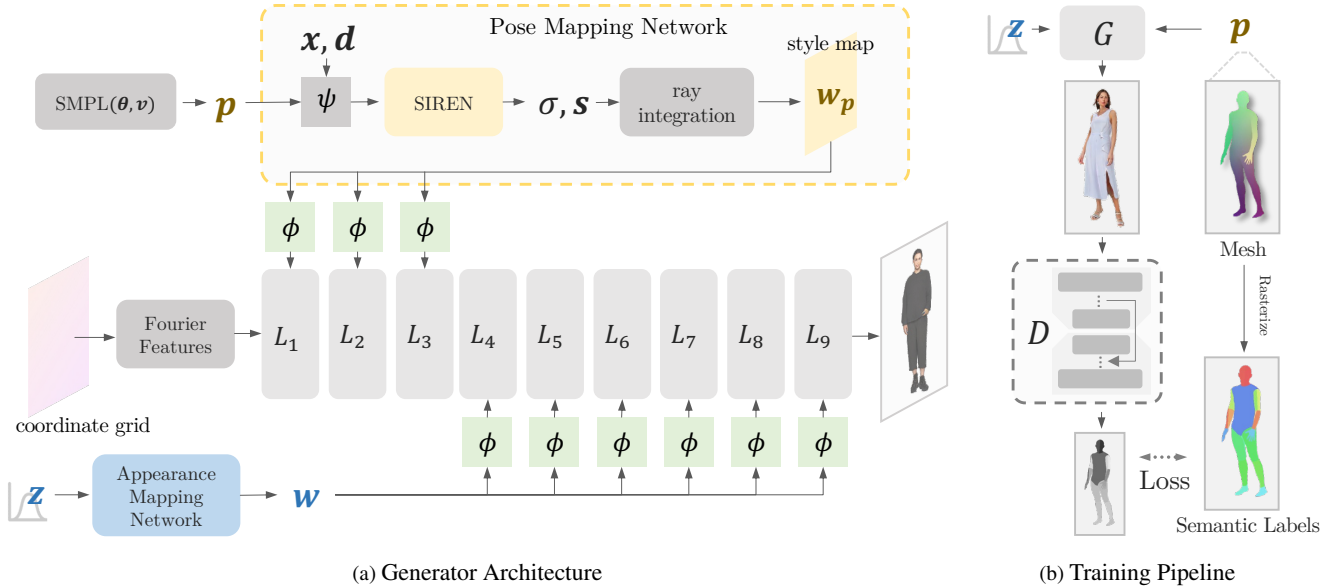


Figure 2: **Architecture.** (a) Our generator uses a pixel-wise independent convolutional backbone modulated by a 3D pose mapping network and an appearance mapping network. (b) Our model is supervised with a segmentation-based GAN loss where the semantic labels are rasterized from the conditioning mesh.

representation [4] and uses a 2D convolutional network to enlarge the image in a feed-forward manner. [20] render images directly with an efficient compositional 3D neural field. Our work is different from these works in the following aspects: 1) We use an equivariant 2D generator modulated by 3D human body prior instead of directly rendering the image from a 3D representation. 2) We use segmentation-based adversarial supervision which encourages not only fidelity but also semantic correspondence.

3. Generator Architecture

We build a generative adversarial network that synthesizes 3D-aware full-body human images with specified pose \mathbf{p} and view-angle \mathbf{v} . Pose is specified as a set of coordinates of the vertices in a posed body mesh $\mathbf{p} = \{\mathbf{v}_i \in \mathbb{R}^3\}_{i=1..6890}$ derived using a parametric human body model SMPL [36] $\mathbf{p} = \text{SMPL}(\boldsymbol{\theta}, \mathbf{v})$. $\boldsymbol{\theta} \in \mathbb{R}^{K \times 4}$ denotes the rotation at each body joint in the quaternion format. The appearance of the generated human is randomly sampled, represented as a vector $\mathbf{z} \in \mathbb{R}^{N_z}$, but required to stay consistent when pose and view-angle vary. To this end, we design a 2D-3D-hybrid generator architecture and a training strategy to achieve partially conditioning on pose.

Architecture. An overview of our generator is presented in Figure 2a. We use a style-based generator but with a major difference from prior works: different layers of the convolutional backbone are separately modulated by two mapping networks, *i.e.*, a pose mapping network that handles 3D human-body geometry and an appearance mapping network which resembles that of other style-based generators. This

design is motivated by the observation that in a StyleGAN2 [27] model trained on human images. The low-level styles correlate strongly with the pose and orientation of the generated human [9] while others correlate more apparently with the appearance, as shown in Figure 3. This inspires us to inject the human geometric prior by calculating the low-level styles from the conditioning posed mesh.

