

# CloSET: Modeling Clothed Humans on Continuous Surface with Explicit Template Decomposition

Hongwen Zhang<sup>1</sup> Siyou Lin<sup>1</sup> Ruizhi Shao<sup>1</sup> Yuxiang Zhang<sup>1</sup> Zerong Zheng<sup>1</sup>  
 Han Huang<sup>2</sup> Yandong Guo<sup>2</sup> Yebin Liu<sup>1</sup>  
<sup>1</sup>Tsinghua University <sup>2</sup>OPPO Research Institute

## Abstract

Creating animatable avatars from static scans requires the modeling of clothing deformations in different poses. Existing learning-based methods typically add pose-dependent deformations upon a minimally-clothed mesh template or a learned implicit template, which have limitations in capturing details or hinder end-to-end learning. In this paper, we revisit point-based solutions and propose to decompose explicit garment-related templates and then add pose-dependent wrinkles to them. In this way, the clothing deformations are disentangled such that the pose-dependent wrinkles can be better learned and applied to unseen poses. Additionally, to tackle the seam artifact issues in recent state-of-the-art point-based methods, we propose to learn point features on a body surface, which establishes a continuous and compact feature space to capture the fine-grained and pose-dependent clothing geometry. To facilitate the research in this field, we also introduce a high-quality scan dataset of humans in real-world clothing. Our approach is validated on two existing datasets and our newly introduced dataset, showing better clothing deformation results in unseen poses. The project page with code and dataset can be found at <https://www.liuyebin.com/closet>.

## 1. Introduction

Animating 3D clothed humans requires the modeling of pose-dependent deformations in various poses. The diversity of clothing styles and body poses makes this task extremely challenging. Traditional methods are based on either simple rigging and skinning [4, 20, 33] or physics-based simulation [14, 23, 24, 49], which heavily rely on artist efforts or computational resources. Recent learning-based methods [10, 36, 39, 59] resort to modeling the clothing deformation directly from raw scans of clothed humans. Despite the promising progress, this task is still far from being solved due to the challenges in clothing representations, generalization to unseen poses, and data acquisition, *etc.*

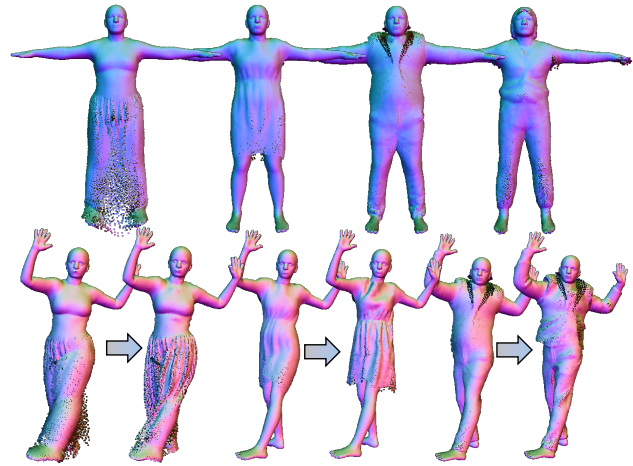


Figure 1. Our method learns to decompose garment templates (top row) and add pose-dependent wrinkles upon them (bottom row).

For the modeling of pose-dependent garment geometry, the representation of clothing plays a vital role in a learning-based scheme. As the relationship between body poses and clothing deformations is complex, an effective representation is desirable for neural networks to capture pose-dependent deformations. In the research of this line, meshes [2, 12, 38], implicit fields [10, 59], and point clouds [36, 39] have been adopted to represent clothing. In accordance with the chosen representation, the clothing deformation and geometry features are learned on top of a fixed-resolution template mesh [8, 38], a 3D implicit sampling space [10, 59], or an unfolded UV plane [2, 12, 36, 39]. Among these representations, the mesh is the most efficient one but is limited to a fixed topology due to its discretization scheme. The implicit fields naturally enable continuous feature learning in a resolution-free manner but are too flexible to satisfy the body structure prior, leading to geometry artifacts in unseen poses. The point clouds enjoy the compact nature and topology flexibility and have shown promising results in the recent state-of-the-art solutions [36, 39] to represent clothing, but the feature learning on UV planes still leads to discontinuity artifacts between body parts.

