

SCANimate: Weakly Supervised Learning of Skinned Clothed Avatar Networks

Shunsuke Saito^{1*} Jinlong Yang¹ Qianli Ma^{1,2} Michael J. Black¹
¹Max Planck Institute for Intelligent Systems, Tübingen, Germany ²ETH Zürich
{ssaito, jyang, qma, black}@tuebingen.mpg.de




Figure 1: **SCANimate**. Given a set of raw scans with multiple poses containing self-intersections, holes, and noise (left), SCANimate automatically aligns all scans to a canonical pose (middle) and learns a Scanimat, a fully animatable avatar that produces pose-dependent deformations and texture without garment-specific templates or mesh registration (right).

Abstract

We present *SCANimate*, an end-to-end trainable framework that takes raw 3D scans of a clothed human and turns them into an animatable avatar. These avatars are driven by pose parameters and have realistic clothing that moves and deforms naturally. *SCANimate* does not rely on a customized mesh template or surface mesh registration. We observe that fitting a parametric 3D body model, like *SMPL*, to a clothed human scan is tractable while surface registration of the body topology to the scan is often not, because clothing can deviate significantly from the body shape. We also observe that articulated transformations are invertible, resulting in geometric cycle-consistency in the posed and unposed shapes. These observations lead us to a weakly supervised learning method that aligns scans into a canonical pose by disentangling articulated deformations without template-based surface registration. Furthermore, to complete missing regions in the aligned scans while modeling pose-dependent deformations, we introduce a locally pose-aware implicit function that learns to complete and model geometry with learned pose correctives. In contrast to commonly used global pose embeddings, our local pose conditioning significantly reduces long-range spurious correlations and improves generalization to unseen poses, especially when training data is limited. Our method can be applied to pose-aware appearance modeling to generate a fully textured avatar. We demonstrate our approach on various clothing

types with different amounts of training data, outperforming existing solutions and other variants in terms of fidelity and generality in every setting. The code is available at <https://scanimate.is.tue.mpg.de>.

1. Introduction

Parametric models of 3D human bodies are widely used for the analysis and synthesis of human shape, pose, and motion. While existing models typically represent “minimally clothed” bodies [4, 26, 43, 52, 66], many applications require realistically clothed bodies. Our goal is to make it easy to produce a realistic 3D avatar of a clothed person that can be reposed and animated as easily as existing models like *SMPL* [43]. In particular, the model must support clothing that moves and deforms naturally, with detailed 3D wrinkles, and the rendering of realistically textured images.

To that end, we introduce *SCANimate* (*Skinned Clothed Avatar Networks for animation*), which creates high-quality animatable clothed humans, called *Scanimats*, from raw 3D scans. *SCANimate* has the following properties: (1) we learn an articulated clothed human model directly from raw scans, completely eliminating the need for surface registration of a custom template or synthetic clothing simulation data, (2) our parametric model retains the complex and detailed deformations of clothing present in the original scans such as wrinkles and sliding effects of garments with arbitrary topology, (3) a *Scanimat* can be animated directly using *SMPL* pose parameters, and (4) our approach predicts

*Currently at Facebook Reality Labs.

pose-dependent clothing deformations based on local pose parameters, providing generalization to unseen poses.

Recent data-driven approaches have shown promise for learning parametric models of clothed humans from real-world observations [39, 46, 54, 56]. However, these approaches typically limit the supported clothing types and topology because they require accurate surface registration of a common template mesh to 3D training scans [39, 46, 56]. Concurrent work by Ma et al. [45] learns clothing deformation without surface registration, yet it is unclear if the method works on raw scans with noise and holes. Learning from real-world observations is essentially challenging because raw 3D scans are un-ordered point clouds with missing data, changing topology, multiple clothing layers, and sliding motions between the body and garments. Although one can learn from synthetic data generated by physics-based clothing simulation [23, 25, 54], the results are less realistic, the data preparation is time consuming and non-trivial to scale to the real-world clothing.

To address these issues, SCANimate learns directly from raw scans of people in clothing. Body scanning is becoming common, and scans can be obtained from a variety of devices. Scans contain high-frequency details, capture varied clothing topology, and are inherently realistic. To make learning from scans possible, we make several contributions: canonicalization, implicit skinning fields, cycle consistency, and implicit shape learning.

Canonicalization and Implicit Skinning Fields. The first step involves transforming the raw scans to a common pose so we can learn to model pose-dependent surface deformations (e.g. bulging, stretching, wrinkling, and sliding), i.e. pose “correctives”. But we are not seeking a traditional “registration” of the scans to a common mesh topology, since this is, in general, not feasible with clothed bodies. Instead, we learn continuous functions of 3D space that allow us to transform posed scans to a canonical pose and back again.

The key idea is to build this on linear blend skinning (LBS), which traditionally defines weights on the surface of a mesh that encode how much each vertex is influenced by the rotation of a body joint. To deal with raw scans of unknown topology, we extend this notion by defining skinning weights *implicitly* everywhere in 3D space. Specifically, given a 3D location \mathbf{x} , we regress a continuous vector function g represented by a neural network, $g(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}^J$, which defines the skinning weights. An inverse LBS function uses the regressed skinning weights to “undo” the pose of the body and transforms the points into the canonical space. As this representation makes no assumptions about the topology or resolution of input scans, we can *canonicalize* arbitrary non-watertight meshes. Furthermore, we can easily generate animations of the parametric clothed avatar by applying forward LBS to the clothed body in the canonical pose with the learned pose correctives.

Cycle Consistency. Despite the desirable properties of canonicalization, learning the skinning function is ill-posed since we do not have ground truth training data that specifies the weights. To address this, we exploit two key observations. First, as demonstrated in previous work [27, 69, 72], fitting a parametric human body model such as SMPL [43] to 3D scans is more tractable than surface registration. We leverage SMPL’s skinning weights, which are defined only on the body surface, to regularize our more general skinning function. Second, the transformations between the posed space and the canonical space should be cycle-consistent. Namely, inverse LBS and forward LBS together should form an identity mapping as illustrated in Fig. 3, which provides a self-supervision signal for training the skinning function. After training the skinning function, we obtain the canonicalized scans (all in the same pose).

Learning Implicit Pose Correctives. Given the canonicalized scans, we learn a model that captures the pose-dependent deformations. However a problem remains: the original raw scans often contain holes, and so do the canonicalized scans. To deal with this and with the arbitrary topology of clothing, we use an implicit surface representation [13, 47, 53]. As multiple canonicalized scans will miss different regions, with this approach, they complement each other, while retaining details present in the original inputs. Furthermore, unlike traditional approaches [39, 46, 54, 70], where pose-dependent deformations are conditioned on entire pose parameters, we spatially filter out irrelevant pose features from the input conditions by leveraging the learned skinning weights. In this way, we effectively prune long-range spurious correlations between garment deformations and body joints, achieving plausible pose correctives for unseen poses even from a small number of training scans. The resulting learned Scanimat can be easily reposed and animated with SMPL pose parameters.