3D Pose Mapping Network. To parse the 3D geometric information, the pose mapping network is formulated as a renderable 3D implicit function $f_p(\psi(\mathbf{x}|\mathbf{p}, \boldsymbol{\theta}), \mathbf{d}) = (\sigma, \mathbf{s}_x)$ where $\mathbf{x} \in \mathbb{R}^3$ is a coordinate in the 3D camera space; $\mathbf{d} \in \mathbb{R}^3$ denotes the orientation of the camera ray; $\mathbf{p} \in \mathbb{R}^{6890 \times 3}$ denotes the coordinates of the body vertices in camera space; $\psi(\mathbf{x}|\mathbf{p}, \boldsymbol{\theta})$ transforms the coordinate \mathbf{x} from the camera space to a canonical space, added to facilitate learning. Following literature in 3D human body reconstruction [67, 48, 40], ψ is defined as the inverse process of linear blend skinning,

$$\psi(\mathbf{x}|\mathbf{p}, \boldsymbol{\theta}) = \left(\sum_{k=1}^K \omega_k(\mathbf{x}_p) \mathbf{G}_k(\boldsymbol{\theta}) \right)^{-1} \mathbf{x}$$

$\omega_k(\mathbf{x}, \mathbf{p})$ returns the blend weights corresponding to the vertex nearest to the querying point \mathbf{x} ; $\mathbf{G}_k(\boldsymbol{\theta}) \in \text{SE}(3)$ denotes the cumulative linear transformation at the k^{th} skeleton joint. The outputs of the implicit function are the opacity $\sigma \in \mathbb{R}$ and a style vector $\mathbf{s} \in \mathbb{R}^{N_s}$, which are used to render a 2D low-resolution pose style map $\mathbf{w}_s \in \mathbb{R}^{H_s \times W_s \times N_s}$ via ray integration [41]. The implicit function is parameterized as an MLP with periodic activation [58] to handle low-dimensional coordinate input.



Figure 3: **Style mixing with StyleGAN2.** The source and reference images are sampled with StyleGAN2. Style mixing is to replace certain levels of styles of the source with their counterparts from the reference. *Low*-level styles modulate the first few layers of the StyleGAN2 generator while *high*-level styles modulate the last few layers.

Neural Rendering. We use classic volume rendering technique to aggregate predicted style vectors across space. We start by evenly sampling N points $\{\mathbf{x}_i = \mathbf{o} + t_i \mathbf{d}\}$ within near and far bounds $[t_n, t_f]$ along each camera ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$. \mathbf{o} denotes the camera center. The style vector on each 2D spatial location is then estimated via

$$\mathbf{S}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma(\mathbf{x}_i) \delta_i)) \mathbf{s}(\mathbf{x}_i, \mathbf{d}),$$

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma(\mathbf{x}_j) \delta_j\right),$$

where $\delta_i = |\mathbf{x}_{i+1} - \mathbf{x}_i|$ denotes the distance between adjacent samples.

Appearance Mapping Network. The appearance mapping network is an MLP $f_a(\mathbf{z}) = \mathbf{w} \in \mathbb{R}^{N_w}$, following common practice of style-based generators [26, 27, 25]. The output style vectors \mathbf{w} do not have spatial dimensions. The final output image is synthesized by the convolutional backbone under modulation from both mapping networks.

4. Training

Our training pipeline is illustrated in Figure 2b. For the network to learn to parse the 3D geometric prior and synthesize images with specified pose, a mapping from 3D geometric information to 2D textures needs to be established. Inspired by

[55], we use a U-Net [51] discriminator architecture together with a segmentation-based GAN loss to establish a one-to-many mapping between geometry and textures. Specifically, the U-Net discriminator classifies each pixel as *fake*, *background* or one of 25 semantic classes (e.g. *head*, *torso* ...), unlike the binary classification on entire images used in traditional GAN training. This approach simultaneously enables free appearance sampling and pose conditioning. A popular practices in pose-conditioned human image generation it to use VAE-style supervision, but the quality of generated image is often outperformed by that of GAN-style training [8]. When updating the discriminator, triplets of real images, corresponding SMPL meshes and semantic label maps $(\mathbf{I}, \mathbf{p}_\mathbf{I}, \mathbf{m}_\mathbf{I})$ are needed to calculate the loss.

$$\mathcal{L}_{\text{adv}}^D = -\mathbb{E}_{\mathbf{I}} \left[\sum_c \alpha_c \sum_{x,y} \mathbf{m}_{\mathbf{I}_{x,y,c}} \log D(\mathbf{I})_{x,y,c} \right]$$

$$- \mathbb{E}_{\mathbf{p}, \mathbf{z}} \left[\sum_c \alpha_c \sum_{x,y} \mathbf{m}_{\mathbf{0}_{x,y,c}} \log D(G(\mathbf{p}, \mathbf{z}))_{x,y,c} \right]$$

where x, y and c are subscripts of the height, width and class dimension. α_c denotes class weights, calculated as the inverse of the per-pixel class frequency. $\mathbf{m}_\mathbf{0}$ denotes a semantic label map in which everywhere is 0, corresponding to the *fake* class. When updating the generator, pairs of SMPL meshes and semantic label maps $(\mathbf{p}, \mathbf{m}_\mathbf{p})$ are required.