To model the pose-dependent deformation of clothing, body templates such as SMPL [34] are typically leveraged to account for articulated motions. However, a body template alone is not ideal, since the body template only models the minimally-clothed humans and may hinder the learning of actual pose-dependent deformations, especially in cases of loose clothing. To overcome this issue, recent implicit approaches [59] make attempts to learn skinning weights in the 3D space to complement the imperfect body templates. However, their pose-dependent deformations are typically coarse due to the difficulty in learning implicit fields. For explicit solutions, the recent approach [32] suggests learning coarse templates implicitly at first and then the pose-dependent deformations explicitly. Despite its effectiveness, such a workaround requires a two-step modeling procedure and hinders end-to-end learning.

In this work, we propose CloSET, an end-to-end method to tackle the above issues by modeling Clothed humans on a continuous Surface with Explicit Template decomposition. We follow the spirit of recent state-of-the-art point-based approaches [32, 36, 39] as they show the efficiency and potential in modeling real-world garments. We take steps forward in the following aspects for better point-based modeling of clothed humans. First, we propose to decompose the clothing deformations into explicit garment templates and pose-dependent wrinkles. Specifically, our method learns a garment-related template and adds the pose-dependent displacement upon them, as shown in Fig. 1. Such a garment-related template preserves a shared topology for various poses and enables better learning of pose-dependent wrinkles. Different from the recent solution [32] that needs two-step procedures, our method can decompose the explicit templates in an end-to-end manner with more garment details. Second, we tackle the seam artifact issues that occurred in recent point-based methods [36, 39]. Instead of using unfolded UV planes, we propose to learn point features on a body surface, which supports a continuous and compact feature space. We achieve this by learning hierarchical point-based features on top of the body surface and then using barycentric interpolation to sample features continuously. Compared to feature learning in the UV space [39], on template meshes [8, 38], or in the 3D implicit space [57, 59], our body surface enables the network to capture not only fine-grained details but also long-range part correlations for pose-dependent geometry modeling. Third, we introduce a new scan dataset of humans in real-world clothing, which contains more than 2,000 high-quality scans of humans in diverse outfits, hoping to facilitate the research in this field. The main contributions of this work are summarized below:

- We propose a point-based clothed human modeling method by decomposing clothing deformations into explicit garment templates and pose-dependent wrinkles

in an end-to-end manner. These learnable templates provide a garment-aware canonical space so that pose-dependent deformations can be better learned and applied to unseen poses.

- We propose to learn point-based clothing features on a continuous body surface, which allows a continuous feature space for fine-grained detail modeling and helps to capture long-range part correlations for pose-dependent geometry modeling.
- We introduce a new high-quality scan dataset of clothed humans in real-world clothing to facilitate the research of clothed human modeling and animation from real-world scans.

## 2. Related Work

**Representations for Modeling Clothed Humans.** A key component in modeling clothed humans is the choice of representation, which mainly falls into two categories: implicit and explicit representations.

*Implicit Modeling.* Implicit methods [3, 10–12, 15, 21, 29, 41, 42, 48, 59, 67, 70, 77] represent surfaces as the level set of an implicit neural scalar field. Recent state-of-the-art methods typically learn the clothing deformation field with a canonical space decomposition [9, 30, 45, 59, 65] or part-based modeling strategies [16, 25, 46, 55, 75]. Compared to mesh templates, implicit surfaces are not topologically constrained to specific templates [57, 58], and can model various clothes with complex topology. However, the learning space of an implicit surface is the whole 3D volume, which makes training and interpolation difficult, especially when the numbers of scan data are limited.

*Explicit Modeling.* Mesh surfaces, the classic explicit representation, currently dominate the field of 3D modeling [6–8, 13, 23, 24, 26–28, 38, 43, 49, 60, 62, 64, 66, 69, 72] with their compactness and high efficiency in downstream tasks such as rendering, but they are mostly limited to a fixed topology and/or require scan data registered to a template. Thus, mesh-based representations forbid the learning of a universal model for topologically varying clothing types. Though some approaches have been proposed to allow varying mesh topology [44, 47, 63, 68, 79], they are still limited in their expressiveness. Point clouds enjoy both compactness and topological flexibility. Previous work generates sparse point clouds for 3D representation [1, 19, 31, 74]. However, the points need to be densely sampled over the surface to model surface geometry accurately. Due to the difficulty of generating a large point set, recent methods group points into patches [5, 17, 18, 22]. Each patch maps the 2D UV space to the 3D space, allowing arbitrarily dense sampling within this patch. SCALE [36] successfully applies this idea to modeling clothed humans, but produces notable discontinuity artifacts near patch boundaries.

