

# SCULPT: Shape-Conditioned Unpaired Learning of Pose-dependent Clothed and Textured Human Meshes

Soubhik Sanyal<sup>1</sup> Partha Ghosh<sup>1</sup> Jinlong Yang<sup>1\*</sup>  
 Michael J. Black<sup>1</sup> Justus Thies<sup>1,2</sup> Timo Bolkart<sup>1\*</sup>

<sup>1</sup>Max Planck Institute for Intelligent Systems <sup>2</sup>Technical University of Darmstadt

{soubhik.sanyal, partha.ghosh, jinlong.yang, black, justus.thies, timo.bolkart}@tue.mpg.de




Figure 1. **SCULPT** is a generative model of geometry and appearance of clothed human meshes. The generated textured clothing mesh can be readily inserted into 3D scenes. In the above figure, we generated the clothed humans and placed them on a 3D floor. The scene has a single camera with global directional light settings.

## Abstract

We present **SCULPT**, a novel 3D generative model for clothed and textured 3D meshes of humans. Specifically, we devise a deep neural network that learns to represent the geometry and appearance distribution of clothed human bodies. Training such a model is challenging, as datasets of textured 3D meshes for humans are limited in size and accessibility. Our key observation is that there exist medium-sized 3D scan datasets like CAPE, as well as large-scale 2D image datasets of clothed humans and multiple appearances can be mapped to a single geometry. To effectively learn from the two data modalities, we propose an unpaired learning procedure for pose-dependent clothed and textured human meshes. Specifically, we learn a pose-dependent geometry space from 3D scan data. We represent this as per vertex displacements w.r.t. the SMPL model. Next, we train a geometry conditioned texture generator in an unsupervised way using the 2D image data. We use intermediate activations of the learned geometry model to condition our texture generator. To alleviate entanglement between pose and clothing type, and pose and clothing

appearance, we condition both the texture and geometry generators with attribute labels such as clothing types for the geometry, and clothing colors for the texture generator. We automatically generated these conditioning labels for the 2D images based on the visual question answering model BLIP and CLIP. We validate our method on the **SCULPT** dataset, and compare to state-of-the-art 3D generative models for clothed human bodies. Our code and data can be found at <https://sculpt.is.tue.mpg.de>.

## 1. Introduction

Generating 3D virtual humans that can be articulated with realistically deforming clothing, is a key challenge of content creation for games and movies. Virtual assistants in augmented and virtual reality could be enriched by an automatically generated 3D human appearance. Furthermore, synthetic humans could also play an important role in data generation to comply with privacy and data protection laws.

Recently, we have seen immense progress in the synthesis of virtual 3D humans [4, 16–18, 31, 49]. However, they are almost exclusively based on implicit representations, such as

\*Now at Google.

neural radiance fields [17, 31]. These implicit representations are incompatible with classical rendering frameworks, making it challenging to integrate them into existing applications. To address this, we propose *SCULPT* a generative model of explicit 3D geometry (meshes) and the appearance (texture maps) of clothed humans (Fig. 1). There are several challenges that we need to address to build such a model. To naively train the 3D generative model, a large dataset of 3D scanned humans would be required containing a diverse set of people in a variety of different poses, wearing different outfits. Not only such data is not publicly available, but it would also be very expensive to collect. However, we observed that there are different requirements for learning the geometry and texture of a generative model. For geometry learning, there are medium-scale datasets of 3D-scanned humans that are publicly accessible. Additionally, 2D images depicting humans in various outfits are abundantly available, which can assist in learning appearance or texture. It is also noted that clothing items, such as t-shirts, despite having identical geometries, can vary significantly in colors and patterns. This diversity allows for a wide range of colors and textures to be associated with similar geometrical structures. Based on this observation, we propose to leverage medium-scale datasets of 3D scans to learn the geometry distribution, while large-scale 2D image datasets are used to learn the appearance model. Based on these data sources, and leveraging the foundational SMPL body model [26], we design our explicit generative model.

Specifically, we modify the StyleGAN architecture to output (i) pose-dependent geometry in terms of displacement maps, and (ii) geometry-dependent appearance in terms of a color texture, in the texture space (UV-space) of the SMPL template. Our geometry model is trained using the CAPE dataset [28], which contains pose annotations as well as registered SMPL meshes. We use the trained geometry model to condition our texture generator, which is trained on a large collection of 2D fashion images. Note that we train the texture model in an unsupervised way, only relying on adversarial losses, thereby avoiding the requirement of 3D data paired with 2D images. The process of geometry conditioning plays a crucial role in maintaining coherence between appearance and geometry. In essence, this process ensures that the visual and geometric attributes of the clothed human harmoniously conform to each other. To mitigate dependence of generated clothing type or its color on the body pose, we condition both the texture and the geometry generators with attribute labels. This approach reduces the entanglement between pose and clothing attributes, resulting in more realistic and accurate generation. We automatically generate these labels by using the visual question-answering model BLIP [25], and CLIP [37].