$$\mathcal{L}_{\text{adv}}^G = -\mathbb{E}_{\mathbf{p}, \mathbf{z}} \left[\sum_c \alpha_c \sum_{x,y} \mathbf{m}_{\mathbf{p}_{x,y,c}} \log D(G(\mathbf{p}, \mathbf{z}))_{x,y,c} \right]$$

The segmentation-based GAN loss is sufficient for training the model to generate consistent images with desired pose and view-angle. In practice, we find adding two additional loss terms improves image quality. The first is a perceptual loss [24] which minimizes the difference between feature activations extracted from real and generated images using a pretrained VGG [57] network.

$$\mathcal{L}_{\text{perceptual}}^G = |\text{VGG}(G(\hat{\mathbf{z}}_\mathbf{I})) - \text{VGG}(\mathbf{I})|$$

where $\hat{\mathbf{z}}_\mathbf{I}$ is an appearance latent code corresponding with the groundtruth image \mathbf{I} . These appearance codes are learnable parameters in the model, initialized by inverting a pretrained StyleGAN2 model [9] with e4e [62]. During training, the collection of appearance latent codes are optimized but anchored by a latent loss.

$$\mathcal{L}_{\text{latent}}^G = |\hat{\mathbf{z}}_\mathbf{I} - \mathbf{e4e}(\mathbf{I})|$$

It is reported that gradient penalty on the discriminator facilitates stable training and convergence [39, 27]. We implement an R1 regularization for our segmentation-based GAN loss:

$$\mathcal{L}_{R1}^D = \frac{\gamma}{2} \mathbb{E}_{\mathbf{I}} \|\nabla_{\mathbf{I}} D(\mathbf{I})\|^2$$




Figure 4: **Spatial adaptive batch normalization** first performs a statistics-based per-channel normalization on the feature maps and then a re-normalization with mean and variance values calculated from the style maps.

The total loss for training the generator is the sum of the above terms.

$$\mathcal{L}^G = \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}^G + \lambda_{\text{perceptual}} \mathcal{L}_{\text{perceptual}}^G + \lambda_{\text{latent}} \mathcal{L}_{\text{latent}}^G + \gamma \mathcal{L}_{R1}^D$$

5. Preserving Consistency

Although capable of producing high-quality results, CNN-based generator networks often sacrifice consistency under geometric transformations [4]. This is a crucial challenge in our approach. Our network is carefully designed to preserve the consistency of appearance against varying pose and view-angle.

Pixel-wise independent convolutional backbone. Our backbone network is based entirely on 1×1 convolutions and no upsampling is used. The input to the backbone network is a constant coordinate grid $\mathbf{g} \in \mathbb{R}^{H \times W \times 2}$ containing normalized screen coordinates ranged $[-1, 1]$. Thereby, each pixel is synthesized independent of others. The purpose of this is to eliminate unwanted positional references so that the network is equal-variant to geometric transformations [25].

Spatial-adaptive batch normalization. The spatial adaptive batch normalization which is used to pass the style maps into the backbone in a way that does not impair underlying structure. Let us denote by $\mathbf{f}^{(i)}$ the output feature map of the i^{th} layer of the backbone. A spatial adaptive normalization [46] function $\phi(\mathbf{f}^{(i)}, \mathbf{w}_s) = \hat{\mathbf{f}}^{(i)}$ is needed to inject the style maps rendered by our pose mapping network into the feature maps. As shown in Figure 4, spatial adaptive normalization first performs a statistics-based per-channel normalization on the feature maps and then a re-normalization with mean and variance values calculated from the style maps. In the first normalization step, how the statistics are calculated makes a crucial difference. When the statistics are calculated on individual feature maps in a mini-batch, *i.e.* instance normalization [63], which is used by [26, 46], consistency is impaired. We hypothesize that this is because instance nor-




Figure 5: **Distribution of view-angle (yaw) in the SHHQ dataset.** We sample view-angle from a normal distribution fitted to the dataset for training and evaluation.

malization removes global affine transformation applied by the previous spatial normalization and forces the normalized layer to resort to controlling only the presence of finer features rather than their precise positions, as similarly observed in [25]. We propose to use *spatial adaptive batch normalization*:

$$\phi(\mathbf{f}_{x,y,c}^{(i)}, \mathbf{s}) = \gamma^{(i)}(\mathbf{s})_{x,y,c} \cdot \frac{\mathbf{f}_{x,y,c}^{(i)} - \mathbb{E}_{\mathbf{p},\mathbf{z}}[\mathbf{f}_{x,y,c}^{(i)}]}{\sqrt{\mathbb{V}_{\mathbf{p},\mathbf{z}}[\mathbf{f}_{x,y,c}^{(i)}]}} + \beta^{(i)}(\mathbf{s})_{x,y,c}$$

where x, y and c are the subscripts for the width, height and channel dimension. $\gamma^{(i)}$ and $\beta^{(i)}$ are learned affine transformations. \mathbb{E} and \mathbb{V} estimates the means and variances of the activations of a feature channel across all pixels and all combinations of poses, views and appearances. In practice, this is done by keeping track of the mini-batch statistics of the per-channel activations through exponential moving average. Spatial adaptive batch normalization is used throughout the generation pipeline. For the style vectors \mathbf{w} calculated with the appearance network, the same re-normalization is broadcasted to all spatial locations.