POP [39] further utilizes a single fine-grained UV map for the whole body surface, leading to a more topologically flexible representation. However, the discontinuity of the UV map may lead to seam artifacts in POP. Very recently, FITE [32] suggests learning implicit coarse templates [76] at first and then explicit fine details. Despite the efficacy, it requires a two-step modeling procedure. Concurrently, SkiRT [37] proposes to improve the body template by learning the blend skinning weights with several data terms. In contrast, our method applies regularization to achieve template decomposition in an end-to-end manner and learns the point-based pose-dependent displacement more effectively.

**Pose-dependent Deformations for Animation.** In the field of character animation, traditional methods utilize rigging and skinning techniques to repose characters [4, 20, 33, 34, 50], but they fail to model realistic pose-dependent clothing deformations such as wrinkles and sliding motions between clothes and body. We conclude two key ingredients in modeling pose-dependent clothing: (i) pose-dependent feature learning; (ii) datasets with realistic clothing.

*Pose-dependent Feature Learning.* Some traditional methods directly incorporate the entire pose parameters into the model [28, 38, 49, 71]. Such methods easily overfit on pose parameters and introduce spurious correlations, causing bad generalization to unseen poses. Recent work explores poses conditioning with local features, either with point clouds [36, 39] or implicit surfaces [59, 67], and shows superiority in improving geometry quality and in eliminating spurious correlations. Among them, the most relevant to ours is POP [39], which extracts local pose features by utilizing convolution on a UV position map. Despite its compelling performance, POP suffers from artifacts inherent to its UV-based representation, hence the convolution on the UV map produces discontinuity near UV islands’ boundaries [39]. We address this issue by discarding the UV-based scheme and returning to the actual 3D body surface. We attach the features to a uniform set of points on a T-posed body template, and process them via a PointNet++ [53] structure for pose-dependent modeling. As validated by our experiments, our pose embedding method leads to both qualitative and quantitative improvements.

*Clothed Human Datasets.* Another challenging issue for training an animatable avatar is the need for datasets of clothed humans in diverse poses. There are considerable efforts seeking to synthesize clothed datasets with physics-based simulation [14, 23, 24, 39, 49]. Although they are diverse in poses, there remains an observable domain gap between synthetic clothes and real data. Acquiring clothed scans with realistic details of clothing deformations [38, 51, 56, 73, 78] is crucial for the development of learning-based methods in this field.

### 3. Method

As illustrated in Fig. 2, the proposed method CloSET learns garment-related and pose-dependent features on body surfaces (see Sec. 3.1), which can be sampled in a continuous manner and fed into two decoders for the generation of explicit garment templates and pose-dependent wrinkles (see Sec. 3.2).

#### 3.1. Continuous Surface Features

As most parts of the clothing are deformed smoothly in different poses, a continuous feature space is desirable to model the garment details and pose-dependent garment geometry. To this end, our approach first learns features on top of a body template surface, *i.e.*, a SMPL [34] or SMPL-X [50] model in a T-pose. Note that these features are not limited to those on template vertices as they can be continuously sampled from the body surface via barycentric interpolation. Hence, our feature space is more continuous than UV-based [2, 39] spaces, while being more compact than 3D implicit feature fields [59, 67].

To model pose-dependent clothing deformations, the underlying unclothed body model is taken as input to the geometry feature encoder. For each scan, let  $\mathbf{V}^u = \{\mathbf{v}_n^u\}_{n=1}^N$  denote the posed vertex positions of the fitted unclothed body model, where  $N = 6890$  for SMPL [34] and  $N = 10475$  for SMPL-X [50]. These posed vertices act as the pose code and will be paired with the template vertices  $\mathbf{V}^t = \{\mathbf{v}_n^t\}_{n=1}^N$  of the body model in a T-pose, which shares the same mesh topology with  $\mathbf{V}^u$ . These point pairs are processed by the pose encoder  $\mathcal{F}_p$  to generate the pose-dependent geometry features  $\{\phi_p(\mathbf{v}_n^t) \in \mathbb{R}^{C_p}\}_{n=1}^N$  at vertices  $\mathbf{V}^t$ , *i.e.*,

$$\{\phi_p(\mathbf{v}_n^t) \in \mathbb{R}^{C_p}\}_{n=1}^N = \mathcal{F}_p(\mathbf{V}^t, \mathbf{V}^u). \quad (1)$$