In summary, our main contributions are (1) the first end-to-end trainable framework to build a high-quality parametric clothed human model from raw scans, (2) a novel weakly-supervised formulation with geometric cycle-consistency that disentangles articulated deformations from the local pose correctives without requiring ground-truth training data, and (3) a locally pose-aware implicit surface representation that models pose-dependent clothing deformation and generalizes to unseen poses. Our results show that SCANimate is superior to existing solutions in terms of generality and accuracy. Furthermore, we perform an extensive study to evaluate the technical contributions that are critical for success. The code and example Scanimats can be found at <https://scanimate.is.tue.mpg.de>.

2. Related Work

Parametric Models for Human Bodies and Clothing. Parametric body models [4, 26, 43, 52, 66] learn statistical

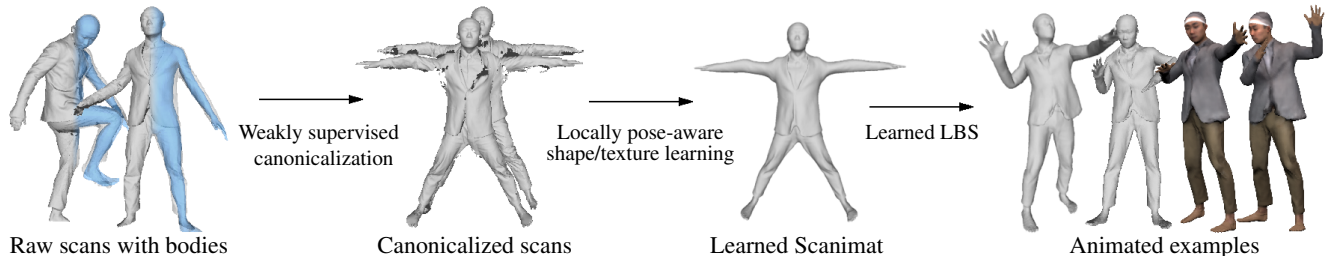


Figure 2: **Overview.** SCANimate learns a pose-aware parametric clothed human model directly from raw scans in a weakly supervised manner. The resulting Scanimats can be animated with SMPL pose parameters, producing realistic pose-dependent deformations and texture.

body shape variations and pose-dependent shape correctives that capture non-linear body deformation and compensate for linear blend skinning artifacts [3, 30, 37, 38, 41]. While these approaches achieve high-fidelity and intuitive control of human body shape and pose, they only focus on bodies without clothing. Similar ideas have been extended to model clothed bodies by introducing additional garment layers [16, 17, 23, 25, 34, 39, 67] or adding displacements or transformations to the base human body mesh [1, 2, 46, 50, 70]. These parametric clothed human models decompose garment deformations into articulated deformations and local deformations such that pose correctives only focus on non-rigid local deformations. Thus, it is essential to obtain the inverse skinning transformation [54] by using the surface registration of a well-defined template [39, 43, 46, 70, 72] or using synthetic simulation data [15, 23, 25, 54]. However, these requirements limit the applicability of the approaches to fairly simple clothing, with a fixed topology, and without complex interactions between garments and the body.

In contrast, our work uses a weakly supervised approach to build a parametric clothed human model from raw scans without the requirement of a template and surface registration. We canonicalize posed scans and learn an implicit surface with arbitrary topology [22] conditioned on pose parameters by leveraging a fitted human body model to the scan data [5, 8, 69, 71, 72]. Moon et al. [49] similarly propose a weakly supervised method for learning a fine-grained hand model from scan data by deforming a fitted base hand model [58]; the approach is non-trivial to extend to human clothing with varying topology.

The most related work to ours is Neural Articulated Shape Approximation (NASA) [18], where the composition of occupancy networks [13, 47] articulated by the fitted SMPL model are directly learned from posed scans in the same spirit as structured implicit functions [20, 21]. Concurrent work, LEAP [48], extends a similar framework to a multi-subject setting. Through an extensive study in Sec. 4.1, we find that the compositional implicit functions proposed in [18] are more prone to artifacts and less generalizable to unseen poses than our LBS-based formulation.

Pose Canonicalization via Inverse LBS. The key to successful canonicalization is learning transformations in the

form of skinning weights in a continuous space. Learning skinning weights for varied topologies has become possible using neural networks with graph convolutions [7, 42, 68]. Given a neutral-posed template, these networks predict skinning weights together with a skeleton [68] or pose-dependent deformations [7]. While they predict skinning weights on a neutral-posed template in a fully supervised manner, our problem requires learning skinning weights, not only on the surface mesh, but in both the canonical and posed space without ground-truth skinning weights.

Extending LBS skinning weights from an underlying body model to the continuous space is used in the data preparation step of ARCH [29] and LoopReg [9]. However, in these approaches, the skinning weights are uniquely determined by the underlining body and not *learnable*. We argue, and experimentally demonstrate, in Sec. 4.1 that jointly learning skinning weights leads to visually pleasing canonicalization while maximizing the reproducibility of input scans by the reconstructed parametric model. Inspired by recent unsupervised methods using cycle consistency [12, 75], we leverage geometric cycle consistency between the canonical space and posed space to learn skinning weights in a weakly supervised manner without requiring any ground-truth training data. Concurrent work, FTP [65], proposes a similar idea but is limited to body modeling; instead, we extend the traditional LBS to the entire 3D space and enable clothing surface modeling.

Reconstructing Clothed Humans. Reconstructing humans from depth maps [14, 64, 71], images [11, 32, 36], or video [33, 35] is also extensively studied. While many works focus on the minimally clothed human body [11, 24, 32, 36, 40], recent approaches show promise in reconstructing clothed human models from RGB inputs using the SMPL mesh with displacements [1, 2, 74], external garment layers [10, 31], depth maps [19, 62], voxels [63, 73], or implicit functions [29, 59, 60]. However, these approaches do not learn, or infer, pose-dependent deformation of garments, and simply apply articulated deformations to the reconstructed shapes. This results in unrealistic pose-dependent deformations that lack garment specific wrinkles. Our work differs by focusing on learning pose-dependent clothing deformation from scans.

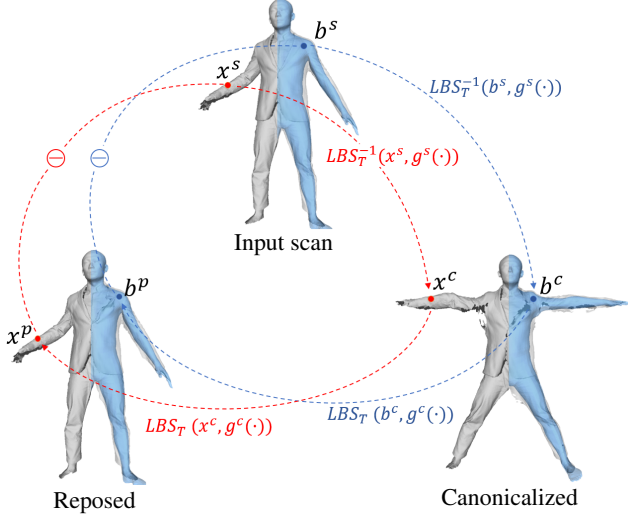


Figure 3: **Canonicalization with cycle consistency.** The geometric cycle consistency loss, with the guidance from the underlining body model, leads to successful canonicalization.

