SCALE: Modeling Clothed Humans with a Surface Codec of Articulated Local Elements

Conference Paper**Author(s):**

Ma, Qianli; Saito, Shunsuke; Yang, Jinlong; Tang, Siyu; Black, Michael J.

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000477982>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/CVPR46437.2021.01582>

SCALE: Modeling Clothed Humans with a Surface Codec of Articulated Local Elements

Qianli Ma^{1,2} Shunsuke Saito^{1†} Jinlong Yang¹ Siyu Tang² Michael J. Black¹

¹Max Planck Institute for Intelligent Systems, Tübingen, Germany ²ETH Zürich

{qma, ssaito, jyang, black}@tuebingen.mpg.de, {siyu.tang}@inf.ethz.ch

Figure 1: **SCALE**. Given a sequence of posed but minimally-clothed 3D bodies, SCALE predicts 798 articulated surface elements (patches, each visualized with a unique color) to dress the bodies with realistic clothing that moves and deforms naturally even in the presence of topological change. The resulting dense point sets have correspondence across different poses, as shown by the consistent patch colors. The result also includes predicted surface normals and texture, with which the point cloud can either be turned into a 3D mesh or directly rendered as realistic images using neural rendering techniques.

Abstract

Learning to model and reconstruct humans in clothing is challenging due to articulation, non-rigid deformation, and varying clothing types and topologies. To enable learning, the choice of representation is the key. Recent work uses neural networks to parameterize local surface elements. This approach captures locally coherent geometry and non-planar details, can deal with varying topology, and does not require registered training data. However, naively using such methods to model 3D clothed humans fails to capture fine-grained local deformations and generalizes poorly. To address this, we present three key innovations: First, we deform surface elements based on a human body model such that large-scale deformations caused by articulation are explicitly separated from topological changes and local clothing deformations. Second, we address the limitations of existing neural surface elements by regressing local geometry from local features, significantly improving the expressiveness. Third, we learn a pose embedding on a 2D parameterization space that encodes posed body geometry, improving generalization to unseen poses by reducing non-local spurious correlations.

We demonstrate the efficacy of our surface representation by learning models of complex clothing from point clouds. The clothing can change topology and deviate from the topology of the body. Once learned, we can animate previously unseen motions, producing high-quality point clouds, from which we generate realistic images with neural rendering. We assess the importance of each technical contribution and show that our approach outperforms the state-of-the-art methods in terms of reconstruction accuracy and inference time. The code is available for research purposes at <https://qianlim.github.io/SCALE>.

1. Introduction

While models of humans in clothing would be valuable for many tasks in computer vision such as body pose and shape estimation from images and videos [9, 15, 31, 32, 35, 36] and synthetic data generation [60, 61, 71, 83], most existing approaches are based on “minimally-clothed” human body models [2, 30, 42, 49, 54, 75], which do not represent clothing. To date, statistical models for clothed humans remain lacking despite the broad range of potential

[†] Now at Facebook Reality Labs.

applications. This is likely due to the fact that modeling 3D clothing shapes is much more difficult than modeling body shapes. Fundamentally, several characteristics of clothed bodies present technical challenges for representing clothing shapes.

The first challenge is that clothing shape varies at different spatial scales driven by global body articulation and local clothing geometry. The former requires the representation to properly handle human pose variation, while the latter requires local expressiveness to model folds and wrinkles. Second, a representation must be able to model smooth cloth surfaces and also sharp discontinuities and thin structures. Third, clothing is diverse and varies in terms of its topology. The topology can even change with the motion of the body. Fourth, the relationship between the clothing and the body changes as the clothing moves relative to the body surface. Finally, the representation should be compatible with existing body models and should support fast inference and rendering, enabling real-world applications.

Unfortunately, none of the existing 3D shape representations satisfy all these requirements. The standard approach uses 3D meshes that are draped with clothing using physics simulation [3, 38, 41]. These require manual clothing design and the physics simulation makes them inappropriate for inference. Recent work starts with classical rigged 3D meshes and blend skinning but uses machine learning to model clothing shape and local non-rigid shape deformation. However, these methods often rely on pre-defined garment templates [8, 37, 45, 53], and the fixed correspondence between the body and garment template restricts them from generalizing to arbitrary clothing topology. Additionally, learning a mesh-based model requires registering a common 3D mesh template to scan data. This is time consuming, error prone, and limits topology change [56]. New neural implicit representations [12, 46, 51], on the other hand, are able to reconstruct topologically varying clothing types [13, 16, 65], but are not consistent with existing graphics tools, are expensive to render, and are not yet suitable for fast inference. Point clouds are a simple representation that also supports arbitrary topology [21, 39, 77] and does not require data registration, but highly detailed geometry requires many points.

A middle ground solution is to utilize a collection of parametric surface elements that smoothly conform to the global shape of the target geometry [20, 25, 80, 82, 84]. As each element can be freely connected or disconnected, topologically varying surfaces can be effectively modeled while retaining the efficiency of explicit shape inference. Like point clouds, these methods can be learned without data registration.

However, despite modeling coherent global shape, existing surface-element-based representations often fail to generate local structures with high-fidelity. The key limiting

factor is that shapes are typically decoded from *global* latent codes [25, 80, 82], i.e. the network needs to learn both the global shape statistics (caused by articulation) and a prior for local geometry (caused by clothing deformation) at once. While the recent work of [24] shows the ability to handle articulated objects, these methods often fail to capture local structures such as sharp edges and wrinkles, hence the ability to model *clothed* human bodies has not been demonstrated.