To learn hierarchical features with different levels of receptive fields, we adopt PointNet++ [53] as the architecture of the pose encoder  $\mathcal{F}_p$ , where vertices  $\mathbf{V}^t$  are treated as the input point cloud in the PointNet++ network, while vertices  $\mathbf{V}^u$  act as the feature of  $\mathbf{V}^t$ . As the template vertices  $\mathbf{V}^t$  are constant, the encoder  $\mathcal{F}_p$  can focus on the feature learned from the posed vertices  $\mathbf{V}^u$ . Moreover, the PointNet++ based  $\mathcal{F}_p$  first abstracts features from the template vertices  $\mathbf{V}^t$  to sparser points  $\{\mathbf{V}_l^t\}_{l=1}^L$  at  $L$  levels, where the number of  $\mathbf{V}_l^t$  decreases with  $l$  increasing. Then, the features at  $\{\mathbf{V}_l^t\}_{l=1}^L$  are further propagated back to  $\mathbf{V}^t$  successively. In this way, the encoder can capture the long-range part correlations of the pose-dependent deformations.

Similar to POP [39], our method can be trained under multi-outfit or outfit-specific settings. When trained with multiple outfits, the pose-dependent deformation should be aware of the outfit type, and hence requires the input of the garment features. Specifically, the garment-related features  $\{\phi_g(\mathbf{v}_n^t) \in \mathbb{R}^{C_g}\}_{n=1}^N$  are also defined on template

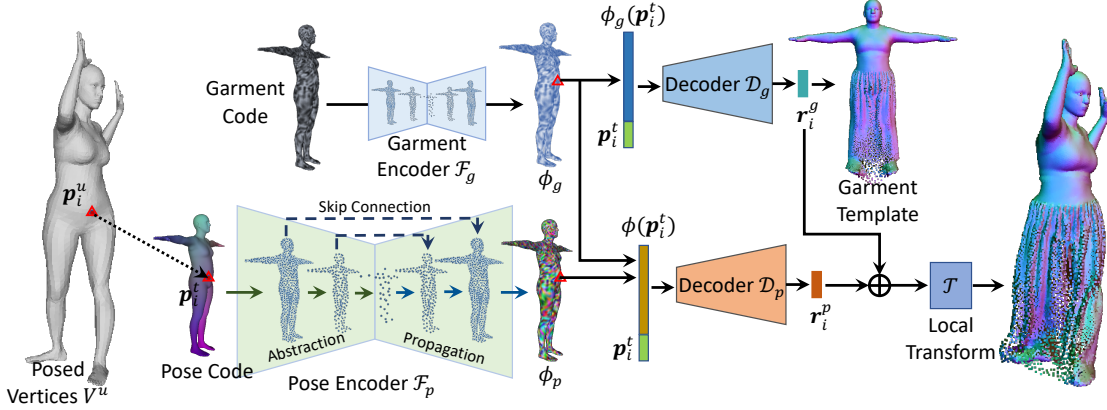


Figure 2. Overview of the proposed method CloSET. Given an input body model, its pose code and garment code are processed hierarchically by point-based pose and garment encoders  $\mathcal{F}_p$  and  $\mathcal{F}_g$  for the learning of surface features  $\phi_p$  and  $\phi_g$ . For any point  $p_i^t$  lying on the template surface, its features  $\phi(p_i^t)$  are sampled from surface features accordingly and fed into two decoders for the prediction of the explicit garment template and pose-dependent wrinkle displacements, which will be combined and transformed to the clothing point cloud.

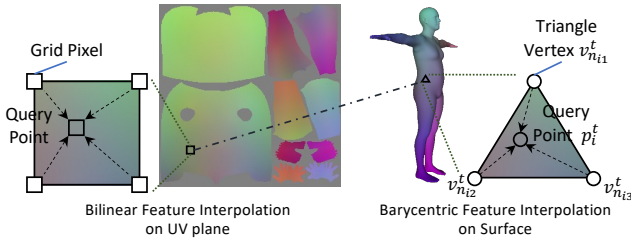


Figure 3. Comparison of the bilinear interpolation on the UV plane and the barycentric interpolation on the surface.

vertices  $\mathbf{V}^t$ , which are learned by feeding the garment code  $\{\phi_{gc}(\mathbf{v}_n^t)\}_{n=1}^N$  to a smaller PointNet++ encoder  $\mathcal{F}_g$ . Note that the garment-related features  $\phi_g$  are shared for each outfit across all poses and optimized during the training. Since both the pose-dependent and garment-related geometry features are aligned with each other, we denote them as the surface features  $\{\phi(\mathbf{v}_n^t)\}_{n=1}^N$  for simplicity. Note that the input of  $\phi_g(p_i^t)$  has no side effect on the results when trained with only one outfit, as the garment features are invariant to the input poses.