6. Experiments

6.1. Dataset

We compare 3D-aware human image generation performance on SHHQ [9], a collection of $\sim 220\text{k}$ real-world full-body human images. All compared models are trained on this dataset. We use an off-the-shelf model PARE [29] to register a SMPL mesh for each of the training images. The semantic label maps are then obtained by rasterizing the conditioning posed mesh (mapping between mesh vertices and semantic labels is known). We augment the semantic label maps with estimations from a 2D semantic segmentation method DensePose [50]. The latter’s results are better aligned with the images compared to the estimations from PARE. During training, pose, view and appearance are independently sampled. Poses are directly sampled from data. The view-angles for training and evaluation are sampled from a normal distribution fitted to the dataset *i.e.* $\text{yaw} \sim \mathcal{N}(0, 0.4)$ as shown in Fig. 5.



(a) Qualitative Results



(b) Appearance Interpolation

(c) Pose Interpolation

Figure 6: **Qualitative Results.** In (a), each row shows two cases separated by the dotted line. For each case we show one identity in two poses and three view-angles. The conditioning mesh is shown on the left of each case. (b) shows two cases of appearance interpolation. (c) shows two cases of pose interpolation.

6.2. Baselines

We compare with a 2D unconditional method StyleGAN2 [27], a 2D conditional method OASIS [55], 3D-aware unconditional methods StyleNeRF [13] and EG3D [4] and two 3D-aware conditional methods EVA3D [20] and ENARF-GAN [45]. Note that we do not compare with fully-3D methods such as [5, 56, 44] because these methods do not support training at a meaningful resolution. To evaluate and compare the pose-conditioning capability, we implement pose conditioning for the unconditional methods by first inverting a human image with desired pose to the style-space and then mixing these style vectors with the ones calculated from the randomly sampled z . For StyleGAN2, the inversion method used is e4e[62]. For StyleNeRF and EG3D, we train a ResNet [17] encoder to infer the latent code given camera parameters.

6.3. Evaluation Metrics

Fidelity. We measure fidelity with Fréchet inception distance [19] and kernel inception distance [3]. Two sets of

fidelity metrics are calculated for unconditional generation and pose-conditional generation, denoted as $(FID-u, KID-u)$ and $(FID-p, KID-p)$, respectively. These values are calculated between 50,000 real and generated images using an alias-free implementation [47].

Accuracy of Pose Conditioning. To evaluate the accuracy of generated pose under conditioning, we use PARE [29] to re-estimate poses from generated images and calculate the 2D mean per-joint position error [22] in normalized screen space ranged $[-1, 1]$. The procedure is: i) We randomly sample 4000 poses from the dataset to serve as ground truth; ii) we use the candidate method to generate 4000 images conditioned on these poses; iii) the 3D poses of generated persons are estimated using PARE and used to calculate the MPJPE (Protocol #1) [?].

Consistency. We measure the consistency of appearance under varying pose and view-angle using the signal restoration metric peak signal-to-noise ratio (PSNR) in decibels (dB), following [25, 66]. Specifically, we perform a reverse mapping from pixels to mesh vertices with meshes estimated using [29]. The PSNR is then calculated between

Table 1: **Quantitative Comparisons** on human images generated by 2D and 3D methods. *KID* values are reported in 10^{-3} units; *Pose* values are reported in 10^{-2} units. The first place and runner-up in each metric are reported in **bold** typeface. Our method achieves the lowest fidelity scores and the highest score of appearance consistency under different body-poses.

Metrics	StyleGAN-Human, 512	OASIS, 256	ENARF-GAN, 128	EVA3D, 256	StyleNeRF, 512	EG3D, 512	Ours, 256	Ours, 512
FID-p ↓	6.68	10.74	30.40	26.91	13.13	17.54	6.61	9.31
KID-p ↓	3.60	8.82	23.66	18.71	7.01	10.71	3.01	5.16
FID-u ↓	2.83	-	-	-	7.60	13.05	-	-
KID-u ↓	1.50	-	-	-	3.96	8.36	-	-
Pose ↓	4.48	1.82	6.66	3.12	6.64	8.39	2.21	2.08
Con-v ↑	-	15.68	19.39	25.95	18.29	19.93	21.83	22.84
Con-p ↑	12.77	13.19	15.67	19.16	13.74	11.91	17.72	18.93

Table 2: **Ablation study** quantitative results. *KID* values are reported in 10^{-3} units; *Pose* values are reported in 10^{-2} units. The first place and runner-up in each metric are reported in **bold** typeface.