During inference, our model has the ability to generate a diverse set of clothed, textured 3D humans that can be controlled by various parameters such as clothing type, and clothing appearance such as color. This level of control over the generated output is a significant advantage of our approach,

as it allows for great flexibility in generating realistically clothed and textured 3D humans. The coarse clothing color can be controlled from text-based inputs, which enables users to specify the desired color for a given piece of clothing without requiring detailed knowledge of 3D modeling or design. Moreover, since our generative model produces 3D meshes it is compatible with existing graphics and game engines, allowing seamless integration into a range of applications. In summary, the contribution of this work is a novel hybrid learning strategy for unsupervised and unpaired learning of a generative model of 3D virtual humans. It is enabled by coupling the appearance and geometry network via language-driven attribute labels.

## 2. Related work

**Generative 3D modeling:** 3D-aware generative models have been receiving a tremendous amount of focus from the computer vision community in recent years [5, 6, 14, 32, 43, 44]. This has resulted in different representations of 3D objects and scenes like implicit functions and different rendering techniques such as volumetric rendering. In the following paragraph, we provide a brief overview of these concepts as they help contextualize our work better.

Chan *et al.* [5] propose a new architecture for generative models, where they build the generator with SIREN networks [45] and perform volumetric rendering to generate 2D images. The SIREN network models the inherent geometry of an image using implicit representations which are rendered to a 2D image via volumetric rendering inspired by NeRF [30]. Or-EI *et al.* [32] combine an SDF-based 3D representation with a style-based 2D generator. The SDF-based representation helps in achieving geometry details but produces low-resolution image features, which are then fed to a 2D-based style generator that produces high-resolution images. Chan *et al.* [6] take an interesting direction by decoupling feature generation and neural rendering with the help of a tri-plane representation. This enables them to leverage the power of superior image generation modules like StyleGAN2 [20] to generate high-quality 2D results. Even though these works have advanced the state-of-the-art in 3D aware image synthesis for relatively simple objects like cars or faces, their generated geometry is not of high quality. Moreover, they are not directly usable for complex 3D deformable and articulable models like clothed humans.

This leads to a more focused set of generative models [4, 7, 8, 16–18, 28, 31, 33, 49] which are concentrated on generative 3D humans. Among these 3D generative models, a stream of work [7, 8, 28, 33] uses 3D data as supervision. Yet they only model geometry and not the texture or appearance of the clothed humans. This is caused by the limited size of 3D datasets containing the large variation of clothing textures along with their corresponding geometry. To get around this limitation, another stream of work uses 2D images as supervision. The concurrent works by Bergman *et al.* [4], Zhang *et al.* [49] and Noguchi *et al.* [31] extend Chan *et al.* [6] for human bodies where they learn




Figure 2. **Overview:** SCULPT consists of two StyleGAN-based generators for geometry ( $\mathcal{G}_{geo}$ ) and appearance ( $\mathcal{G}_{tex}$ ), both acting in the UV space of the SMPL body model. The geometry network  $\mathcal{G}_{geo}$  outputs pose-dependent displacement maps that are added to the SMPL template mesh and is trained using 3D scan data. Based on this model, the appearance generator  $\mathcal{G}_{tex}$  is trained in an unsupervised way using adversarial losses computed on rendered images of the generated synthetic human. It is conditioned on intermediate features of the geometry network. Besides the noise code, both generator networks receive additional attributes for appearance ( $\mathbf{c}_t$ ) and clothing type ( $\mathbf{c}_g$ ) as input. This enhances the connection between appearance and geometry, and it offers a user-friendly control over the generation.

a neural radiance field for a canonical pose that is later reposed. The radiance fields are learned by using a similar concept of tri-planes proposed by Chan et al. [6]. The final human image produced by the network of [4] is blurry however, and the geometry is very smooth and lacks deformations relating to clothing. Similarly, Noguchi et al. [31] also produce blurry results and undesired artifacts. Grigorev et al. [16] propose to generate neural textures using 2D generative models. The neural texture is superimposed on minimally clothed SMPL-X [35] meshes and rendered. The rendered images are then passed via a neural renderer to generate realistic-looking 2D person images. Their method does not alter the geometry and is incompatible with existing infrastructure, e.g., game engines. Hong et al. [17] divide the whole body into parts and learn local radiance fields, which are then integrated to get the final rendered image. They also suffer from undesirable artifacts and lack any control over geometry and texture. In contrast, we learn geometry on top of a fixed topology mesh (SMPL, [26]) and the corresponding texture map for the underlying mesh. This is done via learning from unpaired data using a novel training approach utilizing both 3D and 2D data. Displacements on top of SMPL are sufficient for many kinds of clothes. Moreover, this template is readily compatible with any existing 3D rendering engine. Additionally, unlike prior art, we provide explicit controls over the clothing type and color.

**Generative 2D humans:** If we set aside the goal of having an explicit 3D mesh as output, some 2D methods become relevant. Here, we study a few generative models that can generate

high-quality 2D images for clothed humans [12, 42]. Fu et al. [12] propose a curated dataset of fashion images and train StyleGAN [20, 21]. In this way, they produce high-quality 2D images of clothed humans but lack controllability on the generated pose and other aspects, such as global orientation, clothing type, etc. Sarkar et al. [42] propose a new architecture for controllable human generation. To provide additional pose controllability over the generated humans one can use any of the 2D reposing algorithms [2, 3, 9–11, 15, 24, 27, 36, 36, 40, 40, 41, 46, 48] proposed in the recent years. In summary, these methods propose various architectures and/or training procedures that enable them in synthesizing a human in a new pose given the image of the person in a different pose. Although these methods generate high-quality images, they have no notion of the underlying geometry. Therefore they cannot be used in classical 3D graphics pipelines.

