

VeRi3D: Generative Vertex-based Radiance Fields for 3D Controllable Human Image Synthesis

Xinya Chen¹, Jiaxin Huang¹, Yanrui Bin², Lu Yu¹, Yiyi Liao^{1*}

¹Zhejiang University ²Huazhong University of Science and Technology
 {hust.xinyachen, jaceyh919, binyanrui}@gmail.com, {yu1, yiyi.liao}@zju.edu.cn



Figure 1: **3D Controllable Human Image Synthesis.** Our VeRi3D enables control over shape, pose, and appearance in the first row and part-level editing of the lower body, upper body, and head in the second row. The bottom right illustrates the vertices used to control the image synthesis, where colored vertices indicate the edited part.

Abstract

Unsupervised learning of 3D-aware generative adversarial networks has lately made much progress. Some recent work demonstrates promising results of learning human generative models using neural articulated radiance fields, yet their generalization ability and controllability lag behind parametric human models, i.e., they do not perform well when generalizing to novel pose/shape and are not part controllable. To solve these problems, we propose VeRi3D, a generative human vertex-based radiance field parameterized by vertices of the parametric human template, SMPL.

We map each 3D point to the local coordinate system defined on its neighboring vertices, and use the corresponding vertex feature and local coordinates for mapping it to color and density values. We demonstrate that our simple approach allows for generating photorealistic human images with free control over camera pose, human pose, shape, as well as enabling part-level editing.

1. Introduction

Generating diverse photorealistic renderings of clothed humans has a wide range of applications including visual effects, virtual try-on, VR/AR, and creative image editing.

*Corresponding authors.

Compared to designing 3D human avatars manually which is expensive and time-consuming, learning generative human models from data is a promising alternative to reduce the design effort. To meet the requirements of the aforementioned applications, the generative model should be capable of rendering photorealistic clothed humans with free control over camera pose, human pose, shape, and appearance, ideally at the part-level, such as changing a hoody from a sweater for the same individual.

Existing generative human models do not fulfill these requirements yet. Parametric 3D body models [1, 32, 47], e.g., SMPL, faithfully capture the human pose and shape statistics with explicit control over these attributes but do not model clothing. Based on SMPL, generative models for clothed human avatars have been proposed [6, 9, 35]. However, these methods do not model appearance and require 3D supervision for modeling geometry.

Recently, 3D-aware generative adversarial networks (GANs) have shown great success in learning 3D representations, e.g., neural radiance fields, from single-view 2D supervisions. While most 3D-aware GANs focus on non-articulated objects, e.g., faces and cars [53, 5, 4, 16], there are some recent attempts to learn generative human radiance fields from 2D images [42, 2, 66]. All these methods map a 3D point in the observation space to a canonical space, allowing for learning shape and appearance in a pose-agnostic space. One line of work models the mapping function via a learned blending field [42], yet learning the distribution of the blending field from single-view 2D images is highly challenging and struggles to generalize to out-of-distribution poses. Another line of works model articulation using a fixed surface-based mapping without learning parameters [2, 66]. This indeed simplifies the task of the generator and leads to better generalization, but the surface-based mapping can be inaccurate, e.g., an unoccupied background point in the observation space might be mapped to an occupied foreground point in the canonical space and thus yielding ghosting artifacts. Further, despite providing control over the human pose, the controllability over the shape and fine-grained parts is yet to be explored.

In this work, we propose generative *Vertex-based human Radiance fields*, VeRi3D, aiming for bridging the full controllability of parametric models and the image fidelity of 3D-aware GANs. Our key idea is to parametrize the generative human radiance fields using a set of vertices predefined by a human model, SMPL, and learning distributions over the vertex features to generative diverse human images. Specifically, we transform a 3D point to the individual local coordinate systems of its nearest vertices, mapping the point to a color and a density value based on the feature vectors attached to the nearest vertices and its locations in the local coordinates. Our representation has the following advantages: i) This formulation allows us to enjoy the ben-

efit of using a fixed mapping guided by SMPL pose, leaving the generator to learn a pose-agnostic human representation only. ii) Further, our formulation does not suffer from ghosting artifacts as each point is mapped to an *individual* local coordinate system instead of a *shared* global canonical space as in previous methods [2, 66]. iii) Our formulation is naturally suited for generative models as we can learn a set of pose-agnostic vertex features based on a fixed UV mapping of the vertices using a 2D convolution neural network. iv) We demonstrate that our method naturally enables control over pose and shape by editing the vertex locations. Further, it can easily achieve part-level control by manipulating vertex features of the same part.

In conclusion, our contributions are as follows: 1) We present a 3D-aware GAN method for generating controllable radiance fields of human bodies. 2) Our method introduces a human radiance field representation that enables generalization to novel poses and body shapes, as well as editing of local cloth shape and appearance. 3) We demonstrate high-quality results for unconditional generation and animation of human bodies using several datasets, including Surreal, AIST++ and DeepFashion.