In this work, we extend the surface element representation to create a clothed human model that meets *all* the aforementioned desired properties. We support articulation by defining the surface elements on top of a minimal clothed body model. To densely cover the surface, and effectively model local geometric details, we first introduce a global patch descriptor that differentiates surface elements at different locations, enabling the modeling of hundreds of local surface elements with a single network, and then regress local non-rigid shapes from local pose information, producing folding and wrinkles. Our new shape representation, *Surface Codec of Articulated Local Elements*, or SCALE, demonstrates state-of-the-art performance on the challenging task of modeling the per-subject pose-dependent shape of clothed humans, setting a new baseline for modeling topologically varying high-fidelity surface geometry with explicit shape inference. See Fig. 1.

In summary, our contributions are: (1) an extension of surface element representations to non-rigid articulated object modeling; (2) a revised local elements model that generates local geometry from local shape signals instead of a global shape vector; (3) an explicit shape representation for clothed human shape modeling that is robust to varying topology, produces high-visual-fidelity shapes, is easily controllable by pose parameters, and achieves fast inference; and (4) a novel approach for modeling humans in clothing that does not require registered training data and generalizes to various garment types of different topology, addressing the missing pieces from existing clothed human models. We also show how neural rendering is used together with our point-based representation to produce high-quality rendered results. The code is available for research purposes at <https://qianlim.github.io/SCALE>.

2. Related Work

Shape Representations for Modeling Humans. Surface meshes are the most commonly used representation for human shape due to their efficiency and compatibility with graphics engines. Not only human body models [2, 42, 49, 75] but also various clothing models leverage 3D mesh representations as separate mesh layers [17, 26, 27, 37, 53, 67] or displacements from a minimally clothed body [8, 45, 48, 69, 74, 79]. Recent advances in deep learn-

ing have improved the fidelity and expressiveness of mesh-based approaches using graph convolutions [45], multilayer perceptrons (MLP) [53], and 2D convolutions [29, 37]. The drawback of mesh-based representations is that topological changes are difficult to model. Pan et al. [50] propose a topology modification network (TMN) to support topological change, however it has difficulty learning large topological changes from a single template mesh [85].

To support various topologies, neural implicit surfaces [46, 51] have recently been applied to clothed human reconstruction and registration [6, 65]. The extension of implicit surface models to unsigned distances [14] or probability densities [10] even allows thin structures to be represented with high resolution. Recent work [23] also shows the ability to learn a signed-distance function (SDF) directly from incomplete scan data. Contemporaneous with our work, SCANimate [66], learns an implicit shape model of clothed people from raw scans. Also contemporaneous is SMPLicit [16], which learns an implicit clothing representation that can be fit to scans or 2D clothing segmentations. Despite the impressive reconstruction quality of implicit methods, extracting an explicit surface is time consuming but necessary for many applications.

Surface element representations [20, 25, 73, 80] are promising for modeling clothing more explicitly. These methods approximate various topologies with locally coherent geometry by learning to deform single or multiple surface elements. Recent work improves these patch-based approaches by incorporating differential geometric regularization [4, 19], demonstrating simple clothing shape reconstruction. Although patch-based representations relax the topology constraint of a single template by representing the 3D surface as the combination of multiple surface elements, reconstruction typically lacks details as the number of patches is limited due to the memory requirements, which scale linearly with the number of patches. While SCALE is based on these neural surface representations, we address the limitations of existing methods in Sec. 3.1, enabling complex clothed human modeling.

Articulated Shape Modeling. Articulation often dominates large-scale shape variations for articulated objects, such as hands and human bodies. To efficiently represent shape variations, articulated objects are usually modeled with meshes driven by an embedded skeleton [2, 42, 62]. Mesh-based clothing models also follow the same principle [26, 37, 45, 48, 53], where shapes are decomposed into articulated deformations and non-rigid local shape deformations. While the former are explained by body joint transformations, the latter can be efficiently modeled in a canonical space. One limitation, however, is that registered data or physics-based simulation is required to learn these deformations on a template mesh with a fixed topology. In contrast, recent work on articulated implicit shape model-

ing [18, 47, 66] does not require surface registration. In this work we compare with Deng et al. [18] on a clothed human modeling task from point clouds and show the superiority of our approach in terms of generalization to unseen poses, fidelity, and inference speed.

Local Shape Modeling. Instead of learning 3D shapes solely with a global feature vector, recent work shows that learning from local shape variations leads to detailed and highly generalizable 3D reconstruction [11, 13, 22, 52, 55, 64, 70]. Leveraging local shape priors is effective for 3D reconstruction tasks from 3D point clouds [11, 13, 22, 28, 55, 70] and images [52, 64, 76]. Inspired by this prior work, SCALE leverages both local and global feature representations, which leads to high-fidelity reconstruction as well as robust generalization to unseen poses.

3. SCALE

Figure 2 shows an overview of SCALE. Our goal is to model clothed humans with a topologically flexible point-based shape representation that supports fast inference and animation with SMPL pose parameters [42]. To this end, we model pose-dependent shape variations of clothing using a collection of local surface elements (patches) that are associated with a set of pre-defined locations on the body. Our learning-based local pose embedding further improves the generalization of pose-aware clothing deformations (Sec. 3.1). Using this local surface element representation, we train a model for each clothing type to predict a set of 3D points representing the clothed body shape given an unclothed input body. Together with the predicted point normals and colors, the dense point set can be meshed or realistically rendered with neural rendering (Sec. 3.2).

3.1. Articulated Local Elements

While neural surface elements [25, 80, 82, 84] offer locally coherent geometry with fast inference, the existing formulations have limitations that prevent us from applying them to clothed-human modeling. We first review the existing neural surface elements and introduce our formulation that addresses the drawbacks of the prior work.