### 3. Method

SCULPT is a generative model that takes a geometry code  $\mathbf{z}_g \sim \mathcal{N}(\mathbf{0}, I^{512 \times 512})$ , a texture code  $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, I^{512 \times 512})$ , body pose  $\theta \in \mathbb{R}^{69}$ , clothing geometry type  $\mathbf{c}_g \in \{0, 1\}^6$ , and clothing texture description  $\mathbf{c}_t \in \mathbb{R}^{512}$  as input, and it generates a clothed 3D body mesh  $\mathcal{M} := \{V, \mathcal{C}\}$  with a texture image  $\mathcal{I}_{tex} \in \mathbb{R}^{256 \times 256 \times 3}$  (Fig. 2). Here,  $V \in \mathbb{R}^{6890 \times 3}$  represents a set of 6890 3D vertices, and  $\mathcal{C}$  is a fixed set of triangles represented as a 3-tuple of vertex indices in SMPL [26] mesh topology. Formally, SCULPT is defined as:

$$\mathcal{M}, \mathcal{I}_{tex} = SCULPT(\mathbf{z}_g, \mathbf{z}_t, \theta, \mathbf{c}_g, \mathbf{c}_t). \quad (1)$$

*SCULPT* models clothing geometry as vertex displacements from the minimally clothed SMPL body in the canonical pose. To account for pose articulation effects, the clothing generator is conditioned on pose. The texture generator in turn takes the features from the geometry generator as a conditioning signal. Therefore, the output of *SCULPT* can be fully articulated with SMPL’s pose control, and the mesh is readily usable in existing graphics applications.

### 3.1. Clothing representation

*SCULPT* adapts the SMPL [26] model formulation to clothed bodies with additional parameters  $\mathbf{z}_g$  and  $\mathbf{c}_g$  as:

$$T_P(\beta, \theta, \mathbf{z}_g, \mathbf{c}_g) = T_S(\beta) + B_P(\theta) + V_{geo}(\mathbf{z}_g, \mathbf{c}_g, \theta), \quad (2)$$

where  $T_S(\beta) \in \mathbb{R}^{6890 \times 3}$  denotes the SMPL body in “canonical pose” for the shape parameters  $\beta \in \mathbb{R}^{10}$ .  $B_P(\theta) : \mathbb{R}^{3k} \rightarrow \mathbb{R}^{6890 \times 3}$  are the SMPL pose corrective blend shapes, and  $V_{geo}$  are the pose-dependent clothing vertex displacements. The vertex displacements  $V_{geo}$  are obtained by sampling the UV displacement map output by the clothing geometry generator  $\mathcal{G}_{geo}$  (see Sec. 3.2) at fixed UV coordinates for every vertex.

As the clothing geometry is defined as offsets from the SMPL body in the canonical pose, it can be reposed as:

$$\mathcal{M}(\beta, \theta, \mathbf{z}_g, \mathbf{c}_g) = W(T_P(\beta, \theta, \mathbf{z}_g, \mathbf{c}_g), \mathbf{j}(\beta), \theta), \quad (3)$$

where  $W(T_P, \mathbf{j}, \theta)$  is SMPL’s blend skinning function, which rotates the vertices of  $T_P$  around  $k = 23$  joints  $\mathbf{j} \in \mathbb{R}^{3k}$ . The joint locations  $\mathbf{j}$  are defined as a function of the body shape.

### 3.2. Clothing geometry generator

Given a random geometry code  $\mathbf{z}_g \sim \mathcal{N}(\mathbf{0}, I^{512 \times 512})$ , a one-hot clothing type vector  $\mathbf{c}_g \in \{0, 1\}^6$ , and the SMPL body pose  $\theta$ , the clothing geometry generator outputs a UV displacement map  $UV_{geo} \in \mathbb{R}^{256 \times 256 \times 3}$  (see left column of Fig. 2). Formally, the generator is defined as:

$$UV_{geo} = \mathcal{G}_{geo}(\mathbf{z}_g | \mathbf{c}_g, \theta). \quad (4)$$

Following CAPE ([28]),  $\mathbf{c}_g$  are categorical classifications of the clothing type like “short sleeve T-shirt/short trouser”, “short sleeve T-shirt/long trouser”, “long sleeve T-shirt/long trouser”, “long sleeve T-shirt/short trouser”, “shirt/long trouser”, “shirt/short trouser”, provided with the CAPE dataset, and which are represented as one-hot vectors. Intuitively,  $\mathbf{c}_g$  controls the clothing category, while  $\mathbf{z}_g$  and  $\theta$  model clothing variations and pose-dependent deformations within the particular clothing category, respectively (Fig. 3). To sample clothed 3D body meshes from *SCULPT*, we sample vertex displacements  $V_{geo}$  from the generated displacement map  $UV_{geo}$  and evaluate Eq. (3).