3. Method

Figure 2 shows an overview of our pipeline. The input is a set of raw 3D scans of a person in clothing, together with fitted minimally clothed body models. Here we use the SMPL model [43] fit to the scans to obtain body joints and blend skinning weights, which we exploit in learning. Given the input, we first learn bidirectional transformations between the posed space and canonical space by predicting skinning weights as a function of space coordinates (Sec. 3.1). To address the lack of ground truth correspondence of the scan data, we leverage geometric cycle consistency to learn continuous skinning functions. The raw scans are canonicalized with the learned bidirectional transformations. We further learn a locally pose-aware signed distance function, parameterized by a neural network, from canonicalized scans using implicit geometric regularization [22] (Sec. 3.2). For implementation details, including hyper parameters and network architectures, see the Sup. Mat.

3.1. Canonicalization

Instead of a traditional skinning scheme that assigns a skinning weight vector $w \in \mathbb{R}^J$, where J is the number of joints, to each point on a surface, we extend the notion of skinning using a continuous function: we train a model that takes any point in the space as input and outputs its skinning weight vector w . Figure 3 illustrates the principles. We specifically focus on points from two surfaces, the clothing surface X and the body surface B . To be more specific about the canonicalization step, let us first define the posed space and the canonical space. The posed space is defined for each scan, and the canonical space is shared across all the scans. Let $X_i^s = \{x^s \in \mathbb{R}^3\}$ be vertices on the original

scan in the posed space, where i is the frame index of the scans, and $X_i^c = \{x^c \in \mathbb{R}^3\}$ be vertices on the unposed scans in the canonical space, which are not known. Now we seek a mapping function that aligns the posed scans in a canonical pose. While the mapping function can be arbitrarily defined, we observe that this can be formulated as a composition of the *known* rigid transformations of body joints, $T^i = \{T_j^i \in SE(3), j = 1, \dots, J\}$, which come from the fitted SMPL model. More specifically, given a set of blending weights w , we define linear blend skinning (*LBS*) and inverse linear blend skinning (*LBS*⁻¹) functions as follows:

$$\begin{aligned} X_i^p &= LBS_{T_i}(X_i^c, w(X_i^c)) = \left(\sum w_j T_{i,j}\right) X_i^c \\ X_i^c &= LBS_{T_i}^{-1}(X_i^s, w(X_i^s)) = \left(\sum w_j T_{i,j}^{-1}\right) X_i^s, \end{aligned} \quad (1)$$

where $X_i^p = \{x^p \in \mathbb{R}^3\}$ are the vertices of the reposed scans and ideally should have the same value as X_i^s . The *LBS* function maps arbitrary points in the canonical space to the posed space represented by T_i and the inverse *LBS* function maps points in the posed space to the canonical space. In other words, the equations above show that given skinning weights w on vertices, we can not only apply any pose to the canonicalized shapes as in a traditional character animation pipeline [41], but also transform back the posed shapes into the canonical space.

Implicit Skinning Fields. In contrast to traditional applications, where the skinning weights for each point are predefined, either by artists or by automatic methods [6, 29, 70], skinning weights on the raw scan data are not known *a priori*. Fortunately, we can learn them in a weakly supervised manner, such that all the scans can be decomposed into articulated deformations and non-rigid deformations.

To this end, we introduce two neural networks called the *forward skinning net* and the *inverse skinning net*:

$$\begin{aligned} w(x_i^c) &= g_{\Theta_1}^c(x_i^c) : \mathbb{R}^3 \rightarrow \mathbb{R}^J \\ w(x_i^s) &= g_{\Theta_2}^s(x_i^s, z_i^s) : \mathbb{R}^3 \times \mathbb{R}^{Z_s} \rightarrow \mathbb{R}^J, \end{aligned} \quad (2)$$

where z_i^s represents a latent embedding, and Θ_1 and Θ_2 are the learnable parameters of the multilayer perceptrons (MLP), which we omit below for notational brevity. The forward skinning net predicts *LBS* skinning weights of queried 3D locations in the canonical space. Similarly, the inverse skinning net predicts skinning weights in the posed space of each training scan. Notably, this continuous representation is advantageous over other alternatives including fully connected networks and graph convolutional networks [42, 46, 57] as it does not depend on a fixed number of vertices or predefined topology. Empirically we observe that jointly learning z_i^s in an auto-decoding fashion [53] leads to superior performance compared to taking pose parameters as input; see Sup. Mat. for discussion.

By combining Eq. 1 and 2, we can compute the mappings

between the canonical and posed spaces via:

$$\begin{aligned} \mathbf{x}_i^p &= LBS_{\mathbf{T}_i}(\mathbf{x}_i^c, g^c(\mathbf{x}_i^c)) \\ \mathbf{x}_i^c &= LBS_{\mathbf{T}_i}^{-1}(\mathbf{x}_i^s, g^s(\mathbf{x}_i^s, \mathbf{z}_i^s)). \end{aligned} \quad (3)$$

Note that these functions are differentiable.

Learning Skinning. To successfully train $g^c(\cdot)$ and $g^s(\cdot)$ without ground truth weights on the scans, we leverage two key observations: (1) the regions close to the human body model are highly correlated with the nearest body parts where ground-truth skinning weights are available; (2) any points in the posed space should be mapped back to the same points after reapplying LBS to the canonicalized points. To utilize (1), we use the underlying SMPL body model’s LBS skinning weights as guidance for the canonical and posed space. More specifically, $g^s(\cdot)$ and $g^c(\cdot)$ at points on the scans are loosely guided by the nearest neighbor point on the body model and its SMPL skinning weights, propagating skinning weights from body models to the input scans.

Most importantly, observation (2) plays a central role in the success of the weakly supervised learning. It allows us to formulate cycle consistency constraints, updating both $g^c(\cdot)$ and $g^s(\cdot)$ such that wrongly associated skinning weights that break the cycle consistency are highly penalized. Our evaluation in Sec. 4.1 shows that the cycle consistency constraints are critical to decompose articulated deformations. Note that the jointly learned $g^c(\cdot)$ is used to learn and animate the pose-aware clothed human model (see Sec. 3.2).