Config.	FID-p ↓	KID-p ↓	Pose ↓	Con-v ↑	Con-p ↑
VAE	8.00	4.02	9.54	19.22	15.18
FF	9.42	4.57	2.60	21.63	17.89
Up	8.28	5.42	1.89	16.03	13.16
IN	8.90	4.04	1.93	17.73	13.36
Full _{mixed}	7.98	3.95	2.46	24.80	18.54
Full _{all}	8.15	4.25	2.28	22.23	18.60
Full	6.61	3.01	2.21	21.83	17.72

corresponding vertices on these colored meshes. The *Con-v* and *Con-p* respectively denote PSNR calculated under view change and pose change.

6.4. Qualitative Results

Figure 6a shows 3D-aware human image generation results in multiple poses and view-angles. The appearances of the humans generated with our method is consistent under view and pose changes. Figure 6b and 6c show the interpolation results. Our model is able to generate smooth interpolation of appearance and pose.

6.5. Comparisons

Quantitative results are presented in Table 1. Qualitative results are shown in Figure 7. Our method demonstrates the best overall performance, with an FID comparable to the 2D state-of-the-art method StyleGAN-Human, competitive pose conditioning accuracy and top-level consistency. Using GAN inversion methods, we can produce images with StyleGAN2, StyleNeRF and EG3D that match the condition pose closely but not exactly. This suggests that the representations learned by these methods are entangled in appearance and pose. Style mixing is not enough to achieve clean separation in the two domains. Our method is able

to generate images that accurately matches the conditioning pose. Relative to leading 3D-aware, pose-conditional approaches like EVA3D and ENARF-GAN, our technique exhibits superior capability in producing high-fidelity images. EVA3D showcases remarkable consistency, it is plausible that their single-stage rendering framework, circumventing any upsampling or the use of any 2D convolutional network, contributes to this performance. However, this choice comes at a cost of high GPU memory consumption. Besides, all of the compared baselines except EVA3D show different extents of inconsistencies under view and pose variation. This is another indicator of entanglement of appearance and pose in the learned latent spaces of the baselines. In contrast, the appearance of the humans generated by our method remains consistent under pose and view variation.

6.6. Ablation Study

We compare our 3d-aware generator with different ablated models that use previous methods instead of our designed modules to measure the effectiveness of each module. Specifically, to examine the effects of passing the 3D geometric prior into the 2D backbone as styles, i.e. the 2D-3D hybrid generator, we compare with the feed-forward approach used by [13, 4, 68], denoted as (*Feed Forward*). To evaluate the contribution of the segmentation-based GAN loss, we train a model in which this loss is replaced with the traditional GAN loss with binary discrimination. This results in a model trained under VAE-style supervision as in [21], denoted as *VAE*. To examine the effectiveness of the pixel-independent backbone design, we compare with a variant that uses traditional upsampling convolution network like in [55, 27], which is denoted as *Upsampling*. To examine the effectiveness of spatial adaptive batch normalization, we compare with a variant that uses spatial adaptive instance normalization instead, denoted as *Instance Norm*.

Based on different ways of injecting the styles into the convolutional backbone, we derive three variants of the full model. The default model, *Full*, introduces pose style maps \mathbf{w}_p into the initial three layers and appearance styles \mathbf{w} into the subsequent six. *Full_{all}* adds the elementwise sum of \mathbf{w}_p and \mathbf{w} to every layer, while *Full_{mixed}* incorporates \mathbf{w}_p into the first three and \mathbf{w} throughout all layers.

All models in this experiment are trained and evaluated in the resolution of 256×128 . Quantitative results are presented in Table 2 and qualitative results are presented in the supplementary materials. Our full models perform best in fidelity and view-consistency. The *Feed-Forward* demonstrates favorable consistency, but with no small sacrifice in image quality. In other terms, passing the geometric prior in the form of low-level styles has a small toll on pose-accuracy and consistency, but completely within the acceptable range. Meanwhile, the image quality is greatly improved. By examining the results of the *VAE* setting, we



Figure 7: **Qualitative Comparison.** We show two cases separated by the dotted line. For each case, the first row shows unconditional generation results, the second row shows pose-conditional generation results. We show three view angles, from -30° to 30° .

can see that segmentation-based GAN loss is crucial to pose accuracy. It also improves fidelity, possibly because this loss helps the model to make better use of the 3D Human Prior. The *Upsampling* and *Instance Norm* configurations show the best accuracy of generated pose but with impaired image consistency in different extents. These two configurations show that the pixel-wise independent backbone and spatial adaptive batch normalization indeed preserves consistency. Finally, the variants based on different ways of style injection shows that is beneficial to inject the pose styles into only the initial layers of the network.

6.7. Internal Representation

Figure 8 visualizes typical internal representations from our networks. One interesting observation is that in the model that uses batch normalization, we are able to find certain feature activations that are consistent across different poses, view-angles and appearances, such as the ones depicted. However, such activations are absent in the model with instance normalization. This may be because the instance normalization model mainly modulates the existence of image features rather than their location, as also noted by [25]. We hypothesize that batch normalization preserves an internal coordinate system that helps map different image textures to the 3D surface, which improves image consistency.