The generator  $\mathcal{G}_{geo}$  follows a StyleGAN3 [21] architecture, which is trained with an adversarial loss from a global discriminator. We compute the displacement maps from the




Figure 3. **Pose control:** The figure presents three meshes of an individual. The initial two meshes depict the individual in different types of clothing (long-long and short-short) but maintaining the same pose, while the third mesh illustrates the individual in a different pose, wearing the same type of clothing as depicted in the first mesh. The pose-dependent clothing deformations are visible in the zoomed-in images on the side, which are produced by the geometry generator. It is observed that identical poses result in similar deformations, as indicated by the color-coded bounding boxes, while a change in pose leads to distinct clothing deformations in the geometry.

scan registrations provided by the CAPE dataset, which are treated as real samples for the discriminator. For the fake examples, we combine random noise vectors for  $\mathbf{z}_g$  and  $\mathbf{c}_g$  with randomly sampled CAPE data poses. During training, the output of the generator is masked by the segmentation mask of the SMPL UV map and passed to the discriminator as fake samples. Following standard conditional GAN training, the discriminator is also given  $\mathbf{c}_g$  and  $\theta$  as inputs.

### 3.3. Texture generator

Given a random texture code  $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, I^{512 \times 512})$ , our texture generator  $\mathcal{G}_{tex}$  generates a UV texture image. To control the generated texture  $\mathcal{G}_{tex}$ , it is conditioned on clothing texture descriptors  $\mathbf{c}_t$ . Additionally we condition the texture generator with the categorical clothing type  $\mathbf{c}_g$  and intermediate features of the clothing geometry generator  $\mathcal{G}_{geo}$ . Formally, the texture generator is:

$$\mathcal{I}_{tex} = \mathcal{G}_{tex}(\mathbf{z}_t | \mathbf{c}_g, \mathbf{c}_t, \mathcal{G}_{geo}). \quad (5)$$

We train the texture generator from a collection of 2D fashion images, obtained from fashion websites. For each training image, as detailed in Sec. 3.4, we automatically extract clothing color descriptors  $\mathbf{c}_t$  and clothing types  $\mathbf{c}_g$  with the help of a visual question answering system – BLIP ([25]) and CLIP ([37]). Realistic clothing appearance, however, consists of more than coarse clothing colors. Namely, it can contain varying color patterns. We capture such variations in our generator using an additional latent vector  $\mathbf{z}_t$ . The coarse clothing geometry type  $\mathbf{c}_g$  and the clothing color descriptor  $\mathbf{c}_t$  are concatenated, and provided as input condition to  $\mathcal{G}_{tex}$ . This is sufficient to generate good visual quality texture but the generated texture is inconsistent with the generated geometry. For instance, a short-sleeved shirt covers a smaller skin region of the arms as compared to a long-sleeved one, therefore the texture generator





Figure 4. **Pose control:** Varying pose while keeping other factors fixed. Each row contains two different identities, each in two poses. Texture and geometry meshes are shown side-by-side.




Figure 5. **Cloth-type control:** Varying  $c_g$  while keeping other factors fixed. Each row contains two separate identities. For each identity, we show two different clothing types consecutively. Texture and geometry meshes are shown side-by-side.

must be conditioned using the geometry information. Although the clothing category  $c_g$  loosely correlates with the generated texture and geometry, it is not enough. For instance, without explicit information on the clothing part boundaries (e.g. boundary between shirt and trousers, or boundary between cloth and skin) the texture network is unable to generate a clothing appearance that conforms to the clothing geometry. To account for the correlation of generated geometry and texture, we condition  $\mathcal{G}_{tex}$  on intermediate features of  $\mathcal{G}_{geo}$ . Specifically, during a forward pass,  $\mathcal{G}_{geo}$  takes the batch of  $\mathbf{z}_g, c_g$ , and  $\theta$ , and it generates features at different synthesis blocks. Since  $\mathcal{G}_{tex}$  and  $\mathcal{G}_{geo}$  share StyleGAN’s model architecture, we progressively add at each level, excluding the mapping network, the feature blocks of  $\mathcal{G}_{geo}$  to the feature blocks of  $\mathcal{G}_{tex}$ . See the left half of Fig. 2 for a visual representation of this technique. This effectively passes the signals from the geometry network to the texture generator.

Finally, we combine the generated  $\mathcal{I}_{tex}$  and  $\mathcal{M}$ , and render the textured mesh with a differentiable texture renderer [38]. We use a global and a patch-based discriminator simultaneously to train  $\mathcal{G}_{tex}$ . While the global discriminator acts on the whole rendered clothed human image, the patch-based discriminator is a local discriminator that acts on the random  $64 \times 64$  patches extracted from the image. Following the conditional GAN training scheme,  $c_g$  and  $c_t$  are provided as inputs to both the discriminators. Empirically, we find that the local discriminator helps improve the image quality of the body parts, whereas the global one ensures consistency of the entire structure.