Review: Neural Surface Elements. The original methods that model neural surface elements [25, 80] learn a function to generate 3D point clouds as follows:

$$f_w(\mathbf{p}; \mathbf{z}) : \mathbb{R}^D \times \mathbb{R}^Z \rightarrow \mathbb{R}^3, \quad (1)$$

where f_w is a multilayer perceptron (MLP) parameterized by weights w , $\mathbf{p} \in \mathbb{R}^D$ is a point on the surface element, and $\mathbf{z} \in \mathbb{R}^Z$ is a global feature representing object shape. Specifically, f_w maps $\mathbf{p}$ on the surface element to a point on the surface of a target 3D object conditioned by a shape code $\mathbf{z}$. Due to the inductive bias of MLPs, the resulting 3D point clouds are geometrically smooth within the ele-

Figure 2: **Method overview.** Given a posed, minimally-clothed body, we define a set of points on its surface and associate a local element (a square patch) with each of them. The body points’ positions in $\mathbb{R}^3$ are recorded on a 2D UV positional map, which is convolved by a UNet to obtain pixel-aligned local pose features z_k . The 2D coordinates $\mathbf{u}_k = (u_k, v_k)$ on the UV map, the pose features z_k , and the 2D coordinates within the local elements $\mathbf{p} = (p, q)$ are fed into a shared MLP to predict the deformation of the local elements in the form of residuals from the body. The inferred local elements are finally articulated by the known transformations of corresponding body points to generate a posed clothed human. Each patch and its corresponding body point is visualized with the same color.

ment [73]. While this smoothness is desirable for surface modeling, to support different topologies, AtlasNet [25] requires multiple surface elements represented by individual networks, which increases network parameters and memory cost. The cost is linearly proportional to the number of patches. As a result, these approaches limit expressiveness for topologically complex objects such as clothed humans.

Another line of work represents 3D shapes using a collection of local elements. PointCapsNet [84] decodes a *local* shape code $\{z_k\}_{k=1}^K$, where K is the number of local elements, into *local* patches with separate networks:

$$f_{w_k}(\mathbf{p}; z_k) : \mathbb{R}^D \times \mathbb{R}^Z \rightarrow \mathbb{R}^3. \quad (2)$$

While modeling local shape statistics instead of global shape variations improves the generalization and training efficiency for diverse shapes, the number of patches is still difficult to scale up as in AtlasNet for the same reason.

Point Completion Network (PCN) [82] uses two-stage decoding: the first stage predicts a coarse point set of the target shape, then these points are used as basis points for the second stage. At each basis location $\mathbf{b}_k \in \mathbb{R}^3$, points $\mathbf{p}$ from a local surface element (a regular grid) are sampled and fed into the second decoder as follows:

$$f_w(\mathbf{b}_k, \mathbf{p}; \mathbf{z}) : \mathbb{R}^3 \times \mathbb{R}^D \times \mathbb{R}^Z \rightarrow \mathbb{R}^3. \quad (3)$$

Notably, PCN utilizes a single network to model a large number of local elements, improving the expressiveness with an arbitrary shape topology. However, PCN relies on a global shape code $\mathbf{z}$ that requires learning global shape statistics, resulting in poor generalization to unseen data samples as demonstrated in Sec. 4.3.

Articulated Local Elements. For clothed human modeling, the shape representation needs to be not only expressive but also highly generalizable to unseen poses. These

requirements and the advantages of the prior methods lead to our formulation:

$$g_w(\mathbf{u}_k, \mathbf{p}; z_k) : \mathbb{R}^{D_1} \times \mathbb{R}^{D_2} \times \mathbb{R}^Z \rightarrow \mathbb{R}^3, \quad (4)$$

where $\mathbf{u}_k \in \mathbb{R}^{D_1}$ is a global patch descriptor that provides inter-patch relations and helps the network g_w distinguish different surface elements, and $\mathbf{p} \in \mathbb{R}^{D_2}$ are the local (intra-patch) coordinates within each surface element. Importantly, our formulation achieves higher expressiveness by efficiently modeling a large number of local elements using a single network as in [82] while improving generality by learning local shape variations with z_k .

Moreover, unlike the existing methods [24, 25, 80, 84], where the networks directly predict point locations in $\mathbb{R}^3$, our network $g_w(\cdot)$ models *residuals* from the minimally-clothed body. To do so, we define a set of points $\mathbf{t}_k \in \mathbb{R}^3$ on the posed body surface, and predict a local element (in the form of residuals) for each body point: $\mathbf{r}_{k,i} = g_w(\mathbf{u}_k, \mathbf{p}_i; z_k)$, where $\mathbf{p}_i$ denotes a sampled point from a local element. In particular, an $\mathbf{r}_{k,i}$ is relative to a *local coordinate system*¹ that is defined on $\mathbf{t}_k$. To obtain a local element’s position in the world coordinate $\mathbf{x}_{k,i}$, we apply articulations to $\mathbf{r}_{k,i}$ by the known transformation $\mathbf{T}_k$ associated with the local coordinate system, and add it to $\mathbf{t}_k$:

$$\mathbf{x}_{k,i} = \mathbf{T}_k \cdot \mathbf{r}_{k,i} + \mathbf{t}_k. \quad (5)$$

Our network $g_w(\cdot)$ also predicts surface normals as an additional output for meshing and neural rendering, which are also transformed by $\mathbf{T}_k$. The residual formulation with explicit articulations is critical to clothed human modeling as the network $g_w(\cdot)$ can focus on learning local shape variations, which are roughly of the same scale. This leads to the successful recovery of fine-grained clothing deformations

¹See SupMat. for more details on the definition of the local coordinates.

as shown in Sec. 4. Next, we define local and global patch descriptors as well as the local shape feature z_k .