7. Implementation Details

Training Details. Since both photometric (perceptual loss) and adversarial supervision (segmentation-based GAN loss) are used when training our model, we design two different types of training iterations, *i.e.* partially-conditional iterations and conditional iterations, which are executed alternately. For partially-conditional iterations, we sample appearance codes, poses and view-angles independently. The appearance codes are sampled from a normal distribution $\mathbf{z} \sim \mathcal{N}(0, 1)$. The poses are randomly sampled from all training data. The view-angles are sampled from a uniform distribution $\mathbf{p} \sim \mathcal{N}(0, 0.4)$ based on training data statistics. Only the segmentation-based GAN loss and the $R1$ regularization are calculated for this kind of iteration. For conditional iterations, we sample quadruplets of image, appearance code, pose and view-angle $(\mathbf{I}, \mathbf{z}_\mathbf{I}, \mathbf{p}_\mathbf{I}, \mathbf{v}_\mathbf{I})$. The appearance codes $\mathbf{z}_\mathbf{I}$ are sampled from the code pool $\mathbf{Z} \in \mathbb{R}^{N_\mathbf{I} \times N_z}$ which is a trainable parameter of the model. It contains $N_\mathbf{I}$ entries of appearance codes, each corresponds with an image in the training dataset, initialized by inverting a pretrained StyleGAN2 model [9] with e4e [62]. The pose $\mathbf{p}_\mathbf{I}$ and view-angle $\mathbf{v}_\mathbf{I}$ are estimated using PARE [29] Both the perceptual loss and the segmentation-based GAN loss are calculated for conditional iterations. When calculating the segmentation-based GAN loss, we augment the ground truth semantic maps that are calculated with PARE [29] estimation results with the ones estimated by DensePose [50] with a 50% probability. This is because the results from PARE does not align with

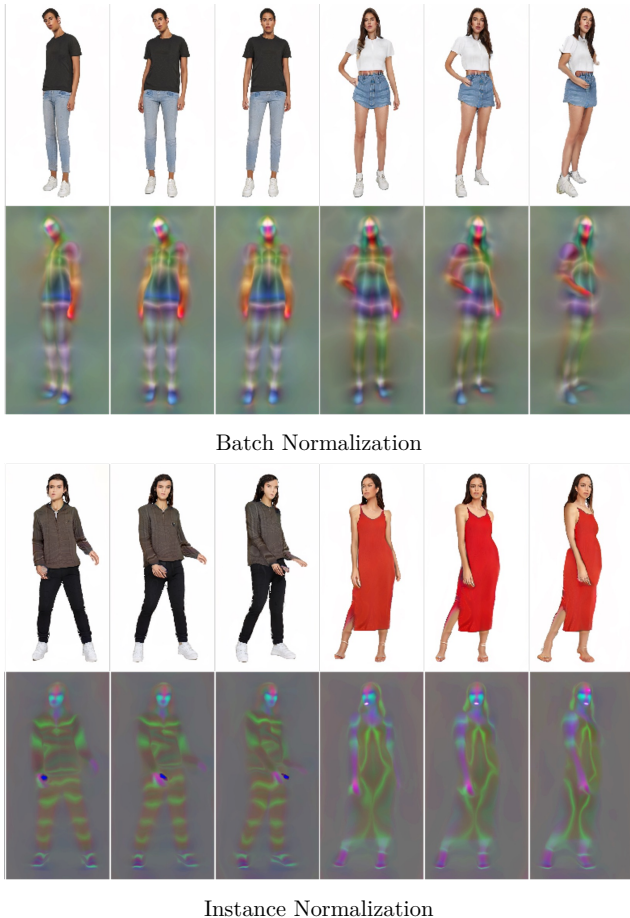


Figure 8: **Internal Representation Visualization.** Two identities in two different poses are visualized for each model. For the same model, the same feature channel are shown for the two identities.

the input image as well as the results from DensePose.

Hyperparameters. Our models are trained 300k steps using an Adam [28] optimizer. The learning rate is set as 0.0001 for the generator backbone and 0.0004 for the discriminator. The learning rate of the two mapping networks is set as 5×10^{-6} . The batch size is set as 32. Spectral normalization [42] on model parameters is used to improve training stability. The weights of the generator’s loss terms are $\lambda_{\text{adv}} = 1$, $\lambda_{\text{percept}} = 1$, $\lambda_{\text{latent}} = 0.05$ and $\gamma = 0.25$. After training for 130k steps, learning rates are decreased by half; λ_{percept} and λ_{latent} are set as 0. We use lazy execution of $R1$ regularization to speedup training. The $R1$ term is updated once every 4 iterations. Our 256×128 resolution model takes 3 days to train on 4 A100 GPUs. Our 512×256 model requires 5 days on 8 A100 GPUs.