### 3.4. Obtaining clothing texture descriptions

We utilize the language-to-image models CLIP [37] and BLIP [25] to automatically label the fashion images with clothing type  $c_g$  and clothing color descriptors  $c_t$ . Specifically, to compute,  $c_g$  we pass one fashion image through CLIP [37] and compute its image features. Then we use augmented prompts for the categorical clothing types, e.g., “the person is

wearing a t-shirt”, “the person is wearing a shirt”, “the person is wearing long pants” etc. as input to CLIP’s text feature extractor. Finally, we use the scores provided by CLIP for each text input and the corresponding fashion image and select the features corresponding to the text prompt with the highest score as its clothing type descriptor  $c_g$ . To compute the color descriptor  $c_t$ , we pass the fashion images to BLIP [25] and query its visual question answering (VQA) model with questions such as, “What is the color of the upper body clothing of the person wearing in the image?”. BLIP then outputs a textual description of the clothing color. We augment this text with the following sentence, “The color of the upper body clothing is [BLIP output text] and the color of the pants is [BLIP output text]”. This is then passed as text input to CLIP to get the text-based feature, which then serves as  $c_t$ . At inference time, we replace the BLIP generated output with a fixed textual description of the clothing colors, and use this as input to CLIP, to compute  $c_t$ .

### 3.5. Training and dataset details

SCULPT is implemented in PyTorch [34] and optimized with ADAM [23] with a learning rate of 0.001. The geometry, texture generators and the discriminators follow the StyleGAN3-t [21] architecture. The fake examples for the texture discriminators of  $\mathcal{G}_{tex}$  are obtained by rendering the generated clothed body mesh  $\mathcal{M}$  with the generated texture map  $\mathcal{I}_{tex}$  using PyTorch3D [38]. All generators and discriminators are trained using a non-saturating GAN loss using R1 regularisation [20, 29] and the adaptive augmentation technique from StyleGAN-Ada [19]. The training process follows two stages: (1) The geometry generator is trained on the CAPE dataset [28] until the FID converges. The CAPE dataset [28] consists of SMPL registrations to the scans of 41 subjects wearing different types of clothing. The dataset comes with six different clothing type variations, namely “short-short”, “short-long”, “long-long”, “long-short”, “shirt-long” and “short-short”. Here, the term “long-short” rep-



Figure 6. **Cloth-color control:**  $c_t$  is varied while keeping other factors fixed. Each row consists of two different clothing geometries and differently colored garments for that geometry.

resents that the person is wearing a round neck shirt/t-shirt with long sleeves and short pants. Similarly, the term “short-long” represents that the person is wearing a round-neck shirt/t-shirt with short sleeves and long pants. The same terminology is followed for the other labels. We encode these six clothing types as a 6-dimensional one-hot vector. Each registered mesh is unposed, i.e., effects of pose articulation and translation are removed, and the vertex offsets from the minimally-clothed SMPL body are represented in the UV space as a displacement map. In total, we compute 63069 displacement maps from these registered meshes. (2) Then the geometry model is kept fixed and used for training the texture generator until the FID converges. The texture generator is trained on a curated dataset of fashion images obtained from the catalog images uploaded in the website of Zalando [1]. We collected 16362 fashion images, normalized the images (human centered in the middle) and removed the background. We run MODNet [22] on the aligned images to get the segmentation masks for the foreground body and use the pose regressors provided by ICON [47] to estimate the SMPL pose and shape of the bodies. We compute the clothing type and clothing color information utilizing CLIP and BLIP as described in Sec. 3.4. More details on the dataset statistics can be found in the Sup. Mat. PDF file.

## 4. Experiments

We conduct a three-part evaluation of *SCULPT*. Initially, we assess its controllability properties. Subsequently, we compare it with state-of-the-art methods. Finally, we execute a detailed series of ablation studies.

**Controllability of *SCULPT*:** In Fig. 4, we show generation results by varying the body pose  $\theta$ , while keeping all other control parameters of *SCULPT* fixed. This leads to pose-dependent deformation of the clothing geometry, which is visible in the side-by-side comparison of textured and textureless mesh for



Figure 7. **Cloth-texture fine control:** Varying  $z_{tex}$  while keeping other factors fixed. Each row consists of two different clothing geometries, each with textures generated for the same color condition but with different  $z_{tex}$ .



Figure 8. **Viewpoint changes:** Rotating the textured mesh.

each pose-pair of one identity. Additionally, we show the pose-dependent deformations only in the geometry in Fig. 3. Notice that the appearance of the clothes changes in tandem with the geometry. In Fig. 5, we vary the clothing type,  $c_g$  keeping all other parameters fixed. As intended, the clothing type changes from long sleeves and long pants to short sleeves and long pants when  $c_g$  is changed accordingly (Row 1, identity 2). Fig. 6 shows that for each clothing geometry, we can generate different colored garments by varying  $c_t$ , fixing the other factors.

Our model offers further fine-grained control over the appearance and allows for fine changes in the texture of the clothing, i.e., it can generate different shades of the same coarse color, and different patterns on the same t-shirt/shirt, etc. This is achieved by varying  $z_{tex}$  in Fig. 7. It is worth noting that varying  $z_{tex}$  changes the texture patterns for identity 1 in row 1 of Fig. 7 and it generates different shades of the same color. In contrast to 2D generative modeling, our output is a textured mesh that can be rendered under arbitrary viewpoints, see Fig. 8.