2. Related Work

GAN-based Image Synthesis: Generative adversarial networks [14] have achieved 2D image synthesis with high visual fidelity [36, 21, 23, 24, 22, 52, 64]. As these unconditional methods have no explicit control over the generated images, a line of works proposes conditional human image synthesis by transferring a 2D human pose [34, 55, 12, 11, 56, 71, 13, 31], a semantic segmentation map [26] or a UV feature map [15, 51] to an RGB image. Despite enabling controllability, the shape and appearance consistency given different input conditions is not guaranteed due to the 2D-to-2D generation. In contrast to these 2D image synthesis methods, we focus on learning to generate 3D representations that can be rendered to novel viewpoints.

3D-Aware Image Synthesis: Recently, 3D-aware image synthesis has attracted growing attention by lifting the generator to the 3D space. A key to 3D-aware image synthesis is the choice of the underlying 3D representation. Early methods attempt to learn 3D voxel grid [40, 17] and mesh [29], yet these discretized representations limit the image fidelity. More recent methods exploit neural radiance fields [38] for 3D-aware image synthesis [53, 5, 4, 16, 20, 63, 70, 44, 62, 45, 10, 54]. Albeit achieving photorealistic 3D-aware image synthesis, these methods focus on non-articulated objects, e.g., faces and cars.

3D-Aware Human Image Synthesis: There are a few attempts to address the task of 3D-aware human image synthesis leveraging neural radiance fields [67, 42, 2, 66, 19].

3D-SGAN [67] leverages a semantic radiance field to render a semantic map, combined with a 2D convolutional network for converting the semantic label to a textured image. Despite demonstrating photorealistic human image synthesis, the 2D texture generator hinders multi-view consistency. Our work is more related to ENARF [42] that generates a human radiance field in a canonical space and obtains textured 2D images via volume rendering. ENARF proposes to learn a blend weight field to map humans of different poses to a canonical space. However, this learned weight field does not generalize well to unseen poses. GNARF [2] and AvatarGen [66] use surface-based mapping that enables better generalization to novel poses. Our experimental results show that surface-based mapping may suffer from ghosting artifacts and our vertex-based representation instead avoids such a problem. More importantly, in contrast to these methods [42, 2, 66] that all leverage the tri-plane-based representation, we propose to learn features directly on the human vertices and learn a generative vertex-based radiance field. This enables generalization to novel poses and shapes. A more recent work [19] also enables control over the shape and pose using a compositional neural radiance field. In contrast to these methods, our method additionally enables part-level editing for free.

Articulated Human Representations: The performance of 3D-aware human generative models is closely related to the choice of articulated human representation. A line of works represents the human body as a mesh controlled by a set of pose and shape parameters [33]. While presenting compact representations, they neglect human appearance and cloth. More recently, coordinate-based MLPs have advanced many fields, including reconstruction [37, 46, 43] and novel view synthesis [39, 60]. There are many attempts to leverage implicit neural representations for human reconstruction [49, 48, 59, 69, 7, 61, 35, 30, 27, 41]. Among these methods, earlier methods [49, 35] model humans in observation space, increasing the variation to be learned. Modeling humans in a global canonical space avoids this problem, yet learning a backward weight field struggles to generalize to novel poses [48, 59, 69], and learning a forward weight field is computationally expensive due to the iterative root finding [7, 27]. Our proposed vertex-based representation is more closely related to deterministic surface-based warping methods [61, 30]. In contrast to all aforementioned representations focusing on reconstruction, our representation enables generalization to novel poses without expensive computation and does not suffer from ghosting artifacts, hence well-suited for generative models.

3. Method

Given a set of 2D images, we aim to learn a generative human radiance field with explicit control over cam-

era pose, human pose and shape, as well as enabling control over part-level shape and appearance. In the following, we first briefly review generative radiance fields of non-articulated objects in Section 3.1. Next, we introduce our vertex-based generator in Section 3.2 and its versatile controllability in Section 3.3. Finally, we describe the training of the model in Section 3.4.

3.1. Generative Radiance Fields

Generator: Generative Radiance Fields [53, 5] learn a conditional neural radiance fields from a set of 2D image collections using adversarial training. Conditioned on a noise $\mathbf{z} \in \mathbb{R}^M$ sampled from a Gaussian distribution, the generator G_θ generates a radiance field that maps a 3D point \mathbf{x} to a corresponding RGB value \mathbf{c} and density value σ :

$$G_\theta : \mathbb{R}^{L_x} \times \mathbb{R}^M \rightarrow \mathbb{R}^3 \times \mathbb{R}^+ \quad (\gamma(\mathbf{x}), \mathbf{z}) \mapsto (\mathbf{c}, \sigma) \quad (1)$$

where $\gamma(\mathbf{x})$ is the positional encoding operation that maps the given input to higher-dimensional features and L_x denotes the dimension of $\gamma(\mathbf{x})$. Note that we do not consider the viewing direction here for simplicity as articulated humans are usually Lambertian [4, 54].

Volume Rendering: The generative radiance field is rendered to a 2D image via the volume rendering operation. Given a camera pose ξ sampled from a camera pose distribution p_ξ , the generator is queried on $H \times W$ rays with N sampling points on each ray. Let $\{\mathbf{c}_i, \sigma_i\}_{i=1, \dots, N}$ denote the queried color and density values of one ray, the pixel color \mathbf{c} at the corresponding ray is obtained via volume rendering:

$$\begin{aligned} \pi : (\mathbb{R}^3 \times \mathbb{R}^+)^N &\rightarrow \mathbb{R}^3 & \{(\mathbf{c}_i, \sigma_i)\} &\mapsto \mathbf{c} \\ \mathbf{c} = \sum_{i=1}^N T_i \alpha_i \mathbf{c}_i & \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j) & \alpha_i &= 1 - \exp(-\sigma_i \delta_i) \end{aligned}$$

where T_i and α_i denote the transmittance and alpha value at a sample point \mathbf{x}_i , δ_i denotes the distance interval between two adjacent sample points.