Our final objective function is defined as:

$$\begin{aligned} E_{cano}(\Theta_1, \Theta_2, \{\mathbf{z}_i^s\}) = \\ \sum_i (\lambda_B E_B + \lambda_S E_S + E_C + E_R), \end{aligned} \quad (4)$$

where E_B and E_S are body-guided loss functions, E_C is based on cycle consistency, and E_R is a regularization term. E_B ensures $g^c(\cdot)$ and $g^p(\cdot)$ predict SMPL skinning weights on the body surface by

$$\begin{aligned} E_B = \sum_{\mathbf{b}_i^c \in \mathbf{B}_i^c} \|g^c(\mathbf{b}_i^c) - \mathbf{w}'(\mathbf{b}_i^c)\| \\ + \sum_{\mathbf{b}_i^s \in \mathbf{B}_i^s} \|g^s(\mathbf{b}_i^s, \mathbf{z}_i^s) - \mathbf{w}'(\mathbf{b}_i^s)\|, \end{aligned} \quad (5)$$

where $\mathbf{B}_i^c = \{\mathbf{b}_i^c \in \mathbb{R}^3\}$ and $\mathbf{B}_i^s = \{\mathbf{b}_i^s \in \mathbb{R}^3\}$ are vertices on the canonical and posed body surfaces, and $\mathbf{w}'(\cdot)$ are the SMPL LBS weights. See Sup. Mat. for details to obtain \mathbf{B}_i^c . Similarly, E_S is the regression loss between the predicted weights and the LBS weights on the nearest neighbor body vertex:

$$\begin{aligned} E_S = \sum_{\mathbf{x}_i^s \in \mathbf{X}_i^s} (\|g^s(\mathbf{x}_i^s, \mathbf{z}_i^s) - \mathbf{w}'(\arg\min_{\mathbf{b}_i^s \in \mathbf{B}_i^s} d(\mathbf{x}_i^s, \mathbf{b}_i^s))\| \\ + \|g^c(\mathbf{x}_i^c) - \mathbf{w}'(\arg\min_{\mathbf{b}_i^s \in \mathbf{B}_i^s} d(\mathbf{x}_i^s, \mathbf{b}_i^s))\|). \end{aligned} \quad (6)$$

Note that this nearest neighbor assignment is also used in

[29] for training data preparation. However, in Sec. 4.1, we show that this alone is prone to inaccurate assignments, causing severe artifacts.

We facilitate cycle consistency with two terms. $E_{C'}$ directly constrains the consistency of skinning weights between the canonical space and the posed space, and $E_{C''}$ facilitates cycle consistency on the vertices of the posed meshes as follows:

$$E_C = \lambda_{C'} E_{C'} + \lambda_{C''} E_{C''} \quad (7)$$

$$E_{C'} = \sum_{\mathbf{x}_i^s \in \mathbf{X}_i^s} \|g^s(\mathbf{x}_i^s, \mathbf{z}_i^s) - g^c(\mathbf{x}_i^c)\| \quad (8)$$

$$E_{C''} = \sum_{\mathbf{x}_i^s \in \mathbf{X}_i^s} \|\mathbf{x}_i^p - \mathbf{x}_i^s\|. \quad (9)$$

Notice that cycle consistency can hold only if we start from the posed space since points in the canonical space can be mapped to the same location in case of self-intersection.

Lastly, our regularization term consists of a sparsity constraint E_{Sp} , a smoothness term E_{Sm} , and a statistical regularization on the latent code E_Z as follows:

$$E_R = \lambda_{Sp} E_{Sp} + \lambda_{Sm} E_{Sm} + \lambda_Z E_Z, \quad (10)$$

$$E_{Sp} = \sum_{\mathbf{x}_i^s} |g^s(\mathbf{x}_i^s, \mathbf{z}_i^s)|^\beta \quad \beta = 0.8, \quad (11)$$

$$E_{Sm} = \sum_{e \in \mathbf{E}/\mathbf{C}} \|g^s(e_1, \mathbf{z}_i^s) - g^s(e_2, \mathbf{z}_i^s)\|, \quad (12)$$

$$E_Z = \|\mathbf{z}_i^p\|_2^2, \quad (13)$$

where $e = (e_1, e_2)$, \mathbf{E} is the set of edges on the triangulated scans and we mask out concave regions \mathbf{C} so that skinning weights are not propagated across merged body parts due to self-intersection (See Sup. Mat. for details.).

After training, we canonicalize all the scans by applying the inverse LBS transform (Eq. 3) to all vertices on the scans. By eliminating triangles with large distortion (see Sup. Mat. for details), we obtain the canonical scans used to learn a pose-aware parametric clothed human model.

3.2. Locally Pose-aware Implicit Shape Learning

Given the canonicalized partial scans together with the learned skinning weights, we learn a parametric clothed human model with pose-aware deformations. To this end, we base our shape representation on an implicit surface representation [13, 47, 53] as it supports arbitrary topology with fine details. However, real scans have holes and such partial scans cause difficulty obtaining ground truth occupancy labels since the meshes are not water-tight. To handle partial scans as input, we learn a signed distance function $f_\Phi(\mathbf{x})$ based on a multilayer perceptron (for brevity, we omit the network parameters Φ), using implicit geometric regularization (IGR) [22] by minimizing the following

objective function:

$$E_{shape}(\Phi) = \sum_i (E_{LS} + \lambda_{igr} E_{IGR} + \lambda_o E_O) \quad (14)$$

$$E_{LS} = \sum_{\mathbf{x} \in \mathbf{X}'_i} (|f(\mathbf{x})| + \|\nabla_{\mathbf{x}} f(\mathbf{x}) - \mathbf{n}(\mathbf{x})\|), \quad (15)$$

$$E_{IGR} = \mathbb{E}_{\mathbf{x}} (\|\nabla_{\mathbf{x}} f(\mathbf{x})\| - 1)^2, \quad (16)$$

$$E_O = \mathbb{E}_{\mathbf{x}} (\exp(-\alpha \cdot |f(\mathbf{x})|)) \quad \alpha \gg 1, \quad (17)$$

where E_{LS} ensures the zero level-set of the predicted SDF lies on the given points with its surface normal aligned with that of the input scans, $\mathbf{n}(\mathbf{x})$. E_{IGR} is the Eikonal regularization term that regularizes the function f to satisfy the Eikonal equation $\|\nabla_{\mathbf{x}} f(\cdot)\| = 1$. E_O regularizes off-surface SDF values from being close to the level-set surface as in [61]. Remarkably, this formulation does not require ground truth signed distance for non-surface points and naturally fills in the missing regions by leveraging the inductive bias of multilayer perceptrons as shown in [22].