**Comparisons with SOTA:** We compare *SCULPT* with the state-of-the-art (SOTA) methods, EG3D [6] and EVA3D [17] quantitatively in Tab. 1 and qualitatively in Fig. 9. We also provide qualitative comparisons of *SCULPT* with two additional SOTA methods, namely GET3D [14] and StylePeople [16] in Fig. 10. We compare to these methods only qualitatively because an official implementation of training code of StylePeople [16] is not publicly available and GET3D [14] requires images of people in a canonical pose to train, which is not available for our dataset.

We find that although *SCULPT* and EG3D both methods generate high-quality images. However, the underlying geometry, generated by EG3D, is of low quality. We hypothesize

Method	FID ↓	KID ↓	Precision/Recall ↑
EG3D	7.38	0.0036	0.79/0.14
EVA3D	44.11	0.0387	0.30/0.03
SCULPT	9.85	0.0063	0.53/0.22

Table 1. **Quantitative comparison:** We evaluate our model using the standard FID, KID, and precision and recall ([39]) metrics against the most recently proposed similar methods. Our rendering quality is comparable to the state-of-the-art methods.

that the highly articulated human body is significantly more complex to model as compared to human and animal faces, which have less articulation. This renders the training of EG3D difficult, leading to undesirable solutions. Moreover, EG3D does not provide any of the highly desirable control parameters that our model provides.

Recently, EVA3D [17] has been proposed to add controllable articulation to 3D aware generative models of the human body. For a fair comparison, we used EVA3D’s publicly available implementation and DeepFashion experiment parameters and trained it with our 256X256 texture data. Note that the original EVA3D model was trained on a different dataset with 512X512 resolution, leading to different results. By doing so, EVA3D improves the quality of the generated geometry as compared to EG3D. But their generated texture is of lower quality than SCULPT, which is evident from Tab. 1 and Fig. 9. Furthermore, as the generated geometry is represented as an implicit function, using it in existing graphics engines requires converting it into meshes, which is time-consuming and often does not preserve the rendering quality.

While GET3D’s geometric representation can model loose-fitting clothes such as skirts, it lacks articulation and pose control (it generates the body in a canonical pose). While SCULPT is more limited in topology due to the explicit representation, it enables control over complex, articulated human figures, and generates better textures (Fig. 10). Adopting different explicit topologies for different clothing types could greatly enhance the types of clothes SCULPT can model. This is left for future work.

In contrast to StylePeople, which employs 2D neural rendering with generated neural textures, our approach estimates accurate clothing geometry that is compatible with standard 3D renderers. Contrary to StylePeople, SCULPT incorporates a geometry branch. The feature outputs from each Style-Block in the geometry generator are added with those from the corresponding ones in the texture generator, resulting in a texture that is consistent with the generated geometry. This fundamental difference in our architecture allows us to generate better-quality renders compared to StylePeople as can be seen in Fig. 10. In summary, SCULPT outperforms the SOTA methods in terms of geometric representation (StylePeople), generated geometry quality (EG3D, EVA3D, StylePeople), articulation (GET3D, EG3D), texture quality (all), and final human image production (all) (Fig. 9 and Fig. 10).




Figure 9. **Qualitative comparison:** We compare with EG3D (left) and EVA3D (middle). Our rendered humans (right) have comparable quality with EG3D whereas our geometry surpasses both.



Figure 10. **Additional qualitative comparisons.** From left to right: StylePeople [16], GET3D [14], SCULPT (each with two results). Images are taken directly from the respective publications.

**Ablation experiments:** To better understand the contribution of different components in *SCULPT*, we perform ablation experiments as shown in Tab. 2. These experiments involve altering the choice of discriminator combinations, patch sizes, and the conditioning of the texture network by the intermediate activations of the geometry network. We train the geometry network conditioned texture network with only one discriminator at a time in cases (b), (c), and (d). However, we observe that the global discriminator alone is not capable of generating sharp results (b). On the other hand, the patch discriminators work relatively well in improving local parts of the body but lack global correspondence. Among the local discriminators, we compare the performance of two patch sizes,  $64 \times 64$  (d) and  $32 \times 32$  (c), and find that  $64 \times 64$  performs better. We hypothesize that as the granularity of the local discriminator increases, its field of view diminishes, leading to a higher likelihood of encountering white backgrounds devoid of human elements, particularly in comparison to a  $64 \times 64$  patch. This phenomenon potentially impairs the model’s overall performance. Ganokratana et al. [13] observed a similar pattern, where their discriminator using  $64 \times 64$  patches outperformed the one with  $32 \times 32$  patches. We then add the global discriminator along with the local discriminator in cases (e) and



<b>Ablated</b>	(a)	(b)	(c)	(d)	(e)	(f)
<b>FID</b> ↓	28.2	31.6	24.5	16.1	19.2	9.85

Table 2. **Ablation study:** (a) full model without geometry conditioning; (b) full model trained with only global discriminator; (c) full model trained with only local discriminator of patch size  $32 \times 32$ ; (d) full model trained with only local discriminator of patch size  $64 \times 64$ ; (e) full model trained with both global and local discriminator of patch size  $32 \times 32$ ; (f) full model trained with both global and local discriminator of patch size  $64 \times 64$ .