Local Descriptor. Each local element approximates a continuous small region on the target 2-manifold. Following [82], we evenly sample M points on a 2D grid and use them as a local patch descriptor: $\mathbf{p} = (p_i, q_i) \in \mathbb{R}^2$, with $p_i, q_i \in [0, 1]$, $i = 1, 2, \dots, M$. Within each surface element, all sampled points share the same global patch descriptor $\mathbf{u}_k$ and patch-wise feature z_k .

Global Descriptor. The global patch descriptor $\mathbf{u}_k$ in Eq. (4) is the key to modeling different patches with a single network. While each global descriptor needs to be unique, it should also provide proximity information between surface elements to generate a globally coherent shape. Thus, we use 2D location on the UV positional map of the human body as a global patch descriptor: $\mathbf{u}_k = (u_k, v_k)$. While the 3D positions of a neutral human body can also be a global descriptor as in [24], we did not observe any performance gain. Note that $\mathbf{T}_k$ and $\mathbf{t}_k$ in Eq. (4) are assigned based on the corresponding 3D locations on the UV positional map.

Pose Embedding. To model realistic pose-dependent clothing deformations, we condition the proposed neural network with pose information from the underlying body as the local shape feature z_k . While conditioning every surface element on global pose parameters θ is possible, in the spirit of prior work [37, 45, 53, 78], we observe that such global pose conditioning does not generalize well to new poses and the network learns spurious correlations between body parts (a similar issue was observed and discussed in [49] for parametric human body modeling). Thus, we introduce a learning-based pose embedding using a 2D positional map, where each pixel consists of the 3D coordinates of a unique point on the underlying body mesh normalized by a transformation of the root joint. This 2D positional map is fed into a UNet [63] to predict a 64-channel feature z_k for each pixel. The advantage of our learning-based pose embedding is two-fold: first, the influence of each body part is clothing-dependent and by training end-to-end, the learning-based embedding ensures that reconstruction fidelity is maximized adaptively for each outfit. Furthermore, 2D CNNs have an inductive bias to favor local information regardless of theoretical receptive fields [44], effectively removing non-local spurious correlations. See Sec. 4.3 for a comparison of our local pose embedding with its global counterparts.

3.2. Training and Inference

For each input body, SCALE generates a point set $\mathbf{X}$ that consists of K deformed surface elements, with M points sampled from each element: $|\mathbf{X}| = KM$. From its corresponding clothed body surface, we sample a point set $\mathbf{Y}$ of size N (i.e., $|\mathbf{Y}| = N$) as ground truth. The network is

trained end-to-end with the following loss:

$$\mathcal{L} = \lambda_d \mathcal{L}_d + \lambda_n \mathcal{L}_n + \lambda_r \mathcal{L}_r + \lambda_c \mathcal{L}_c, \quad (6)$$

where $\lambda_d, \lambda_n, \lambda_r, \lambda_c$ are weights that balance the loss terms.

First, the Chamfer loss $\mathcal{L}_d$ penalizes bi-directional point-to-point $L2$ distances between the generated point set $\mathbf{X}$ and the ground truth point set $\mathbf{Y}$ as follows:

$$\mathcal{L}_d = d(\mathbf{x}, \mathbf{y}) = \frac{1}{KM} \sum_{k=1}^K \sum_{i=1}^M \min_j \|\mathbf{x}_{k,i} - \mathbf{y}_j\|_2^2 + \frac{1}{N} \sum_{j=1}^N \min_{k,i} \|\mathbf{x}_{k,i} - \mathbf{y}_j\|_2^2. \quad (7)$$

For each predicted point $\mathbf{x}_{k,i} \in \mathbf{X}$, we penalize the $L1$ difference between its normal and that of its nearest neighbor from the ground truth point set: $\mathcal{L}_n =$

$$\frac{1}{KM} \sum_{k=1}^K \sum_{i=1}^M \left\| \mathbf{n}(\mathbf{x}_{k,i}) - \mathbf{n}(\arg\min_{\mathbf{y}_j \in \mathbf{Y}} d(\mathbf{x}_{k,i}, \mathbf{y}_j)) \right\|_1, \quad (8)$$

where $\mathbf{n}(\cdot)$ denotes the unit normal of the given point. We also add $L2$ regularization on the predicted residual vectors to prevent extreme deformations:

$$\mathcal{L}_r = \frac{1}{KM} \sum_{k=1}^K \sum_{i=1}^M \|\mathbf{r}_{k,i}\|_2^2. \quad (9)$$

When the ground-truth point clouds are textured, SCALE can also represent RGB color inference by predicting another 3 channels, which can be trained with an $L1$ reconstruction loss:

$$\mathcal{L}_c = \frac{1}{KM} \left\| \mathbf{c}(\mathbf{x}_{k,i}) - \mathbf{c}(\arg\min_{\mathbf{y}_j \in \mathbf{Y}} d(\mathbf{x}_{k,i}, \mathbf{y}_j)) \right\|_1, \quad (10)$$

where $\mathbf{c}(\cdot)$ represents the RGB values of the given point.

Inference, Meshing, and Rendering. SCALE inherits the advantage of existing patch-based methods for fast inference. Within a surface element, we can sample arbitrarily dense points to obtain high-resolution point clouds. Based on the area of each patch, we adaptively sample points to keep the point density constant. Furthermore, since SCALE produces oriented point clouds with surface normals, we can apply off-the-shelf meshing methods such as Ball Pivoting [5] and Poisson Surface Reconstruction (PSR) [33, 34]. As the aforementioned meshing methods are sensitive to hyperparameters, we present a method to directly render the SCALE outputs into high-resolution images by leveraging neural rendering based on point clouds [1, 57, 81]. In Sec. 4, we demonstrate that we can render the SCALE outputs using SMPLpix [57]. See SupMat for more details on the adaptive point sampling and our neural rendering pipeline.