Challenges: Despite achieving superior performance on non-articulated objects, e.g., human faces and cars, it is non-trivial to extend generative radiance fields to articulated objects. The naïve approach of generating articulated humans in the observation space (e.g., EG3D) leads to unsatisfying performance and a lack of controllability. Existing attempts for generative human radiance fields map a 3D point in the observation space to a pose-agnostic canonical space, either achieved by learning a blend weight field [42] or using a fixed surface-based warping [2, 66]. Learning generative blend weight fields from single-view 2D images using adversarial loss is quite challenging. Further, the learned backward weight fields struggle to generalize to novel poses

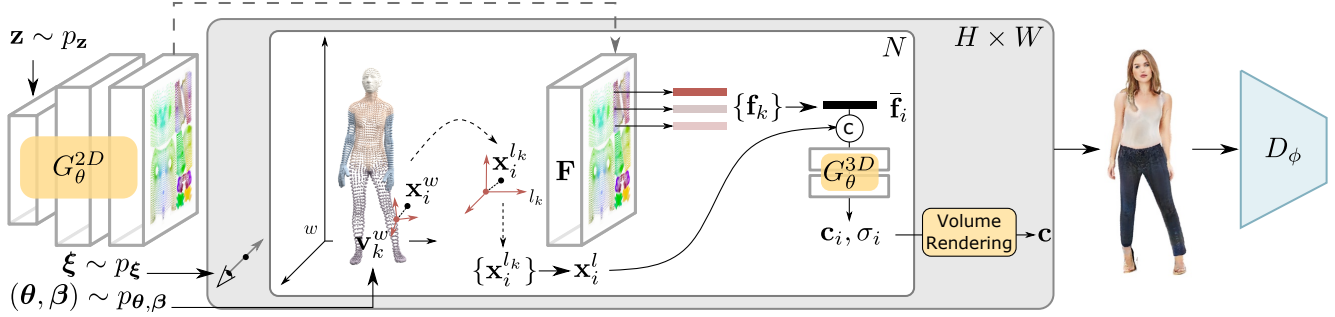


Figure 2: **Generative VeRi3D.** We model the generative human radiance fields based on a set of vertices and their corresponding feature vectors. Given a sampled camera pose ξ and sampled human pose and shape (θ, β) , we sample a 3D point \mathbf{x}_i^w along a ray in the observation space and retrieve its K nearest neighbors $\{\mathbf{v}_k^w\}_{k=1, \dots, K}$. For each neighboring vertex \mathbf{v}_k^w , we transform \mathbf{x}_i^w to the local coordinate system of \mathbf{v}_k^w and obtain $\mathbf{x}_i^{l_k}$. Next, we combine the weighted feature vectors $\{\mathbf{f}_k\}$ and the local coordinate $\{\mathbf{x}_i^{l_k}\}$ to generate the color \mathbf{c} and density σ at this point. This mapping is performed on N sample points on each ray and $H \times W$ pixels to obtain an output image, whereas the feature vectors are generated once using a 2D convolutional network.

out of the dataset distribution. On the other hand, surface-based warping methods avoid the challenging task of learning the blend weight fields and enable better generalization. However, it may suffer from inaccurate surface mapping and thus yielding ghosting artifacts. We propose a simple approach that maps a 3D point to the local coordinate systems of its nearest vertices. Our formulation avoids learning the weight field and do not suffer from ghosting artifacts. Moreover, our formulation allows for generalizing to novel shapes and poses, as well as editing local parts.

3.2. VeRi3D as Generative Human Fields

Fig. 2 gives an overview of our method. We train our generative VeRi3D on 2D image collections and their corresponding camera parameters and human pose distributions using adversarial training. We refer to the space under a sampled SMPL pose as the *observation space*. Given a set of vertices of SMPL, we first define a local coordinate system for each vertex conditioned on the pose. Next, we transform a sampled 3D point in the world coordinate of the observation space to the local coordinate systems of its neighboring SMPL vertices, representing each 3D point using the vertex feature and the local coordinates. This formulation enables generalization to novel shapes and poses. The vertex features are generated using a 2D convolutional neural network based on a fixed UV mapping.

More formally, let \mathbf{z} denote a random noise vector, (θ, β) a sampled SMPL pose and shape, and ξ a sampled camera pose. We define the world coordinate system of the observation space as w . Further, let $\mathbf{V}^w \in \mathbb{R}^{M \times 3}$ denote M vertices of the SMPL model in the observation space where each vertex \mathbf{v}^w is associated with a feature vector \mathbf{f} . Let $T_a^b \in SE(3)$ denote a rigid transformation, i.e., $\mathbf{x}^b = T_a^b \mathbf{x}^a$.