To learn pose-dependent deformations of clothing, we could condition the function f with the pose features $\theta \in \mathbb{R}^{J \times 4}$ (we use quaternions as in [52]). However, the straightforward approach of concatenating the pose features with Cartesian coordinates, namely $f(\mathbf{x}, \theta)$, suffers from overfitting due to the limited pose variations in the training data and spurious correlations between joints. Since the relationship between body joints and clothing deformation tends to be non-local [67], it is also important to limit the influence of irrelevant joints to reduce spurious correlations [52]. Thus, we need an attention mechanism to associate spatial locations with only the relevant pose features. To this end, we modify the function f :

$$f(\mathbf{x}, (W \cdot g^c(\mathbf{x})) \circ \theta), W \in \mathbb{R}^{J \times J}, \quad (18)$$

where $g^c(\cdot)$ is the skinning network learned in Sec. 3.1, W is the weight map that converts skinning weights into pose attention weights, and \circ denotes element-wise product. Specifically, if we want a 3D point that is skinned to the n^{th} joint with non-zero skinning weights to pay attention to the m^{th} joint, $W_{m,n}$ and $W_{n,m}$ are set to 1, otherwise, they are set to 0. The weight map is essential because the movement of one joint will be propagated to regions associated with neighboring body joints (e.g. raising the shoulders lifts up an entire T-shirt). In this paper, we set $W_{n,m} = 1$ when n^{th} joint is within 4-ring neighbors of m^{th} joint in the kinematic tree. By reducing spurious correlations, our formulation significantly reduces over-fitting artifacts given a set of unseen poses, demonstrating better generalization ability even with a small number of input scans (see Sec. 4.1).

4. Experimental Results

Dataset and metric. For evaluation and comparison with baseline methods, we use the CAPE dataset [46], which

includes raw 3D scan sequences and SMPL model fits.

We evaluate generalization to unseen poses with both pose interpolation (denoted as Int. in tables) and extrapolation tasks (denoted as Ex. in tables). The motion sequences are randomly split into training (80%) and test (20%) sets, where the test sequences are used to evaluate extrapolation. For the training sequences, we choose every 10th frame starting from the first frame as training scans and every 10th frame with 5 frame strides from the training sequences for the interpolation evaluation. We perform Marching Cubes [44] to the predicted implicit surface in canonical space as in Eq 18 and then pose it by forward LBS in Eq. 1 to get the resulting meshes. For quantitative evaluation, we use scan-to-mesh distance D_{s2m} (cm) and surface normal consistency D_n , where a nearest neighbor vertex on the resulting meshes is used to compute the average L2 norm.

In addition, we conduct a perceptual study to assess the plausibility score, P , of generated garment shapes and deformations. Workers on Amazon Mechanical Turk (AMT) are given a pair of side-by-side images or videos showing a rendered result from our approach and another approach; the left-right order of the results is randomized. The task is to choose the result with the most realistic clothing. We continue this N times and compute the probability of the other approach being favored $P = M/N$, where M is how many times the users chose the other method over ours. In other words, we set our approach as baseline with a constant score $P = 0.5$; for other approaches, if $P < 0.5$, ours achieves higher fidelity. The perceptual score for image and video pairs is denoted as P_i and P_v , respectively. While P_i focuses on the plausibility of static clothing, P_v reveals the temporal consistency and realism of pose-dependent clothing deformations. Note that we provide only the perceptual scores for the extrapolation task as numerical evaluation is difficult due to the stochasticity of clothing deformations.

4.1. Evaluation

Canonicalization. The goal of canonicalization is to disentangle articulated deformations from other non-rigid deformations for effective shape learning. We choose two baseline approaches to replace our canonicalization module. The first, as used by [29], copies skinning weights on the clothed scans from the nearest neighbor body vertex. The other approach is based on weighted correspondences by interpolating skinning weights from the k -nearest neighbors (we use $k = 6$) in the spirit of [70]. This reduces the impact of a wrong clothing-body association that limits the performance of single nearest neighbor assignment.

Figure 4 shows that the two baseline methods break the cycle consistency with wrong associations of the skinning weights, resulting in noticeable artifacts. The inaccurate canonicalization results are propagated to the parametric model learning, substantially degrading the quality of reconstructed avatars as shown in Tab. 1. Our approach with

Table 1: Quantitative comparison of canonicalization. As the perceptual score is pair-wise and compared against ours, we put 0.5 for the proposed approach throughout the tables. D_{s2m} is in centimeters throughout the tables.

	Ours	NN [29]	KNN	
Int.	$D_{s2m} \downarrow$	0.570	1.25	1.25
	$D_n \downarrow$	0.253	0.301	0.299
	$P_i \uparrow$	0.5	0.374	0.396
	$P_v \uparrow$	0.5	0.435	0.431
Ex.	$P_i \uparrow$	0.5	0.262	0.312
	$P_v \uparrow$	0.5	0.392	0.449




Figure 4: **Qualitative comparison on canonicalization.** Top: canonicalization results. Bottom: reposed canonicalization results. Compared with our method, the baseline methods suffer from severe artifacts.

cycle consistency successfully normalizes the input scans into a canonical pose while retaining coherent geometric details, enabling the parametric modeling of clothed avatars. **Locally Pose-aware Shape Learning.** We evaluate our local pose representation using the learned skinning weights for pose-dependent shape learning and compare against commonly used global pose conditioning [18, 39, 46, 54, 70]. To this end, we replace the second input of Eq. 18 with the global pose parameter, θ , as a baseline. To assess the generalization ability, both models are trained on 100%, 50%, 10% and 5% of the original training set.

Table 2 shows that our local pose conditioning achieves better reconstruction accuracy and fidelity for both interpolation and extrapolation. Note that the performance of global pose conditioning drastically degrades when the training data is reduced to less than 10%, suffering from severe overfitting. In contrast, our approach keeps roughly equivalent reconstruction accuracy even when only 5% of the original training data is used, exhibiting few noticeable artifacts (see Fig. 5).

Comparison with SoTA. We compare the proposed method with two state-of-the-art methods that also learn an articulated parametric human model with pose correctives from real world scans [18, 46]. CAPE [46] learns pose-dependent deformations on a fixed mesh topology using

Table 2: Quantitative evaluation of the importance of locality in the pose conditioning on different sizes of training data.

Train size (%)	100	50	10	5	
Local pose conditioning (Ours)					
Int.	$D_{s2m} \downarrow$	0.570	0.663	0.699	0.732
	$D_n \downarrow$	0.253	0.253	0.261	0.268
	$P_i \uparrow$	0.5	0.476	0.466	0.398
	$P_v \uparrow$	0.5	0.453	0.435	0.425
Ex.	$P_i \uparrow$	0.5	0.429	0.359	0.359
	$P_v \uparrow$	0.5	0.408	0.408	0.343
Global pose conditioning					
Int.	$D_{s2m} \downarrow$	0.768	0.786	1.54	2.38
	$D_n \downarrow$	0.253	0.256	0.293	0.354
	$P_i \uparrow$	0.424	0.393	0.350	0.252
	$P_v \uparrow$	0.468	0.457	0.363	0.301
Ex.	$P_i \uparrow$	0.417	0.401	0.291	0.192
	$P_v \uparrow$	0.436	0.382	0.311	0.203