4. Experiments

4.1. Experimental Setup

Baselines. To evaluate the efficacy of SCALE’s novel neural surface elements, we compare it to two state-of-the-art methods for clothed human modeling using meshes (CAPE [45]) and implicit surfaces (NASA [18]). We also compare with prior work based on neural surface elements: AtlasNet [25] and PCN [82]. Note that we choose a minimally-clothed body with a neutral pose as a surface element for these approaches as in [24] for fair comparison. To fully evaluate each technical contribution, we provide an ablation study that evaluates the use of explicit articulation, the global descriptor $\mathbf{u}_k$, the learning-based pose embedding using UNet, and the joint-learning of surface normals.

Datasets. We primarily use the CAPE dataset [45] for evaluation and comparison with the baseline methods. The dataset provides registered mesh pairs (clothed and minimally clothed body) of multiple humans in motion wearing common clothing (e.g. T-shirts, trousers, and a blazer). In the main paper we choose *blazerlong* (blazer jacket, long trousers) and *shortlong* (short T-shirt, long trousers) with subject 03375 to illustrate the applicability of our approach to different clothing types. The numerical results on other CAPE subjects are provided in the SupMat. In addition, to evaluate the ability of SCALE to represent a topology that significantly deviates from the body mesh, we synthetically generate point clouds of a person wearing a skirt using physics-based simulation driven by the motion of the subject 00134 in the CAPE dataset. The motion sequences are randomly split into training (70%) and test (30%) sets.

Metrics. We numerically evaluate the reconstruction quality of each method using Chamfer Distance (Eq. (7), in m^2) and the $L1$ -norm of the unit normal discrepancy (Eq. (8)), evaluated over the 12,768 points generated by our model. For CAPE [45], as the mesh resolution is relatively low, we uniformly sample the same number of points as our model on the surface using barycentric interpolation. As NASA [18] infers an implicit surface, we extract an iso-surface using Marching Cubes [43] with a sufficiently high resolution (512^3), and sample the surface. We sample and compute the errors three times with different random seeds and report the average values.

Implementation details. We use the SMPL [42] UV map with a resolution of 32×32 for our pose embedding, which yields $K = 798$ body surface points (hence the number of surface elements). For each element, we sample $M = 16$ square grid points, resulting in 12,768 points in the final output. We uniformly sample $N = 40,000$ points from each clothed body mesh as the target ground truth scan. More implementation details are provided in the SupMat.

Figure 3: **Qualitative comparison with mesh and implicit methods.** Our method produces coherent global shape, salient pose-dependent deformation, and sharp local geometry. The meshed results are acquired by applying PSR [34] to SCALE’s point+normal prediction. The patch color visualization assigns a consistent set of colors to the patches, showing correspondence between the two bodies.

4.2. Comparison: Mesh and Implicit Surface

Block I of Tab. 1 quantitatively compares the accuracy and inference runtime of SCALE, CAPE [45] and NASA [18]. CAPE [45] learns the shape variation of articulated clothed humans as displacements from a minimally clothed body using MeshCNN [59]. In contrast to ours, by construction of a mesh-based representation, CAPE requires registered templates to the scans for training. While NASA, on the other hand, learns the composition of articulated implicit functions without surface registration, it requires watertight meshes because the training requires ground-truth occupancy information. Note that these two approaches are unable to process the skirt sequences as the thin structure of the skirt is non-trivial to handle using the fixed topology of human bodies or implicit functions.

For the other two clothing types, our approach not only achieves the best numerical result, but also qualitatively demonstrates globally coherent and highly detailed reconstruction results as shown in Fig. 3. On the contrary, the mesh-based approach [45] suffers from a lack of details and fidelity, especially in the presence of topological change. Despite its topological flexibility, the articulated implicit function [18] is outperformed by our method by a large margin, especially on the more challenging *blazerlong* data (22% in Chamfer- L_2). This is mainly due to the artifacts caused by globally incoherent shape predictions for unseen poses, Fig. 3. We refer to the SupMat for extended qualitative comparison with the baselines.

The run-time comparison illustrates the advantage of fast

Table 1: Results of pose dependent clothing deformation prediction on unseen test sequences from the 3 prototypical garment types, of varying modeling difficulty. Best results are in **boldface**.

Methods / Variants	Chamfer- L_2 ($\times 10^{-4} m^2$) $\downarrow$			Normal diff. ($\times 10^{-1}$) $\downarrow$			Inference time (s) $\downarrow$
	<i>blazerlong</i>	<i>shortlong</i>	<i>skirt</i>	<i>blazerlong</i>	<i>shortlong</i>	<i>skirt</i>	
I. Comparison with SoTAs, Sec. 4.2							
CAPE [45]	1.96	1.37	–	1.28	1.15	–	0.013
NASA [18]	1.37	0.95	–	1.29	1.17	–	12.0
Ours (SCALE, full model)	1.07	0.89	2.69	1.22	1.12	0.94	0.009 (+ 1.1)
II. Global vs Local Elements, Sec. 4.3							
a). global z [24] + AN [25]	6.20	1.54	4.59	1.70	1.86	2.46	–
b). global z [24] + AN [25] + Arti.	1.32	0.99	2.95	1.46	1.23	1.20	–
c). global z [24] + PCN [82] + Arti.	1.43	1.09	3.07	1.59	1.40	1.32	–
d). pose param θ + PCN [82] + Arti.	1.46	1.04	2.90	1.61	1.39	1.33	–
III. Ablation Study: Key Components, Sec. 4.4							
e). w/o Arti. $\mathbf{T}_k$	1.92	1.43	2.71	1.70	1.53	0.96	–
f). w/o $\mathbf{u}_k$	1.08	0.89	2.71	1.22	1.12	0.94	–
g). UNet $\rightarrow$ PointNet, with $\mathbf{u}_k$	1.46	1.03	2.72	1.34	1.16	0.95	–
h). UNet $\rightarrow$ PointNet, w/o $\mathbf{u}_k$	1.92	1.42	3.06	1.69	1.52	1.39	–
i). w/o Normal pred.	1.08	0.99	2.72	–	–	–	–