Figure 11. **Geometry conforming with texture:** Geometric features like clothing boundaries or wrinkles in different body areas (highlighted in blue boxes) are consistent with texture.

(f), which improves the overall performance. The  $64 \times 64$  patch size also performs best in the combined discriminator strategy as can be seen by comparing (c) and (e), and (d) and (f). Finally, we build a baseline where we train the texture generator without conditioning from the geometry network with the dual discriminator strategy, as shown in (a). However, the baseline performs poorly compared to our full model (f), as the geometry and texture do not conform to each other. Figure 11 shows that the geometry network’s conditioning of the texture network allows the texture to conform to the clothing geometry, as observed in the clothing boundaries and wrinkles of the different identities.

**Limitations and discussion:** Similar to existing generative clothed human body models, *SCULPT* generates poor-quality textures in the backside region because of a dataset bias towards frontal and near-frontal views, and to “typical” fashion poses in the fashion image dataset used for training. The dataset bias can be observed from the dataset statistics provided in the **Sup. Mat. PDF file**. The quality also degrades for challenging, unseen body poses Fig. 12. The existing datasets further lack diversity in age, race, skin tone, and gender. The predominance of male examples is due to the CAPE dataset bias towards tight clothing and the prevalence of male subjects in our texture training data. To overcome the limitation in view- and pose diversity, one can train our model on multi-view datasets or videos of subjects in the same clothing but in varying poses and visible from different viewpoints. *SCULPT* sometimes generates hand textures with the same color as the clothing. The issue partly arises from about 30% of the training set’s fashion images showing hands in pockets, as observed from manual image



Figure 12. **Limitations:** The texture quality degrades for out-of-distribution poses (columns 1-3) and for the body’s back (right).

examination of a hundred images. However, the CAPE dataset used for geometry training lacks hands-in-pockets instances, leading to model ambiguity in recognizing hand positions in pockets. Augmenting the dataset with annotations specifying hand positions could mitigate this. Finally, our model also has topological limitations in modeling loosely fitting clothes such as skirts or long dresses. To handle loosely fitting clothing such as skirts, our method would require a corresponding mesh template that offers the correct topology. Given the limited range of typical clothing typologies, it could be feasible to design a few different typologies to accommodate such clothing types but this is outside of the scope of this work. Instead, we focus on developing a novel method for learning geometry and texture without paired 3D training data. We use SMPL as a template for our explicit representation (mesh) which is compatible with current graphics engines.

## 5. Conclusion

We have introduced *SCULPT*, a generative model that creates 3D virtual humans using explicit geometry (mesh) and appearance (texture maps). *SCULPT* represents clothing geometry as offsets from the SMPL body’s vertices, and includes a generative model for texture maps based on clothing type and appearance. This approach effectively combines traditional graphics elements such as meshes, forward rendering and texture maps with modern 3D-aware generative models. Our texture model is trained from unpaired 2D-3D data, making it easy to use and retrain on new data. The clothing geometry is learned using a dataset of 3D meshes. The trained model offers generative control over clothing geometry and appearance, with semantic control over clothing type and color, while still retaining SMPL’s pose articulation and body-shape variation. Compared to previous 3D-aware generative models, *SCULPT* offers greater control and produces higher quality geometry and textures.

**Acknowledgments:** The authors would like to thank Nikos and Joachim for their valuable tips on Blender renderings, Benjamin for IT support, and Tsvetelina, Taylor, and Tomasz for their support in conducting the perceptual study.

**Disclosure:** [https://files.is.tue.mpg.de/black/CoI\\_CVPR\\_2024.txt](https://files.is.tue.mpg.de/black/CoI_CVPR_2024.txt)