Given a sampled 3D point \mathbf{x}^w in the observation space,

we first retrieve its K nearest neighbors from the vertices $\{\mathbf{v}_k^w \in \mathcal{N}(\mathbf{x}^w)\}_{k=1, \dots, K}$ and transform the 3D points to the local coordinate systems defined on $\{\mathbf{v}_k^w\}$. Next, \mathbf{x}^w is mapped to a color and a density value based on its local coordinates determined by $\{\mathbf{v}_k^w\}$ and their corresponding feature vectors $\{\mathbf{f}_k\}$.

Vertex-Based Local Coordinate Systems: For each vertex \mathbf{v}_k^w , we define a local coordinate system $T_{l_k}^w$ with its origin at the vertex, allowing for transforming a 3D point in the world coordinate system to the local coordinate system via:

$$\mathbf{x}^{l_k} = (T_{l_k}^w)^{-1} \mathbf{x}^w = T_{l_k}^w \mathbf{x}^w, \quad (2)$$

where l_k denotes the local coordinate system of the k th neighbor. Note that \mathbf{x}^{l_k} should be pose-agnostic, e.g., the color and density of a point should be invariant to the given pose. Thus, we predefine a fixed local coordinate system $T_{l_k}^w$ under the T-pose and transform it to $T_{l_k}^w$ via linear blend skinning. Concretely, the translation of $T_{l_k}^w$ is $\bar{\mathbf{v}}_k^w$ defined under the T-pose. The rotation is determined by the normal vector of $\bar{\mathbf{v}}_k^w$ and a fixed up direction $[0, 0, 1]^T$. The vertex normal is calculated as the weighted average of the normals of all the faces that the vertex is part of. This vertex normal is taken as the z -axis of the rotation matrix. Next, we define the rest two axes using a fixed vector in the world coordinate system $[0, 0, 1]^T$. The x -axis is the cross product of the z -axis and $[0, 0, 1]^T$, and the y -axis is the cross product of the other two axes.

Next, we transform $T_{l_k}^w$ to $T_{l_k}^w$ based on linear blend skinning using the skinning weight of the vertex:

$$T_{l_k}^w = \left(\sum_{p=1}^P W(\bar{\mathbf{v}}_k^w)_p G_p(\theta) \right) T_{l_k}^w \quad (3)$$

where P is the number of parts defined in SMPL, $G_p(\theta) \in SE(3)$ is a transformation matrix of the p th part given pose θ , and $W(\bar{\mathbf{v}}_k^w)_p$ is the blend weight of the p th part.

Vertex-based Radiance Fields: Our vertex-based radiance fields map each point \mathbf{x}^w to a color and a density value based on its K nearest vertices. This design is similar to Point-NeRF [60] but we focus on learning generative radiance fields instead of overfitting to a single scene or object. Concretely, after obtaining local coordinates $\{\mathbf{x}^{l_k}\}_{k=1\dots K}$, we query the color and density of the point based on the feature vectors $\{\mathbf{f}_k\}$ and $\{\mathbf{x}^{l_k}\}$. Here, $\{\mathbf{f}_k\}$ provides the information to distinguish different local coordinate systems, and $\{\mathbf{x}^{l_k}\}$ further distinguishes points within the same local coordinate systems. This brings several advantages: 1) Our local coordinates are attached to the SMPL vertices, thus naturally enabling pose and shape control; 2) This representation enables part-level control as we directly learn a set of features associated with the SMPL vertices.

More specifically, we simply accumulate the feature vectors of the K nearest vertices as follows:

$$\bar{\mathbf{f}} = \sum_k \frac{p_k}{\sum p_k} \mathbf{f}_k \quad (4)$$

where p_k is the inverse-distance weight widely used in scattered data interpolation. The intuition behind this is that closer vertices should contribute more to the output.

Regarding the local coordinates $\{\mathbf{x}^{l_k}\}$, we average their directions and summarize the statistics of their norms to be order-invariant wrt. the K nearest neighbors. The statistics contain minimum, maximum, mean, and variance of K norms. We concatenate the 3-dimensional averaged directions and 4-dimensional norm statistics to a vector $\mathbf{x}^l \in \mathbb{R}^7$ to indicate the local coordinate information.

We use an MLP to map the accumulated feature vector $\bar{\mathbf{f}}$ and the local coordinate \mathbf{x}^l to the color and RGB values:

$$G_\theta^{3D} : \mathbb{R}^{L_F} \times \mathbb{R}^{L_{x^l}} \rightarrow \mathbb{R}^3 \times \mathbb{R}^+ \quad (\bar{\mathbf{f}}, \gamma(\mathbf{x}^l)) \mapsto (\mathbf{c}, \sigma) \quad (5)$$

where $\gamma(\cdot)$ denotes positional encoding.

Vertex Feature Generator: We leverage the UV mapping to project SMPL vertices $\mathbf{V} \in \mathbb{R}^{M \times 3}$ to a 2D map, enabling us to learn features on SMPL vertices leveraging a 2D CNN. In practice, we use a style-based 2D CNN to learn to generate a 2D feature map $\mathbf{F} \in \mathbb{R}^{H_F \times W_F \times L_F}$ given a sampled noise vector \mathbf{z} :

$$G_\theta^{2D} : \mathbb{R}^M \rightarrow \mathbb{R}^{H_F \times W_F \times L_F} \quad \mathbf{z} \mapsto \mathbf{F} \quad (6)$$

Note that the projected SMPL vertices may not fall on integer locations. We use bilinear interpolation to query a feature $\mathbf{f} \in \mathbb{R}^{L_F}$ from \mathbf{F} for each vertex.