Inversion-based baselines For unconditional generative models like StyleGAN-Human, StyleNeRF and EG3D, we follow the inversion-and-manipulation manner to explicitly control the pose/view of the human image. Specifically, for

StyleGAN-Human, we adopt the pre-trained e4e [62] model to invert the source and target images to w_s and w_t , respectively. We then replace the injection style w_s in the first six synthesis layers with w_t to control the pose and global orientation of the human. 3D-aware generative models require the extrinsic camera as a conditional input to control the image’s view. Therefore, we construct an encoder consisting of $\log_2 N$ ResBlock layers and a fully connected layer to predict the style latent w and camera parameters c . We adopt a self-supervised learning scheme by using synthesized images as input and corresponding style latent w and random camera c as ground-truth. After training, we can edit the pose/view of the source image by replacing the injection style in the first six synthesis layers with the predicted latent (as in StyleGAN-Human) and using predicted camera as condition.

8. Limitations

Our model can handle most of the view-angles and poses in the training data, but it fails to generalize beyond them. Exploring better generalization and extrapolation methods could be a promising direction for future work. We also notice that inconsistencies in the appearance are still observable in some cases when the view-angle changes. We believe this can be improved by closely inspecting the CNN backbone from the perspective of signal processing to ensure equivariance. Finally, since our framework uses a disentangled 3D representation for pose, efficient scene-specific 3D representations [43] could be employed for better computational performance.

9. Conclusion

We build a generative adversarial networks 3DHumanGAN that synthesizes full-body human images with photorealistic 3D-awareness. Our generator combines 3D geometric prior of the human body with a 2D-3D hybrid structure and produces consistent appearance across different poses and view-angles with high image quality. Our segmentation-based GAN loss is essential for guiding the generator to parse and condition on the 3D human body prior. We demonstrate through an ablation study and internal representation visualization that our pixel-independent backbone and spatial adaptive batch normalization technique effectively preserve consistency.

References

- [1] Badour Albahar, Jingwan Lu, Jimei Yang, Zhixin Shu, Eli Shechtman, and Jia-Bin Huang. Pose with style: Detail-preserving pose-guided image synthesis with conditional StyleGAN. *ACM TOG*, 2021. 2
- [2] Alexander W Bergman, Petr Kellnhofer, Yifan Wang, Eric R Chan, David B Lindell, and Gordon Wetzstein. Gener-

- ative neural articulated radiance fields. *arXiv preprint arXiv:2206.14314*, 2022. 2
- [3] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018. 6
- [4] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. *arXiv preprint arXiv:2112.07945*, 2021. 1, 2, 3, 5, 6, 7
- [5] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *CVPR*, 2021. 1, 2, 6
- [6] Wei Cheng, Su Xu, Jingtian Piao, Chen Qian, Wayne Wu, Kwan-Yee Lin, and Hongsheng Li. Generalizable neural performer: Learning robust radiance fields for human novel view synthesis. *arXiv preprint arXiv:2204.11798*, 2022. 2
- [7] Haoye Dong, Xiaodan Liang, Ke Gong, Hanjiang Lai, Jia Zhu, and Jian Yin. Soft-gated warping-gan for pose-guided person image synthesis. 2018. 2
- [8] Anna Frühstück, Krishna Kumar Singh, Eli Shechtman, Niloy J Mitra, Peter Wonka, and Jingwan Lu. InsetGAN for full-body image generation. *arXiv preprint, arXiv:2203.07293*, 2022. 2, 4
- [9] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen-Change Loy, Wayne Wu, and Ziwei Liu. Stylegan-human: A data-centric odyssey of human generation. *arXiv preprint, arXiv:2204.11823*, 2022. 2, 3, 4, 5, 8
- [10] Matheus Gadelha, Subhansu Maji, and Rui Wang. 3d shape induction from 2d views of multiple objects. In *3DV*, 2017. 2
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *NeurIPS*, 2014. 2
- [12] Artur Grigorev, Karim Iskakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. Stylepeople: A generative model of fullbody human avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5151–5160, 2021. 2
- [13] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. In *ICLR*, 2022. 1, 2, 6, 7
- [14] Xintong Han, Xiaojun Hu, Weilin Huang, and Matthew R Scott. Clothflow: A flow-based model for clothed person generation. In *ICCV*, 2019. 2
- [15] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. VITON: An image-based virtual Try-On network. In *CVPR*, 2018. 2
- [16] Zekun Hao, Arun Mallya, Serge Belongie, and Ming-Yu Liu. Gancraft: Unsupervised 3d neural rendering of minecraft worlds. In *ICCV*, 2021. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [18] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Escaping plato’s cave: 3d shape from adversarial rendering. In *ICCV*, 2019. 2
- [19] 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. *Advances in neural information processing systems*, 30, 2017. 6
- [20] Fangzhou Hong, Zhaoxi Chen, Yushi Lan, Liang Pan, and Ziwei Liu. Eva3d: Compositional 3d human generation from 2d image collections. *arXiv preprint arXiv:2210.04888*, 2022. 2, 3, 6
- [21] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. *arXiv preprint arXiv:2112.05637*, 2021. 1, 2, 7
- [22] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *TPAMI*, 2013. 6
- [23] Yuming Jiang, Shuai Yang, Haonan Qiu, Wayne Wu, Chen Change Loy, and Ziwei Liu. Text2human: Text-driven controllable human image generation. *ACM Transactions on Graphics (TOG)*, 41(4):1–11, 2022. 2
- [24] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 4
- [25] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. 2021. 2, 4, 5, 6, 8
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2, 4, 5
- [27] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 2, 3, 4, 6, 7
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 9
- [29] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *CVPR*, 2021. 5, 6, 8
- [30] Kathleen M Lewis, Srivatsan Varadharajan, and Ira Kemelmacher-Shlizerman. TryOnGAN: Body-aware Try-On via layered interpolation. *ACM TOG*, 2021. 2
- [31] Yining Li, Chen Huang, and Chen Change Loy. Dense intrinsic appearance flow for human pose transfer. In *CVPR*, 2019. 2
- [32] Yiyi Liao, Katja Schwarz, Lars Mescheder, and Andreas Geiger. Towards unsupervised learning of generative models for 3d controllable image synthesis. In *CVPR*, 2020. 2
- [33] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020. 2
- [34] Wen Liu, Zhixin Piao, Min Jie, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *ICCV*, 2019. 2