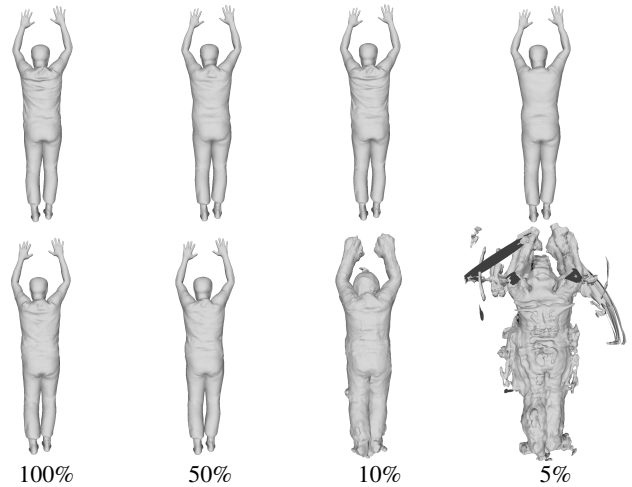


Figure 5: **Evaluation of pose encoding with different sizes of training data.** Top row: our local pose encoding. Bottom row: global pose encoding. While the global pose encoding suffers from severe overfitting artifacts, our local pose encoding generalizes well even if data size is severely limited.

graph convolutions [57], but requires surface registration for training. NASA [18], on the other hand, can be learned without registration but needs to determine occupancy values. We train both methods using registered CAPE data. Table 3 shows that our approach achieves superior reconstruction accuracy and perceptual realism, while Fig. 6 illustrates limitations of the prior methods. As CAPE relies on a template mesh with a fixed topology, the reconstructions are not only less detailed but also fail to capture topological changes such as the lifting up of the jacket. While NASA can model pose-dependent shapes using articulated implicit functions, discontinuities and ghosting artifacts are visible, as the implicit functions of each body part are learned independently, which limits generalization to unseen poses. In contrast, our approach can produce highly detailed and

Table 3: Comparison with the state-of-the-art pose-aware shape modeling methods.

	Ours	CAPE [46]	NASA [18]
$D_{s2m} \downarrow$	0.570	0.970	1.12
Int. $D_n \downarrow$	0.253	0.308	0.289
$P_i \uparrow$	0.5	0.268	0.432
$P_v \uparrow$	0.5	0.455	0.457
Ex. $P_i \uparrow$	0.5	0.214	0.343
$P_v \uparrow$	0.5	0.422	0.395

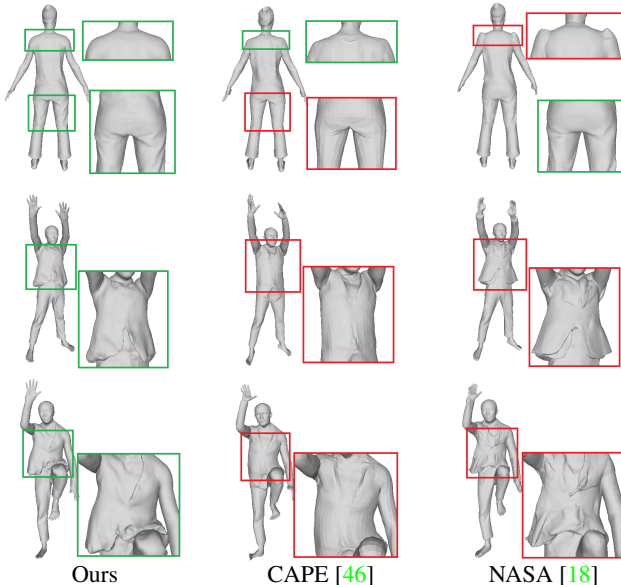


Figure 6: **Comparison with the SoTA methods.** We show qualitative results on the extrapolation task, illustrating the advantages of our method as well as the limitations of the existing approaches.

globally coherent pose-dependent deformations without template-registration.

Learning a Fully Textured Avatar. We extend our pose-aware shape modeling to appearance modeling by predicting texture fields [51, 59]; see Sup. Mat. for details. Figure 7 shows that high-resolution texture can be modeled without 2D texture mapping, which illustrates another advantage of eliminating the template-mesh requirement.

5. Discussion and Future Work

We introduced SCANimate, a fully automatic framework to create high-quality avatars (Scanimats), with realistic clothing deformations, driven by pose parameters, that are directly learned from raw 3D scans. Our experiments show that decomposing articulated deformations from scanned data is now possible in a weakly supervised manner by combining body-guided supervision with cycle-consistency regularization. Previously, the difficulty of accurate and coherent surface registration limited the field from analysing and modeling complex clothing deformations involving



Figure 7: **Textured Scanimats.** Our method can be extended to texture modeling, allowing us to automatically build a Scanimat with high-resolution realistic texture.

multiple garments from real-world observations. Our approach enables, for the first time, learning of physically plausible clothing deformations from raw scans, unlocking the possibility of realistic avatar learning from data.

Limitations and Future Work. The current representation works well for clothing that is topologically similar to the body. The method may fail for clothing, like skirts, that deviates significantly from the body; see Sup. Mat. for an example. Clothing wrinkles tend to be stochastic; that is, for a specific pose, they may differ depending on the preceding sequence of poses. The current model, however, is deterministic. Future work should factor the surface texture into albedo, shape, and lighting enabling more realistic relighting of Scanimats. Additionally, an adversarial texture loss [28] could improve visual quality. Here we model a person in a single garment. Learning a generative model with clothing variety should be possible but will require training data of varied clothing in varied poses. Most exciting is the idea of fitting Scanimats to, or even learning them from, images or videos. Finally, extending this approach to model hand articulation and facial expressions should be possible using expressive body models like SMPL-X [55].

Acknowledgements: Q. Ma was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 276693517 SFB 1233. **Disclosure:** MJB has received research gift funds from Adobe, Intel, Nvidia, Facebook, and Amazon. While MJB is a part-time employee of Amazon, his research was performed solely at, and funded solely by, Max Planck. MJB has financial interests in Amazon, Datagen Technologies, and Meshcapade GmbH.