3.3. VeRi3D for Controllable Image Synthesis

Pose Control: Our model enables controlling poses using the pose parameter θ . Thanks to our vertex-based formulation, our method inherits the generalization ability of SMPL, thus enabling generalizing to poses out of the training images.

Shape Control: Since our representation is shape-agnostic, we can control the body shape of generated humans by manipulating the body mesh surface. By controlling the shape parameter β of SMPL, i.e., the coefficients of 10 orthogonal basis obtained by principal component analysis (PCA), we can obtain different body shapes.

Part-level Control: As we directly learn a set of features on SMPL vertices, we can control the part-level appearance and geometry by changing the vertices' features. These vertices \mathbf{V} are segmented to different body parts based on the SMPL skinning weights. For vertices of each part, we apply PCA to the vertices features and get the principal component. Specifically, we randomly sample 1k latent code to compute PCA component and select the top 30 components with the largest eigenvalues. This enables part control by changing the principal coefficients.

3.4. Discriminator and GAN Training

Discriminator: Following [25], we adopt a normal 2D convolutional neural network as the discriminator. Our discriminator is not conditioned on the camera pose or the human pose.

GAN Training: Given a latent code \mathbf{z} , a camera pose ξ , human pose and shape (θ, β) , and a real image \mathbf{I} sampled from the real-data distribution $p_{\mathcal{D}}$, we train our VeRi3D using non-saturated GAN loss with R1 regularization [36].

$$\begin{aligned} \mathcal{L} = & \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0,1), \xi \sim p_\xi, (\theta, \beta) \sim p_{\theta, \beta}} [f(D_\phi(G_\theta(\mathbf{z}, \xi, \theta, \beta)))] \\ & + \mathbb{E}_{\mathbf{I} \sim p_{\mathcal{D}}} [f(-D_\phi(\mathbf{I})) + \lambda \|\nabla D_\phi(\mathbf{I})\|^2] \end{aligned}$$

4. 3D-Aware Human Synthesis

Datasets: We conduct experiments on three datasets: DeepFashion [65], AIST++ [28] and Surreal [57]. DeepFashion is a real-world fashion dataset. We use the data provided by [19]. The dataset filters out images with partial observations and inaccurate SMPL estimations, resulting in 8K images for training. AIST++ is a video dataset capturing 30 performers in dance motion. Each frame is annotated with the ground truth camera and fitted SMPL body model. We use a segmentation model [8] to remove backgrounds and get 90k images in total. Surreal contains 6M images of synthetic humans created by the SMPL body model in various poses rendered in indoor scenes.



Figure 3: **Qualitative Comparison on DeepFashion.** We show each sample from three viewing directions.



Figure 4: **Qualitative Comparison on AIST++.** We provide the corresponding real image as a reference for each input pose. Note that “Tri. + Surf.” sometimes fails to follow the input pose (1st and 2nd) and may have ghosting artifacts (3rd).

Evaluation Metrics: We adopt Fréchet Inception Distance (FID) [18] to measure image quality. Following ENARF, we use the Percentage of Correct Keypoints (PCKh@0.5) to evaluate the correctness of generated poses.

Baselines: We compare our method against a state-of-the-art 3D-aware image synthesis model, EG3D [3], and a state-of-the-art 3D-aware human image synthesis model, ENARF [42]. As EG3D does not model human articulation, we provide the SMPL pose and shape parameters as input conditions to EG3D. We also compare our method with a surface-based baseline. Following AvatarGen [66], we map

a 3D point in the observation space to a canonical space based on SMPL skinning weights. We use a tri-plane to represent the canonical space, where we sample feature for each point in canonical space to predict density and color. We refer to this method as “Tri. + Surf.”. We implement this baseline in our framework to ensure a fair comparison.

Implementation Details: We use a feature dimension of $256 \times 256 \times 64$. The local coordinate vector \mathbf{x}^l has a length of 7 which includes direction and distance statistics. Similar to [19], we only sample points around the human body and increase the sampling of side view for the Deep-

Method	Deepfashion			Surreal			AIST++		
	Res.	FID10k ↓	PCKh@0.5 ↑	Res.	FID10k ↓	PCKh@0.5 ↑	Res.	FID10k ↓	PCKh@0.5 ↑
EG3D [3]	512	27.2	-	128	19.5	-	256	38.2	-
ENARF [42]	128	60.3	0.991	128	21.3	0.855	128	95.7	0.916
Tri. + Surf.	512	23.7	0.936	128	8.0	0.883	256	32.4	0.949
Ours	512	21.4	0.988	128	6.7	0.984	256	32.3	0.973

Table 1: **Quantitative Results** on Deepfashion, Surreal, and AIST++ dataset.



Figure 5: **Pose Control**. We compare pose controllability with ENARF by applying poses of the AIST++ dataset to the generative radiance fields trained on DeepFashion.



Figure 6: **Part Control**. Our method enables controlling the head, upper body, and lower body independently, with cloth color, style, hairstyle, and facial changes.