**Continuous Feature Interpolation.** In the implicit modeling solutions [42, 57, 59], features are learned in a spatially continuous manner, which contributes to the fine-grained modeling of clothing details. To sample continuous features in our scheme, we adopt barycentric interpolation on the surface features  $\phi$ . As illustrated in Fig. 3, for any point  $p_i^t = \mathbf{V}^t(\mathbf{b}_i)$  lying on the template surface, where  $\mathbf{b}_i = [n_{i1}, n_{i2}, n_{i3}, b_{i1}, b_{i2}, b_{i3}]$  denotes the corresponding vertex indices and barycentric coordinates in  $\mathbf{V}^t$ . Then, the corresponding surface features can be retrieved via barycentric interpolation, *i.e.*,

$$\phi(p_i^t) = \sum_{j=1}^3 (b_{ij} * \phi(\mathbf{v}_{n_{ij}}^t)). \quad (2)$$

In this way, the point features are not limited to those learned on the template vertices and are continuously defined over the whole body surface without the seaming discontinuity issue in the UV plane.

### 3.2. Point-based Clothing Deformation

Following previous work [36, 39], our approach represents the clothed body as a point cloud. For any point on the surface of the unclothed body model, the corresponding features are extracted from surface features to predict its displacement and normal vector.

**Explicit Template Decomposition.** Instead of predicting the clothing deformation directly, our method decomposes the deformations into two components: garment-related template displacements and pose-dependent wrinkle displacements. To achieve this, the garment-related template is learned from the garment-related features and shared across all poses. Meanwhile, the learning of pose-dependent wrinkles are conditioned on both garment-related and pose-dependent features. Specifically, for the point  $p_i^u = \mathbf{V}^u(\mathbf{b}_i)$  at the unclothed body mesh, it has the same vertex indices and barycentric coordinates  $\mathbf{b}_i$  as the point  $p_i^t$  on the template surface. The pose-dependent and garment-related features of the point  $p_i^u$  are first sampled according to Eq. (2) based on  $p_i^t = \mathbf{V}^t(\mathbf{b}_i)$  and then further fed into the garment decoder  $\mathcal{D}_g$  and the pose decoder  $\mathcal{D}_p$  for displacement predictions, *i.e.*,

$$\begin{aligned} \mathbf{r}_i^g &= \mathcal{D}_g(\phi_g(p_i^t), p_i^t), \\ \mathbf{r}_i^p &= \mathcal{D}_p(\oplus(\phi_g(p_i^t), \phi_p(p_i^t)), p_i^t), \end{aligned} \quad (3)$$

where  $\oplus$  denotes the concatenation operation,  $\mathbf{r}_i^g$  and  $\mathbf{r}_i^p$  are the displacements for garment templates and pose-dependent wrinkles, respectively. Finally,  $\mathbf{r}_i^g$  and  $\mathbf{r}_i^p$  will be added together as the clothing deformation  $\mathbf{r}_i = \mathbf{r}_i^g + \mathbf{r}_i^p$ .

**Local Transformation.** Similar to [36,39], the displacement  $\mathbf{r}_i$  is learned in a local coordinate system. It is further transformed to the world coordinate system by applying the following transformation, *i.e.*,  $\mathbf{x}_i = \mathcal{T}_i \mathbf{r}_i + \mathbf{p}_i^u$ , where  $\mathcal{T}_i$  denotes the local transformation calculated based on the unclothed body model. Following [36,39], the transformation matrix  $\mathcal{T}_i$  is defined at the point  $\mathbf{p}_i^u$  on the unclothed body model, which naturally supports the barycentric interpolation. Similarly, the normal  $\mathbf{n}_i$  of each point is predicted together with  $\mathbf{r}_i$  from  $\mathcal{D}$  and transformed by  $\mathcal{T}_i$ .

### 3.3. Loss Functions

Following previous work [36,39], the point-based clothing deformation is learned with the summation of loss functions:  $\mathcal{L} = \mathcal{L}_{data} + \lambda_{rgl} \mathcal{L}_{rgl}$ . where  $\mathcal{L}_{data}$  and  $\mathcal{L}_{rgl}$  denote the data and regularization terms respectively, and the weight  $\lambda_{rgl}$  balances the loss terms.