inference with the explicit shape representations. CAPE directly generates a mesh (with 7K vertices) in 13ms. SCALE generates a set of 13K points within 9ms; if a mesh output is desired, the PSR meshing takes 1.1s. Note, however, that the SCALE outputs can directly be neural-rendered into images at interactive speed, see Sec. 4.5. In contrast, NASA requires densely evaluating occupancy values over the space, taking 12s to extract an explicit mesh.

4.3. Global vs. Local Neural Surface Elements

We compare existing neural surface representations [24, 25, 82] in Fig. 4 and block II of Tab. 1. Following the original implementation of AtlasNet [25], we use a global encoder that provides a global shape code $z \in \mathbb{R}^{1024}$ based on PointNet [58]. We also provide a variant of AtlasNet [25] and PCN [82], where the networks predict residuals on top of the input body and then are articulated as in our approach. AtlasNet with the explicit articulation (b) significantly outperforms the original AtlasNet without articulation (a). This shows that our newly introduced articulated surface elements are highly effective for modeling articulated objects, regardless of neural surface formulations. As PCN also efficiently models a large number of local elements using a single network, (c) and (d) differ from our approach only in the use of a global shape code z instead of local shape codes. While (c) learns the global code in an end-to-end manner, (d) is given global pose parameters θ *a priori*. Qualitatively, modeling local elements with a global shape code leads to noisier results. Numerically, our method outperforms both approaches, demonstrating the importance of modeling local shape codes. Notably, an-

other advantage of modeling local shape codes is its parameter efficiency. The global approaches often require high dimensional latent codes (e.g. 1024), leading to the high usage of network parameters (1.06M parameters for the networks above). In contrast, our local shape modeling allows us to efficiently model shape variations with significantly smaller latent codes (64 in SCALE) with nearly half the trainable parameters (0.57M) while achieving the state-of-the-art modeling accuracy.

4.4. Ablation Study

We further evaluate our technical contributions via an ablation study. As demonstrated in Sec. 4.3 and Tab. 1 (e), explicitly modeling articulation plays a critical role in the success of accurate clothed human modeling. We also observe a significant degradation by replacing our UNet-based pose embedding with PointNet, denoted as (g) and (h). This indicates that the learning-based pose embedding with a 2D CNN is more effective for local feature learning despite the conceptual similarity of these two architectures that incorporate spatial proximity information. Interestingly, the lack of a global descriptor derived from the UV map, denoted as (f), has little impact on numerical accuracy. As the similar ablation study between (g) and (h) shows significant improvement by adding u_k in the case of the PointNet architecture, this result implies that our UNet local encoder implicitly learns the global descriptor as part of the local codes z_k . As shown in Tab. 1 (i), another interesting observation is that the joint training of surface normals improves reconstruction accuracy, indicating that the multi-task learning of geometric features can be mutually beneficial.

Figure 4: **Qualitative results of the ablation study.** Points are colored according to predicted normals. Our full model produces globally more coherent and locally more detailed results compared to the baselines. Note the difference at the bottom of the blazer (upper row) and the skirt (lower row).

Figure 5: **Neural Rendering of SCALE.** The dense point set of textured (upper row) or normal-colored (lower row) predictions from SCALE (the left image in each pair) can be rendered into realistic images with a state-of-the-art neural renderer [57].

4.5. Neural Rendering of SCALE

The meshing process is typically slow, prone to artifacts, and sensitive to the choice of hyperparameters. To circumvent meshing while realistically completing missing regions, we show that generated point clouds can be directly rendered into high-resolution images with the help of the SMPLpix [57] neural renderer, which can generate e.g. a 512×512 image in $42ms$. Figure 5 shows that the dense point clouds generated by SCALE are turned into complete images in which local details such as fingers and wrinkles are well preserved. Note that we show the normal color-coded renderings for the synthetic skirt examples, since they lack ground-truth texture information.

5. Conclusion

We introduce SCALE, a highly flexible explicit 3D shape representation based on pose-aware local surface elements with articulation, which allows us to faithfully model a clothed human using point clouds without relying on a fixed-topology template, registered data, or watertight scans. The evaluation demonstrates that efficiently modeling a large number of local elements and incorporating explicit articulation are the key to unifying the learning of complex clothing deformations of various topologies.

Limitations and future work. While the UV map builds a correspondence across all bodies, a certain patch produced by SCALE is not guaranteed to represent semantically the same region on the cloth in different poses. Jointly optimizing explicit correspondences [7, 72] with explicit shape representations like ours remains challenging yet promising. Currently, SCALE models clothed humans in a subject-specific manner but our representation should support learning a unified model across multiple garment types. While we show that it is possible to obviate the meshing step by using neural rendering, incorporating learnable triangulation [40, 68] would be useful for applications that need meshes.