Fashion dataset. We use the Adam optimizer, with a batch size of 8 for training. The learning rate of the generator is 0.0001. The learning rate of the discriminator is 0.0002. For the DeepFashion dataset, we train 620k iterations. For the Surreal dataset, we train 640k iterations. For the AIST++ dataset, we train 400k iterations.

4.1. 3D-Aware Human Image Synthesis

Table 1 shows our quantitative comparison to the baselines. Note that our method achieves superior performance in terms of FID. While EG3D leads to reasonable FID scores, it is not controllable in terms of the human pose and thus we do not evaluate its PCKh@0.5. Our method outperforms ENARF which maps a 3D point to the canonical space based on the learned weight fields. ENARF performs better on datasets with smaller pose variation, i.e., Surreal and DeepFashion, yet the performance degenerates on the AIST++ dataset which captures dancing videos. Regard-

ing the surface mapping-based approach “Tri. + Surf.”, its FID scores are comparable to our proposed method, but the generated images are of lower accuracy in terms of pose correctness. Compared with ENARF, our method is superior on Surreal and AIST++ and similar on the DeepFashion dataset in terms of PCKh@0.5. Note that both ENARF and our method can maintain the input pose accurately on DeepFashion as its poses are of lower variation.

We show the qualitative comparison to ENARF and Tri. + Surf. in Fig. 3 on the DeepFashion dataset. Aligned with the quantitative results, our method produces images with higher visual fidelity compared to ENARF. Despite achieving similar visual quality on DeepFashion, the Tri. + Surf. baseline sometimes leads to inaccurate poses and ghosting artifacts when training on datasets with larger pose variations. This can be seen in Fig. 4, which shows our qualitative comparison to Tri. + Surf. on the AIST++ dataset.



Figure 7: **Shape Control.** We manipulate the coefficient of the first principle component for the left four images and the second principle component for the right ones.

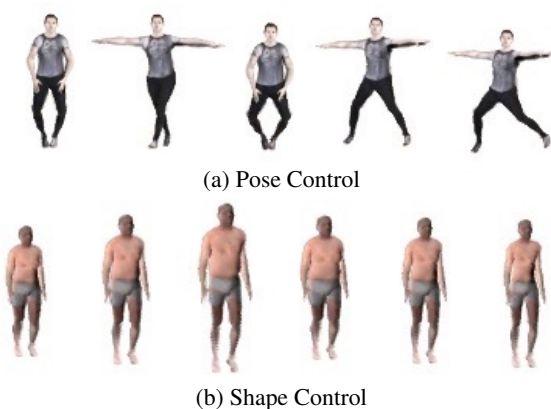


Figure 8: **Qualitative Results** on Surreal Dataset.

4.2. Controllability

Pose Control: Our vertex-based representation enables generalization to out-of-distribution novel poses. To evaluate our generalization ability, we apply the poses of the AIST++ dataset to the human radiance fields learned on the DeepFashion dataset in Fig. 5 and compare with ENARF. As ENARF uses a learned blend weight field, it produces artifacts at the joints given such out-of-distribution poses. We further show pose control results of our model trained on Surreal and AIST++ in Fig. 8a and Fig. 9a, respectively. The results suggest that our method achieves convincing pose control results on various datasets.

Shape Control: Fig. 7 demonstrate our shape controllability, achieved by manipulating the SMPL shape parameters θ . Note that our method generalizes well to the unseen type of shapes. Similarly, we show shape control results of our models trained on Surreal and AIST++ in Fig. 8b and Fig. 9b.

Part Control: We demonstrate our part editing results in

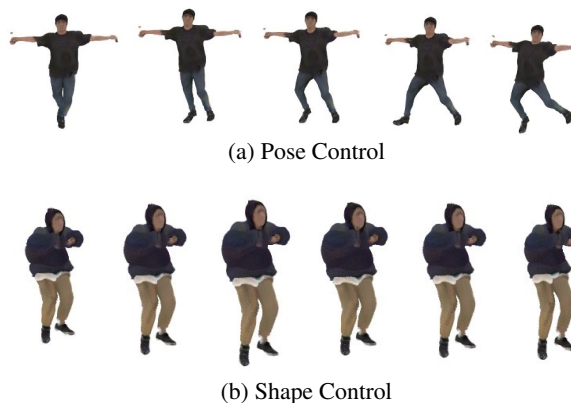


Figure 9: **Qualitative Results** on AIST++ Dataset.

Fig. 6. As can be seen, we can edit the head, the upper body, and the lower body independently, with clothes color, clothes style, hairstyle, and facial changes. Please refer to our supplementary for more results.

5. Reconstruction

Datasets: We conduct experiments on ZJU-MoCap dataset [50] to evaluate the effectiveness of our vertex-based representation when supervised by a reconstruction loss. ZJU-MoCap records human motion using 21 synchronized cameras and provides accurate human poses. We use one camera for training and the remaining ones for testing novel view synthesis. We divide the motion sequence frames into “training pose” and “novel pose”, and use the former to train the model and the latter to test the performance under the novel pose.

Evaluation Metrics: We evaluate PSNR, SSIM [58], and LPIPS [68] to compare the quality of synthesized images under novel view and novel pose settings.