References

- [1] Thiemo Alldieck, Marcus A. Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 1175–1186. Computer Vision Foundation / IEEE, 2019. 3
- [2] Thiemo Alldieck, Marcus A. Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3D people models. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 8387–8397. IEEE Computer Society, 2018. 3
- [3] Brett Allen, Brian Curless, and Zoran Popovic. Articulated body deformation from range scan data. *ACM Trans. Graph.*, 21(3):612–619, 2002. 3
- [4] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: shape completion and animation of people. *ACM Trans. Graph.*, 24(3):408–416, 2005. 1, 2
- [5] Alexandru O. Balan and Michael J. Black. The naked truth: Estimating body shape under clothing. In *Computer Vision - ECCV 2008, 10th European Conference on Computer Vision, Marseille, France, October 12-18, 2008, Proceedings, Part II*, volume 5303 of *Lecture Notes in Computer Science*, pages 15–29. Springer, 2008. 3
- [6] Ilya Baran and Jovan Popovic. Automatic rigging and animation of 3D characters. *ACM Trans. Graph.*, 26(3):72, 2007. 4
- [7] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. DeePSD: Automatic deep skinning and pose space deformation for 3D garment animation. *CoRR*, abs/2009.02715, 2020. 3
- [8] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Combining implicit function learning and parametric models for 3D human reconstruction. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part II*, volume 12347 of *Lecture Notes in Computer Science*, pages 311–329. Springer, 2020. 3
- [9] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. LoopReg: Self-supervised learning of implicit surface correspondences, pose and shape for 3D human mesh registration. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. 3
- [10] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-Garment Net: Learning to dress 3D people from images. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5419–5429. IEEE, 2019. 3
- [11] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter V. Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: automatic estimation of 3D human pose and shape from a single image. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V*, volume 9909 of *Lecture Notes in Computer Science*, pages 561–578. Springer, 2016. 3
- [12] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. PairedCycleGAN: Asymmetric style transfer for applying and removing makeup. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 40–48. IEEE Computer Society, 2018. 3
- [13] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 5939–5948. Computer Vision Foundation / IEEE, 2019. 2, 3, 5
- [14] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3D shape reconstruction and completion. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 6968–6979. IEEE, 2020. 3
- [15] Enric Corona, Albert Pumarola, Guillem Alenyà, Gerard Pons-Moll, and Francesc Moreno-Noguer. SMPLicit: Topology-aware generative model for clothed people. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. 3
- [16] R. Danerek, Endri Dibra, A. Cengiz Öztireli, Remo Ziegler, and Markus H. Gross. DeepGarment : 3D garment shape estimation from a single image. *Comput. Graph. Forum*, 36(2):269–280, 2017. 3
- [17] Edilson de Aguiar, Leonid Sigal, Adrien Treuille, and Jessica K. Hodgins. Stable spaces for real-time clothing. *ACM Trans. Graph.*, 29(4):106:1–106:9, 2010. 3
- [18] Boyang Deng, John P. Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey E. Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. NASA neural articulated shape approximation. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VII*, volume 12352 of *Lecture Notes in Computer Science*, pages 612–628. Springer, 2020. 3, 7, 8
- [19] Valentin Gabeur, Jean-Sébastien Franco, Xavier Martin, Cordelia Schmid, and Grégory Rogez. Moulding humans: Non-parametric 3D human shape estimation from single images. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 2232–2241. IEEE, 2019. 3
- [20] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas A. Funkhouser. Local deep implicit functions for 3D shape. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 4856–4865. IEEE, 2020. 3
- [21] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T. Freeman, and Thomas A. Funkhouser. Learning shape templates with structured implicit functions. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 7153–7163. IEEE, 2019. 3
- [22] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning

- shapes. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 3789–3799. PMLR, 2020. 3, 4, 5, 6
- [23] Peng Guan, Loretta Reiss, David A. Hirshberg, Alexander Weiss, and Michael J. Black. DRAPE: dressing any person. *ACM Trans. Graph.*, 31(4):35:1–35:10, 2012. 2, 3
- [24] Peng Guan, Alexander Weiss, Alexandru O. Balan, and Michael J. Black. Estimating human shape and pose from a single image. In *IEEE 12th International Conference on Computer Vision, ICCV 2009, Kyoto, Japan, September 27 - October 4, 2009*, pages 1381–1388. IEEE Computer Society, 2009. 3
- [25] Erhan Gundogdu, Victor Constantin, Amrollah Seifoddini, Minh Dang, Mathieu Salzmann, and Pascal Fua. GarNet: A two-stream network for fast and accurate 3D cloth draping. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 8738–8747. IEEE, 2019. 2, 3
- [26] Nils Hasler, Carsten Stoll, Martin Sunkel, Bodo Rosenhahn, and Hans-Peter Seidel. A statistical model of human pose and body shape. *Comput. Graph. Forum*, 28(2):337–346, 2009. 1, 2
- [27] Pengpeng Hu, Nastaran Nourbakhsh Kaashki, Vasile Dadarlat, and Adrian Munteanu. Learning to estimate the body shape under clothing from a single 3-d scan. *IEEE Trans. Ind. Informatics*, 17(6):3793–3802, 2021. 2
- [28] Jingwei Huang, Justus Thies, Angela Dai, Abhijit Kundu, Chiyu Max Jiang, Leonidas J. Guibas, Matthias Nießner, and Thomas A. Funkhouser. Adversarial texture optimization from RGB-D scans. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 1556–1565. IEEE, 2020. 8
- [29] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. ARCH: animatable reconstruction of clothed humans. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 3090–3099. IEEE, 2020. 3, 4, 5, 6, 7
- [30] Doug L. James and Christopher D. Twigg. Skinning mesh animations. *ACM Trans. Graph.*, 24(3):399–407, 2005. 3
- [31] Boyi Jiang, Juyong Zhang, Yang Hong, Jinhao Luo, Ligang Liu, and Hujun Bao. BCNet: Learning body and cloth shape from a single image. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XX*, volume 12365 of *Lecture Notes in Computer Science*, pages 18–35. Springer, 2020. 3
- [32] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 7122–7131. IEEE Computer Society, 2018. 3
- [33] Angjoo Kanazawa, Jason Y. Zhang, Panna Felsen, and Jitendra Malik. Learning 3D human dynamics from video. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 5614–5623. Computer Vision Foundation / IEEE, 2019. 3
- [34] Doyub Kim, Woojong Koh, Rahul Narain, Kayvon Fatahalian, Adrien Treuille, and James F. O’Brien. Near-exhaustive precomputation of secondary cloth effects. *ACM Trans. Graph.*, 32(4):87:1–87:8, 2013. 3
- [35] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. VIBE: video inference for human body pose and shape estimation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 5252–5262. IEEE, 2020. 3
- [36] Nikos Kolotouros, Georgios Pavlakos, Michael J. Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 2252–2261. IEEE, 2019. 3
- [37] Paul G. Kry, Doug L. James, and Dinesh K. Pai. EigenSkin: real time large deformation character skinning in hardware. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, San Antonio, TX, USA, July 21-22, 2002*, pages 153–159. ACM, 2002. 3
- [38] Tsuneya Kurihara and Natsuki Miyata. Modeling deformable human hands from medical images. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, Grenoble, France, August 27-29, 2004*, pages 355–363. The Eurographics Association, 2004. 3
- [39] Zorah Löhner, Daniel Cremers, and Tony Tung. Deep-Wrinkles: Accurate and realistic clothing modeling. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, volume 11208 of *Lecture Notes in Computer Science*, pages 698–715. Springer, 2018. 2, 3, 7
- [40] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, and Peter V. Gehler. Unite the people: Closing the loop between 3D and 2d human representations. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4704–4713. IEEE Computer Society, 2017. 3
- [41] John P. Lewis, Matt Cordner, and Nickson Fong. Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2000, New Orleans, LA, USA, July 23-28, 2000*, pages 165–172. ACM, 2000. 3, 4
- [42] Lijuan Liu, Youyi Zheng, Di Tang, Yi Yuan, Changjie Fan, and Kun Zhou. NeuroSkinning: automatic skin binding for production characters with deep graph networks. *ACM Trans. Graph.*, 38(4):114:1–114:12, 2019. 3, 4
- [43] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: a skinned multi-person linear model. *ACM Trans. Graph.*, 34(6):248:1–248:16, 2015. 1, 2, 3, 4
- [44] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1987, Anaheim, California, USA, July 27-31, 1987*, pages 163–169. ACM, 1987. 6