Acknowledgements: We thank Sergey Prokudin for insightful discussions and the help with the SMPLpix rendering. Qianli Ma is partially funded by 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] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 696–712, 2020. [5](#)
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: shape completion and animation of people. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 408–416, 2005. [1](#), [2](#), [3](#)
- [3] David Baraff and Andrew P. Witkin. Large steps in cloth simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, pages 43–54, 1998. [2](#)
- [4] Jan Bednářik, Shaifali Parashar, Erhan Gundogdu, Mathieu Salzmann, and Pascal Fua. Shape reconstruction by learning differentiable surface representations. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4715–4724, 2020. [3](#)
- [5] Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5(4):349–359, 1999. [5](#)
- [6] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Combining implicit function learning and parametric models for 3D human reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 311–329, 2020. [3](#)
- [7] 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 (NeurIPS)*, pages 12909–12922, 2020. [8](#)
- [8] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-Garment Net: Learning to dress 3D people from images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5420–5430, 2019. [2](#)
- [9] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 561–578, 2016. [1](#)
- [10] Ruojin Cai, Guandao Yang, Hadar Averbuch-Elor, Zekun Hao, Serge Belongie, Noah Snavely, and Bharath Hariharan. Learning gradient fields for shape generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 364–381, 2020. [3](#)
- [11] Rohan Chhabra, Jan E. Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3D reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 608–625, 2020. [3](#)
- [12] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5939–5948, 2019. [2](#)
- [13] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3D shape reconstruction and completion. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 6970–6981, 2020. [2](#), [3](#)
- [14] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 21638–21652, 2020. [3](#)
- [15] Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J. Black. Monocular expressive body regression through body-driven attention. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 20–40, 2020. [1](#)
- [16] 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)*, 2021. [2](#), [3](#)
- [17] R Daněřek, Endri Dibra, Cengiz Öztireli, Remo Ziegler, and Markus Gross. DeepGarment: 3D garment shape estimation from a single image. In *Computer Graphics Forum*, volume 36, pages 269–280, 2017. [2](#)
- [18] Boyang Deng, JP Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Neural articulated shape approximation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 612–628, 2020. [3](#), [6](#), [7](#)
- [19] Zhantao Deng, Jan Bednářik, Mathieu Salzmann, and Pascal Fua. Better patch stitching for parametric surface reconstruction. In *International Conference on 3D Vision (3DV)*, pages 593–602, 2020. [3](#)
- [20] Theo Deprelle, Thibault Groueix, Matthew Fisher, Vladimir Kim, Bryan Russell, and Mathieu Aubry. Learning elementary structures for 3D shape generation and matching. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 7433–7443, 2019. [2](#), [3](#)
- [21] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3D object reconstruction from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2463–2471, 2017. [2](#)
- [22] Kyle Genova, Forrester Cole, Ayneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3D shape. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4857–4866, 2020. [3](#)
- [23] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of Machine Learning and Systems (ICML)*, pages 3569–3579, 2020. [3](#)
- [24] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. 3D-CODED: 3D correspondences by deep deformation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 230–246, 2018. [2](#), [4](#), [5](#), [6](#), [7](#)
- [25] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. A papier-mâché approach to learning 3D surface generation. *Proceedings IEEE*

- Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 216–224, 2018. [2](#), [3](#), [4](#), [6](#), [7](#)
- [26] Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. DRAPE: DRessing Any Person. *ACM Transactions on Graphics (TOG)*, 31(4):35–1, 2012. [2](#), [3](#)
- [27] 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 *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8739–8748, 2019. [2](#)
- [28] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, and Thomas Funkhouser. Local implicit grid representations for 3D scenes. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 6001–6010, 2020. [3](#)
- [29] Ning Jin, Yilin Zhu, Zhenglin Geng, and Ronald Fedkiw. A pixel-based framework for data-driven clothing. In *Computer Graphics Forum*, volume 39, pages 135–144, 2020. [3](#)
- [30] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total Capture: A 3D deformation model for tracking faces, hands, and bodies. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8320–8329, 2018. [1](#)
- [31] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 7122–7131, 2018. [1](#)
- [32] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3D human dynamics from video. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5614–5623, 2019. [1](#)
- [33] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics Symposium on Geometry Processing*, volume 7, 2006. [5](#)
- [34] Michael Kazhdan and Hugues Hoppe. Screened Poisson surface reconstruction. *ACM Transactions on Graphics (TOG)*, 32(3):1–13, 2013. [5](#), [6](#)
- [35] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. VIBE: Video inference for human body pose and shape estimation. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5252–5262, 2020. [1](#)
- [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 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2252–2261, 2019. [1](#)
- [37] Zorah Löhner, Daniel Cremers, and Tony Tung. DeepWrinkles: Accurate and realistic clothing modeling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 698–715, 2018. [2](#), [3](#), [5](#)
- [38] Junbang Liang, Ming C. Lin, and Vladlen Koltun. Differentiable cloth simulation for inverse problems. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 771–780, 2019. [2](#)
- [39] Chen-Hsuan Lin, Chen Kong, and Simon Lucey. Learning efficient point cloud generation for dense 3D object reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 7114–7121, 2018. [2](#)
- [40] Minghua Liu, Xiaoshuai Zhang, and Hao Su. Meshing point clouds with predicted intrinsic-extrinsic ratio guidance. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 68–84, 2020. [8](#)
- [41] Tiantian Liu, Adam W. Bargteil, James F. O’Brien, and Ladislav Kavan. Fast simulation of mass-spring systems. *ACM Transactions on Graphics (TOG)*, 32(6):214:1–214:7, 2013. [2](#)
- [42] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG)*, 34(6):248, 2015. [1](#), [2](#), [3](#), [6](#)
- [43] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *ACM SIGGRAPH Computer Graphics*, volume 21, pages 163–169, 1987. [6](#)
- [44] Wenjie Luo, Yujia Li, Raquel Urtasun, and Richard Zemel. Understanding the effective receptive field in deep convolutional neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4898–4906, 2016. [5](#)
- [45] 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 *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 6468–6477, 2020. [2](#), [3](#), [5](#), [6](#), [7](#)
- [46] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4460–4470, 2019. [2](#), [3](#)
- [47] 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)*, 2021. [3](#)
- [48] Alexandros Neophytou and Adrian Hilton. A layered model of human body and garment deformation. In *International Conference on 3D Vision (3DV)*, pages 171–178, 2014. [2](#), [3](#)
- [49] Ahmed A. A. Osman, Timo Bolkart, and Michael J. Black. STAR: Sparse trained articulated human body regressor. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 598–613, 2020. [1](#), [2](#), [5](#)
- [50] Junyi Pan, Xiaoguang Han, Weikai Chen, Jiapeng Tang, and Kui Jia. Deep mesh reconstruction from single RGB images via topology modification networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9963–9972, 2019. [3](#)
- [51] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019. [2](#), [3](#)
- [52] Despoina Paschalidou, Angelos Katharopoulos, Andreas Geiger, and Sanja Fidler. Neural Parts: Learning expressive 3D shape abstractions with invertible neural networks.