Baselines: We compare our method with a state-of-the-




Figure 10: **Qualitative Comparison** of the reconstruction experiment. Surface-based mapping suffers from ghosting artifacts, i.e., an extra leg appears in the empty space.

art reconstruction method, HumanNeRF [59], using learned weight field, to demonstrate our method’s representation capability and generalization ability. We use their released model to test for novel poses. We also compare with the “Tri. + Surf.” baseline in this reconstruction experiment.

5.1. Reconstruction Results

Table 2 shows our quantitative comparison of the reconstruction experiments. We compute averaged results over six subjects following the evaluation setting of HumanNeRF. Note that our method achieves competitive performance to the state-of-the-art human reconstruction method HumanNeRF. The tri-plane-based baseline achieves similar results on PSNR and SSIM, but degenerates significantly on LPIPS. This may be caused by ghosting artifacts as shown in Fig. 10, i.e., an extra leg appears in the empty space. In contrast, by transforming each 3D point to the local coordinate systems of its nearest neighbors, our method does not suffer from such a problem.

5.2. Ablation Study

Number of Neighboring Vertices: Table 3 shows our quantitative comparison of the reconstruction experiments. We compare different numbers of neighboring vertices. The performance improves as the number of reference neighbors increases to K=3 but then saturates with K=5. Therefore, we adopt K=3 as the default setting in our main experiments.

Method	Novel View			Novel Pose		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
HumanNeRF	30.24	0.9680	31.73	29.68	0.9664	32.87
Tri. + Surf.	29.21	0.9631	36.50	29.96	0.9678	31.44
Ours	30.11	0.9684	30.24	30.25	0.9698	28.53

Table 2: **Reconstruction Results** on ZJU-MoCap.

Method	Novel View			Novel Pose		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
K=1	30.03	0.9683	30.59	30.13	0.9694	28.84
K=3	30.11	0.9684	30.24	30.25	0.9698	28.53
K=5	30.05	0.9684	30.09	30.19	0.9697	28.41
w/o direction	29.72	0.9652	33.43	29.97	0.9673	30.92

Table 3: **Ablation Study** on ZJU-MoCap.

Direction in Local Coordinate System: In the last row of Table 3, we remove the direction of sampled points in the local coordinate system. It means the sampled points with the same distances to the neighboring vertex have the same representation, thus leading to worse performance. This experimental result verifies our design choice.

6. Conclusions

We present VeRi3D, a generative vertex-based human radiance field that offers versatile controllability over camera pose, human pose, shape, and part-level editing. We demonstrate that our vertex-based human representation achieves superior performance on unconditional human image synthesis and is also competitive in the reconstruction setting.

Limitations: Our method relies on accurate pose distribution to be able to respect the input pose. With mismatched pose distribution, e.g., bent legs estimated as straight ones, our method needs to generate bent legs given straight legs as input, thus not respecting the input pose. This could be addressed by jointly refining the pose distribution in future work. Further, our method may have artifacts in less observed regions, e.g., there are sometimes artifacts in the back regions for forward-facing datasets. This may be resolved by adding regularizations to unobservable regions.

Acknowledgements

This work was supported by NSFC under grant U21B2004, 62202418, the Fundamental Research Funds for the Central Universities under Grant No. 226-2022-00145, and the Zhejiang University Education Foundation Qizhen Scholar Foundation. We thank Sheng Miao for proofreading the manuscript and Yue Wang for helpful discussions.

References

- [1] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: shape completion and animation of people. *ACM Trans. on Graphics*, 24(3):408–416, 2005. 2
- [2] Alexander W. Bergman, Petr Kellnhofer, Yifan Wang, Eric R. Chan, David B. Lindell, and Gordon Wetzstein. Generative neural articulated radiance fields. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2, 3
- [3] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J. Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3d generative adversarial networks. In *CVPR*, 2022. 6, 7
- [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, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3
- [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 *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 3
- [6] Xu Chen, Tianjian Jiang, Jie Song, Jinlong Yang, Michael J. Black, Andreas Geiger, and Otmar Hilliges. gdna: Towards generative detailed neural avatars. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [7] Xu Chen, Yufeng Zheng, Michael Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 3
- [8] PaddlePaddle Contributors. Paddleseg, end-to-end image segmentation kit based on paddlepaddle. <https://github.com/PaddlePaddle/PaddleSeg>, 2019. 5
- [9] Enric Corona, Albert Pumarola, Guillem Alenyà, Gerard Pons-Moll, and Francesc Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [10] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv.org*, 2107.02791, 2021. 2
- [11] Haoye Dong, Xiaodan Liang, Ke Gong, Hanjiang Lai, Jia Zhu, and Jian Yin. Soft-gated warping-gan for pose-guided person image synthesis. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems (NIPS)*, 2018. 2
- [12] Patrick Esser, Ekaterina Sutter, and Björn Ommer. A variational u-net for conditional appearance and shape generation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [13] Anna Frühstück, Krishna Kumar Singh, Eli Shechtman, Niloy J. Mitra, Peter Wonka, and Jingwan Lu. Insetgan for full-body image generation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [14] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NIPS)*, 2014. 2
- [15] Artur Grigorev, Karim Iskakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. Stylepeople: A generative model of fullbody human avatars. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [16] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *Proc. of the International Conf. on Learning Representations (ICLR)*, 2022. 2
- [17] Philipp Henzler, Niloy J Mitra, , and Tobias Ritschel. Escaping plato’s cave: 3d shape from adversarial rendering. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 2
- [18] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 6
- [19] Fangzhou Hong, Zhaoxi Chen, Yushi LAN, Liang Pan, and Ziwei Liu. EVA3d: Compositional 3d human generation from 2d image collections. In *International Conference on Learning Representations*, 2023. 2, 3, 5, 6
- [20] Kyungmin Jo, Gyumin Shim, Sanghun Jung, Soyoung Yang, and Jaegul Choo. Cg-nerf: Conditional generative neural radiance fields. *arXiv.org*, 2021. 2
- [21] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2018. 2
- [22] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. 2020. 2
- [25] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. 5
- [26] Christoph Lassner, Gerard Pons-Moll, and Peter V. Gehler. A generative model of people in clothing. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2017. 2
- [27] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhöfer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner.