## References

- [1] Zalando. <https://www.zalando.de>, 2021. 6
- [2] 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 Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 40(6): 218:1–218:11, 2021. 3
- [3] Guha Balakrishnan, Amy Zhao, Adrian V. Dalca, Fredo Durand, and John Guttag. Synthesizing images of humans in unseen poses. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8340–8348, 2018. 3
- [4] Alexander W. Bergman, Petr Kellnhofer, Yifan Wang, Eric R. Chan, David B. Lindell, and Gordon Wetzstein. Generative neural articulated radiance fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 1, 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5799–5809, 2021. 2
- [6] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16102–16112, 2022. 2, 3, 6
- [7] Xu Chen, Tianjian Jiang, Jie Song, Jinlong Yang, Michael J. Black, Andreas Geiger, and Otmar Hilliges. gDNA: Towards generative detailed neural avatars. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20427–20437, 2022. 2
- [8] Enric Corona, Albert Pumarola, Guillem Alenya, Gerard Pons-Moll, and Francesc Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11875–11885, 2021. 2
- [9] Haoye Dong, Xiaodan Liang, Ke Gong, Hanjiang Lai, Jia Zhu, and Jian Yin. Soft-gated warping-GAN for pose-guided person image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 31:472–482, 2018. 3
- [10] Haoye Dong, Xiaodan Liang, Yixuan Zhang, Xujie Zhang, Xiaohui Shen, Zhenyu Xie, Bowen Wu, and Jian Yin. Fashion editing with adversarial parsing learning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8120–8128, 2020.
- [11] Patrick Esser, Ekaterina Sutter, and Björn Ommer. A variational U-Net for conditional appearance and shape generation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8857–8866, 2018. 3
- [12] 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. In *European Conference on Computer Vision (ECCV)*, pages 1–19, 2022. 3
- [13] Thittaporn Ganokratanaa, Supavadee Aramvith, and Nicu Sebe. Unsupervised anomaly detection and localization based on deep spatiotemporal translation network. *IEEE Access*, 8: 50312–50329, 2020. 7
- [14] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. GET3D: A generative model of high quality 3D textured shapes learned from images. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2, 6, 7
- [15] Artur Grigorev, Artem Sevastopolsky, Alexander Vakhitov, and Victor Lempitsky. Coordinate-based texture inpainting for pose-guided human image generation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12135–12144, 2019. 3
- [16] Artur Grigorev, Karim Isakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. StylePeople: A generative model of fullbody human avatars. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5151–5160, 2021. 1, 2, 3, 6, 7
- [17] 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 (ICLR)*, 2023. 2, 3, 6, 7
- [18] Suyi Jiang, Haoran Jiang, Ziyu Wang, Haimin Luo, Wenzheng Chen, and Lan Xu. HumanGen: Generating human radiance fields with explicit priors. *arXiv preprint arXiv:2212.05321*, 2022. 1, 2
- [19] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:12104–12114, 2020. 5
- [20] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8110–8119, 2020. 2, 3, 5
- [21] 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)*, pages 852–863, 2021. 3, 4, 5
- [22] Zhanghan Ke, Jiayu Sun, Kaican Li, Qiong Yan, and Rynson W.H. Lau. MODNet: Real-time trimap-free portrait matting via objective decomposition. In *Proceedings of the AAAI proceedings on Artificial Intelligence*, pages 1140–1147, 2022. 6
- [23] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 5
- [24] Markus Knoche, István Sáráandi, and Bastian Leibe. Reposing humans by warping 3D features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 1044–1045, 2020. 3
- [25] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International proceedings on Machine Learning (ICML)*, pages 12888–12900, 2022. 2, 4, 5
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. 2, 3, 4
- [27] Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image

- generation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 99–108, 2018. 3
- [28] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J. Black. Learning to dress 3D people in generative clothing. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6468–6477, 2020. 2, 4, 5
- [29] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for GANs do actually converge? In *International proceedings on Machine Learning (ICML)*, pages 3481–3490, 2018. 5
- [30] 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 *European Conference on Computer Vision (ECCV)*, pages 405–421. Springer, 2020. 2
- [31] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Unsupervised learning of efficient geometry-aware neural articulated representations. In *European Conference on Computer Vision (ECCV)*, pages 597–614. Springer, 2022. 1, 2, 3
- [32] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. StyleSDF: High-resolution 3D-consistent image and geometry generation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13503–13513, 2022. 2
- [33] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. NPMs: Neural parametric models for 3D deformable shapes. In *International Conference on Computer Vision (ICCV)*, pages 12695–12705, 2021. 2
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8024–8035, 2019. 5
- [35] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. 3
- [36] Albert Pumarola, Antonio Agudo, Alberto Sanfeliu, and Francesc Moreno-Noguer. Unsupervised person image synthesis in arbitrary poses. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8620–8628, 2018. 3
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International proceedings on Machine Learning (ICML)*, pages 8748–8763, 2021. 2, 4, 5
- [38] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 5
- [39] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. *Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018. 7
- [40] 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 *International Conference on Computer Vision (ICCV)*, pages 11138–11147, 2021. 3
- [41] 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. 3
- [42] Kripasindhu Sarkar, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. HumanGAN: A generative model of human images. In *International proceedings on 3D Vision (3DV)*, pages 258–267, 2021. 3
- [43] 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)*, pages 20154–20166, 2020. 2
- [44] Katja Schwarz, Axel Sauer, Michael Niemeyer, Yiyi Liao, and Andreas Geiger. VoxGRAF: Fast 3D-aware image synthesis with sparse voxel grids. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2
- [45] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 7462–7473, 2020. 2
- [46] Sijie Song, Wei Zhang, Jiaying Liu, and Tao Mei. Unsupervised person image generation with semantic parsing transformation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2357–2366, 2019. 3
- [47] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J. Black. ICON: Implicit clothed humans obtained from normals. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13296–13306, 2022. 6
- [48] Hongtao Yang, Tong Zhang, Wenbing Huang, Xuming He, and Fatih Porikli. Towards purely unsupervised disentanglement of appearance and shape for person images generation. In *Proceedings of the 1st International Workshop on Human-Centric Multimedia Analysis*, pages 33–41, 2020. 3
- [49] 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. 1, 2