- In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [53] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. TailorNet: Predicting clothing in 3D as a function of human pose, shape and garment style. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 7363–7373, 2020. 2, 3, 5
- [54] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. 1
- [55] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 523–540, 2020. 3
- [56] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. ClothCap: Seamless 4D clothing capture and retargeting. *ACM Transactions on Graphics (TOG)*, 36(4):73, 2017. 2
- [57] Sergey Prokudin, Michael J. Black, and Javier Romero. SM-PLpox: Neural avatars from 3D human models. In *Winter Conference on Applications of Computer Vision (WACV)*, 2021. 5, 8
- [58] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2017. 7
- [59] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3D faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 725–741, 2018. 6
- [60] Anurag Ranjan, David T. Hoffmann, Dimitrios Tzionas, Siyu Tang, Javier Romero, and Michael J. Black. Learning multi-human optical flow. *International Journal of Computer Vision (IJCV)*, 128(4):873–890, 2020. 1
- [61] Anurag Ranjan, Javier Romero, and Michael J Black. Learning human optical flow. In *British Machine Vision Conference (BMVC)*, page 297, 2018. 1
- [62] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: modeling and capturing hands and bodies together. *ACM Transactions on Graphics (TOG)*, 36(6):245:1–245:17, 2017. 3
- [63] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, 2015. 5
- [64] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2304–2314, 2019. 3
- [65] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 84–93, 2020. 2, 3
- [66] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [67] Igor Santesteban, Miguel A. Otaduy, and Dan Casas. Learning-Based Animation of Clothing for Virtual Try-On. *Computer Graphics Forum*, 38(2):355–366, 2019. 2
- [68] Nicholas Sharp and Maks Ovsjanikov. PointTriNet: Learned triangulation of 3D point sets. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 762–778, 2020. 8
- [69] Garvita Tiwari, Bharat Lal Bhatnagar, Tony Tung, and Gerard Pons-Moll. SIZER: A dataset and model for parsing 3D clothing and learning size sensitive 3D clothing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 12348, pages 1–18, 2020. 2
- [70] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Carsten Stoll, and Christian Theobalt. PatchNets: Patch-based generalizable deep implicit 3D shape representations. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 293–309, 2020. 3
- [71] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 109–117, 2017. 1
- [72] Shaoifei Wang, Andreas Geiger, and Siyu Tang. Locally aware piecewise transformation fields for 3D human mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 8
- [73] Francis Williams, Teseo Schneider, Claudio Silva, Denis Zorin, Joan Bruna, and Daniele Panozzo. Deep geometric prior for surface reconstruction. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10130–10139, 2019. 3, 4
- [74] Donglai Xiang, Fabian Prada, Chenglei Wu, and Jessica K. Hodgins. MonoClothCap: Towards temporally coherent clothing capture from monocular RGB video. In *International Conference on 3D Vision (3DV)*, pages 322–332, 2020. 2
- [75] 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 *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 6184–6193, 2020. 1, 2
- [76] Qiangeng Xu, Weiye Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. DISN: Deep implicit surface network for high-quality single-view 3D reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 492–502, 2019. 3

- [77] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. PointFlow: 3D point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4541–4550, 2019. [2](#)
- [78] Jinlong Yang, Jean-Sébastien Franco, Franck Hétroy-Wheeler, and Stefanie Wuhrer. Analyzing clothing layer deformation statistics of 3D human motions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 237–253, 2018. [5](#)
- [79] Shan Yang, Zherong Pan, Tanya Amert, Ke Wang, Licheng Yu, Tamara Berg, and Ming C Lin. Physics-inspired garment recovery from a single-view image. *ACM Transactions on Graphics (TOG)*, 37(5):1–14, 2018. [2](#)
- [80] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. FoldingNet: Point cloud auto-encoder via deep grid deformation. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 206–215, 2018. [2](#), [3](#), [4](#)
- [81] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM Transactions on Graphics (TOG)*, 38(6), 2019. [5](#)
- [82] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. PCN: Point completion network. In *International Conference on 3D Vision (3DV)*, pages 728–737, 2018. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [83] Siwei Zhang, Yan Zhang, Qianli Ma, Michael J. Black, and Siyu Tang. PLACE: Proximity learning of articulation and contact in 3D environments. In *International Conference on 3D Vision (3DV)*, pages 642–651, 2020. [1](#)
- [84] Yongheng Zhao, Tolga Birdal, Haowen Deng, and Federico Tombari. 3D point capsule networks. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1009–1018, 2019. [2](#), [3](#), [4](#)
- [85] Heming Zhu, Yu Cao, Hang Jin, Weikai Chen, Dong Du, Zhangye Wang, Shuguang Cui, and Xiaoguang Han. Deep Fashion3D: A dataset and benchmark for 3D garment reconstruction from single images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 12346, pages 512–530, 2020. [3](#)