- Tava: Template-free animatable volumetric actors. In *European Conference on Computer Vision*, pages 419–436. Springer, 2022. 3
- [28] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. Learn to dance with aist++: Music conditioned 3d dance generation, 2021. 5
- [29] Yiyi Liao, Katja Schwarz, Lars Mescheder, and Andreas Geiger. Towards unsupervised learning of generative models for 3d controllable image synthesis. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [30] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Trans. Graph.(ACM SIGGRAPH Asia)*, 2021. 3
- [31] Wen Liu, Wenhao Luo, Lin Ma, Zhixin Piao, Min Jie, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [32] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. on Graphics*, 2015. 2
- [33] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. 34(6):248:1–248:16, Oct. 2015. 3
- [34] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Advances in Neural Information Processing Systems (NIPS)*, pages 406–416, 2017. 2
- [35] Qianli Ma, Jinlong Yang, , Siyu Tang, and Michael J Black. The power of points for modeling humans in clothing. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 2, 3
- [36] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *Proc. of the International Conf. on Machine Learning (ICML)*, 2018. 2, 5
- [37] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [38] 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 *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 2
- [39] 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 *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 3
- [40] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 2
- [41] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5762–5772, 2021. 3
- [42] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Unsupervised learning of efficient geometry-aware neural articulated representations. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Proc. of the European Conf. on Computer Vision (ECCV)*, 2022. 2, 3, 6, 7
- [43] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 3
- [44] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Park, and Ira Kemelmacher. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [45] Xingang Pan, Xudong Xu, Chen Change Loy, Christian Theobalt, and Bo Dai. A shading-guided generative implicit model for shape-accurate 3d-aware image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [46] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [47] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [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 *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 3
- [49] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [50] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 8
- [51] Kripasindhu Sarkar, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. Humangan: A generative model of human images. In *2021 International Conference on 3D Vision (3DV)*, pages 258–267. IEEE, 2021. 2
- [52] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In Munkhtset-

- seg Nandigjavi, Niloy J. Mitra, and Aaron Hertzmann, editors, *ACM Trans. on Graphics*, 2022. 2
- [53] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3
- [54] Katja Schwarz, Axel Sauer, Michael Niemeyer, Yiyi Liao, and Andreas Geiger. Voxgraf: Fast 3d-aware image synthesis with sparse voxel grids. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2, 3
- [55] Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, and Nicu Sebe. Deformable gans for pose-based human image generation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [56] Hao Tang, Song Bai, Li Zhang, Philip H. S. Torr, and Nicu Sebe. Xinggan for person image generation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 2
- [57] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [58] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004. 8
- [59] Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3, 9
- [60] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3, 5
- [61] Tianhan Xu, Yasuhiro Fujita, and Eiichi Matsumoto. Surface-aligned neural radiance fields for controllable 3d human synthesis. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [62] Xudong Xu, Xingang Pan, Dahua Lin, and Bo Dai. Generative occupancy fields for 3d surface-aware image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [63] Yinghao Xu, Sida Peng, Ceyuan Yang, Yujun Shen, and Bolei Zhou. 3d-aware image synthesis via learning structural and textural representations. *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [64] Zhuoqian Yang, Shikai Li, Wayne Wu, and Bo Dai. 3dhuman: Towards photo-realistic 3d-aware human image generation. *arXiv preprint*, arXiv:2212.07378, 2022. 2
- [65] Liu Z, Luo P, and Qiu S. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [66] 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.org*, 2022. 2, 3, 6
- [67] Jichao Zhang, Enver Sangineto, Hao Tang, Aliaksandr Siarohin, Zhun Zhong, Nicu Sebe, and Wei Wang. 3d-aware semantic-guided generative model for human synthesis. *Proc. of the European Conf. on Computer Vision (ECCV)*, 2022. 2, 3
- [68] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 8
- [69] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Efficiently generated human radiance field from sparse inputs. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [70] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. CIPS-3D: A 3D-Aware Generator of GANs Based on Conditionally-Independent Pixel Synthesis. *arXiv.org*, 2021. 2
- [71] Zhen Zhu, Tengpeng Huang, Mengde Xu, Baoguang Shi, Wenqing Cheng, and Xiang Bai. Progressive and aligned pose attention transfer for person image generation. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 44(8):4306–4320, 2022. 2