- [45] Qianli Ma, Shunsuke Saito, Jinlong Yang, Siyu Tang, and Michael J. Black. SCALE: Modeling clothed humans with a surface codec of articulated local elements. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. [2](#)
- [46] 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 *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 6468–6477. IEEE, 2020. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [47] Lars M. Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4460–4470. Computer Vision Foundation / IEEE, 2019. [2](#), [3](#), [5](#)
- [48] Marko Mihajlovic, Yan Zhang, Michael J. Black, and Siyu Tang. LEAP: Learning articulated occupancy of people. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. [3](#)
- [49] Gyeongsik Moon, Takaaki Shiratori, and Kyoung Mu Lee. DeepHandMesh: A weakly-supervised deep encoder-decoder framework for high-fidelity hand mesh modeling. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part II*, volume 12347 of *Lecture Notes in Computer Science*, pages 440–455. Springer, 2020. [3](#)
- [50] Alexandros Neophytou and Adrian Hilton. A layered model of human body and garment deformation. In *2nd International Conference on 3D Vision, 3DV 2014, Tokyo, Japan, December 8-11, 2014, Volume 1*, pages 171–178. IEEE Computer Society, 2014. [3](#)
- [51] Michael Oechsle, Lars M. Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 4530–4539. IEEE, 2019. [8](#)
- [52] Ahmed A. A. Osman, Timo Bolkart, and Michael J. Black. STAR: sparse trained articulated human body regressor. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VI*, volume 12351 of *Lecture Notes in Computer Science*, pages 598–613. Springer, 2020. [1](#), [2](#), [6](#)
- [53] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 165–174. Computer Vision Foundation / IEEE, 2019. [2](#), [4](#), [5](#)
- [54] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. TailorNet: Predicting clothing in 3D as a function of human pose, shape and garment style. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 7363–7373. IEEE, 2020. [2](#), [3](#), [7](#)
- [55] 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 *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 10975–10985. Computer Vision Foundation / IEEE, 2019. [8](#)
- [56] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J. Black. ClothCap: seamless 4d clothing capture and retargeting. *ACM Trans. Graph.*, 36(4):73:1–73:15, 2017. [2](#)
- [57] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J. Black. Generating 3D faces using convolutional mesh autoencoders. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part III*, volume 11207 of *Lecture Notes in Computer Science*, pages 725–741. Springer, 2018. [4](#), [7](#)
- [58] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: modeling and capturing hands and bodies together. *ACM Trans. Graph.*, 36(6):245:1–245:17, 2017. [3](#)
- [59] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Hao Li, and Angjoo Kanazawa. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 2304–2314. IEEE, 2019. [3](#), [8](#)
- [60] Shunsuke Saito, Tomas Simon, Jason M. Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 81–90. IEEE, 2020. [3](#)
- [61] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [6](#)
- [62] David Smith, Matthew Loper, Xiaochen Hu, Paris Mavroidis, and Javier Romero. FACSIMILE: fast and accurate scans from an image in less than a second. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5329–5338. IEEE, 2019. [3](#)
- [63] Gül Varol, Duygu Ceylan, Bryan C. Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. BodyNet: Volumetric inference of 3D human body shapes. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part VII*, volume 11211 of *Lecture Notes in Computer Science*, pages 20–38. Springer, 2018. [3](#)
- [64] Lizhen Wang, Xiaochen Zhao, Tao Yu, Songtao Wang, and Yebin Liu. NormalGAN: Learning detailed 3D human from a single RGB-D image. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XX*, volume 12365 of *Lecture Notes in Computer Science*, pages 430–446. Springer, 2020. [3](#)

- [65] Shaofei Wang, Andreas Geiger, and Siyu Tang. Locally aware piecewise transformation fields for 3D human mesh registration. In *Proceedings of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. [3](#)
- [66] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T. Freeman, Rahul Sukthankar, and Cristian Sminchisescu. GHUM & GHUML: generative 3D human shape and articulated pose models. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 6183–6192. IEEE, 2020. [1](#), [2](#)
- [67] Weiwei Xu, Nobuyuki Umetani, Qianwen Chao, Jie Mao, Xiaogang Jin, and Xin Tong. Sensitivity-optimized rigging for example-based real-time clothing synthesis. *ACM Trans. Graph.*, 33(4):107:1–107:11, 2014. [3](#), [6](#)
- [68] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. RigNet: neural rigging for articulated characters. *ACM Trans. Graph.*, 39(4):58, 2020. [3](#)
- [69] Jinlong Yang, Jean-Sébastien Franco, Franck Hétroy-Wheeler, and Stefanie Wuhrer. Estimation of human body shape in motion with wide clothing. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, volume 9908 of *Lecture Notes in Computer Science*, pages 439–454. Springer, 2016. [2](#), [3](#)
- [70] Jinlong Yang, Jean-Sébastien Franco, Franck Hétroy-Wheeler, and Stefanie Wuhrer. Analyzing clothing layer deformation statistics of 3D human motions. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part VII*, volume 11211 of *Lecture Notes in Computer Science*, pages 245–261. Springer, 2018. [2](#), [3](#), [4](#), [6](#), [7](#)
- [71] Tao Yu, Zerong Zheng, Kaiwen Guo, Jianhui Zhao, Qionghai Dai, Hao Li, Gerard Pons-Moll, and Yebin Liu. DoubleFusion: Real-time capture of human performances with inner body shapes from a single depth sensor. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 7287–7296. IEEE Computer Society, 2018. [3](#)
- [72] Chao Zhang, Sergi Pujades, Michael J. Black, and Gerard Pons-Moll. Detailed, accurate, human shape estimation from clothed 3D scan sequences. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5484–5493. IEEE Computer Society, 2017. [2](#), [3](#)
- [73] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. DeepHuman: 3D human reconstruction from a single image. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 7738–7748. IEEE, 2019. [3](#)
- [74] Hao Zhu, Xinxin Zuo, Sen Wang, Xun Cao, and Ruigang Yang. Detailed human shape estimation from a single image by hierarchical mesh deformation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4491–4500. Computer Vision Foundation / IEEE, 2019. [3](#)
- [75] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2242–2251. IEEE Computer Society, 2017. [3](#)