- [35] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. 2
- [36] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 2015. 2, 3
- [37] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. *arXiv preprint, arXiv:1705.09368*, 2017. 2
- [38] Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. In *CVPR*, 2018. 2
- [39] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International conference on machine learning*, pages 3481–3490. PMLR, 2018. 4
- [40] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. Leap: Learning articulated occupancy of people. In *CVPR*, 2021. 3
- [41] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2, 3
- [42] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018. 9
- [43] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 9
- [44] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *CVPR*, 2021. 1, 2, 6
- [45] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Unsupervised learning of efficient geometry-aware neural articulated representations. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*, pages 597–614. Springer, 2022. 6
- [46] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 2, 5
- [47] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11410–11420, 2022. 6
- [48] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Animatable neural radiance fields for human body modeling. In *ICCV*, 2021. 3
- [49] Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H Li, and Ge Li. Deep image spatial transformation for person image generation. 2020. 2
- [50] Iasonas Kokkinos Riza Alp Güler, Natalia Neverova. Densepose: Dense human pose estimation in the wild. 2018. 2, 5, 8
- [51] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 2, 4
- [52] Soubhik Sanyal, Alex Vorobiov, Timo Bolkart, Matthew Loper, Betty Mohler, Larry S Davis, Javier Romero, and Michael J Black. 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*, pages 11138–11147, 2021. 2
- [53] Kripasindhu Sarkar, Vladislav Golyanik, Lingjie Liu, and Christian Theobalt. Style and pose control for image synthesis of humans from a single monocular view. *arXiv preprint, arXiv:2102.11263*, 2021. 2
- [54] Kripasindhu Sarkar, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. HumanGAN: A generative model of human images. In *3DV*, 2021. 2
- [55] Edgar Schönfeld, Vadim Sushko, Dan Zhang, Juergen Gall, Bernt Schiele, and Anna Khoreva. You only need adversarial supervision for semantic image synthesis. In *ICLR*, 2021. 4, 6, 7
- [56] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *NeurIPS*, 2020. 1, 2, 6
- [57] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 4
- [58] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *NeurIPS*, 2020. 3
- [59] Vadim Sushko, Edgar Schönfeld, Dan Zhang, Juergen Gall, Bernt Schiele, and Anna Khoreva. You only need adversarial supervision for semantic image synthesis. *arXiv preprint arXiv:2012.04781*, 2020. 2
- [60] Attila Szabó, Givi Meishvili, and Paolo Favaro. Unsupervised generative 3d shape learning from natural images. *arXiv preprint arXiv:1910.00287*, 2019. 2
- [61] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering. In *CGF*, 2020. 1
- [62] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM TOG*, 2021. 4, 6, 8, 9
- [63] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016. 5
- [64] Zhuoqian Yang, Wentao Zhu, Wayne Wu, Chen Qian, Qiang Zhou, Bolei Zhou, and Chen Change Loy. Transmomo: Invariance-driven unsupervised video motion retargeting. In *CVPR*, 2020. 2
- [65] Jianfeng Zhang, Zihang Jiang, Dingdong Yang, Hongyi Xu, Yichun Shi, Guoxian Song, Zhongcong Xu, Xinchao Wang, and Jiashi Feng. Avatargen: a 3d generative model for animatable human avatars. *arXiv preprint arXiv:2208.00561*, 2022. 2
- [66] Richard Zhang. Making convolutional networks shift-invariant again. In *ICML*, 2019. 6

- [67] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Generalizable neural human radiance field from sparse inputs. *arXiv preprint arXiv:2112.02789*, 2021. [3](#)
- [68] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. CIPS-3D: A 3d-aware generator of gans based on conditionally-independent pixel synthesis. *arXiv preprint, arXiv:2110.09788*, 2021. [1](#), [2](#), [7](#)
- [69] Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Josh Tenenbaum, and Bill Freeman. Visual object networks: Image generation with disentangled 3d representations. *NeurIPS*, 2018. [2](#)
- [70] Zhen Zhu, Tengting Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive pose attention transfer for person image generation. In *CVPR*, 2019. [2](#)