**Data Term.** The data term  $\mathcal{L}_{data}$  is calculated on the final predicted points and normals, *i.e.*,  $\mathcal{L}_{data} = \lambda_p \mathcal{L}_p + \lambda_n \mathcal{L}_n$ . Specifically,  $\mathcal{L}_p$  is the normalized Chamfer distance to minimize the bi-directional distances between the point sets of the prediction and the ground-truth scan:

$$\mathcal{L}_p = Chamfer \left( \{\mathbf{x}_i\}_{i=1}^M, \{\hat{\mathbf{x}}_j\}_{j=1}^{N_s} \right) = \frac{1}{M} \sum_{i=1}^M \min_j \|\mathbf{x}_i - \hat{\mathbf{x}}_j\|_2^2 + \frac{1}{N_s} \sum_{j=1}^{N_s} \min_i \|\mathbf{x}_i - \hat{\mathbf{x}}_j\|_2^2, \quad (4)$$

where  $\hat{\mathbf{x}}_j$  is the point sampled from the ground-truth surface,  $M$  and  $N_s$  denote the number of the predicted and ground-truth points, respectively.

The normal loss  $\mathcal{L}_n$  is an averaged  $L1$  distance between the normal of each predicted point and its nearest ground-truth counterpart:

$$\mathcal{L}_n = L1 \left( \{\mathbf{n}_i\}_{i=1}^M, \{\hat{\mathbf{n}}_i\}_{i=1}^M \right) = \frac{1}{M} \sum_{i=1}^M \|\mathbf{n}_i - \hat{\mathbf{n}}_i\|, \quad (5)$$

where  $\hat{\mathbf{n}}_i$  is the normal of its nearest point in the ground-truth point set.

Note that we do not apply data terms on the garment templates, as we found such a strategy leads to noisy template learning in our experiments.

**Regularization Term.** The regularization terms are added to prevent the predicted deformations from being extremely large and regularize the garment code. Moreover, following the previous implicit template learning solution [30,76], we also add regularization on the pose-dependent displacement  $\mathbf{r}_i^p$  to encourage it to be as small as possible. As the pose-dependent displacement represents the clothing deformation in various poses, such a regularization implies that the pose-invariant deformation should be retained in the template displacement  $\mathbf{r}_i^g$ , which forms

the garment-related template shared by all poses. Overall, the regularization term can be written as follows:

$$\mathcal{L}_{rgl} = \frac{1}{M} \sum_{i=1}^M \|\mathbf{r}_i\|_2^2 + \frac{\lambda_{pd}}{M} \sum_{i=1}^M \|\mathbf{r}_i^p\|_2^2 + \frac{\lambda_{gc}}{N} \sum_{n=1}^N \|\phi_{gc}(\mathbf{v}_n^t)\|_2^2. \quad (6)$$

## 4. Experiments

**Network Architecture.** For a fair comparison with POP [39], we modify the official PointNet++ [53] (PN++) architecture so that our encoders have comparable network parameters as POP [39]. The modified PointNet++ architecture has 6 layers for feature abstraction and 6 layers for feature propagation (*i.e.*,  $L = 6$ ). Since the input point cloud  $\mathbf{V}^t$  has constant coordinates, the farthest point sampling in PointNet++ is only performed at the first forward process, and the sampling indices are saved for the next run. In this way, the runtime is significantly reduced for both training and inference so that our pose and garment encoders can have similar network parameters and runtime speeds to POP. Note that the pose and garment encoders in our method can also be replaced with recent state-of-the-art point-based encoders such as PointMLP [40] and PointNeXt [54]. More details about the network architecture and implementation can be found in the Supp.Mat.

**Datasets.** We use CAPE [38], ReSynth [39], and our newly introduced dataset THuman-CloSET for training and evaluation.

**CAPE** [38] is a captured human dataset consisting of multiple humans in various motions. The outfits in this dataset mainly include common clothing such as T-shirts. We follow SCALE [36] to choose blazerlong (with outfits of blazer jacket and long trousers) and shortlong (with outfits of short T-shirt and long trousers) from subject 03375 to validate the efficacy of our method.

**ReSynth** [39] is a synthetic dataset introduced in POP [39]. It is created by using physics simulation, and contains challenging outfits such as skirts and jackets. We use the official training and test split as [39].

**THuman-CloSET** is our newly introduced dataset, containing high-quality clothed human scans captured by a dense camera rig. We introduce THuman-CloSET for the reason that existing pose-dependent clothing datasets [38,39] are with either relatively tight clothing or synthetic clothing via physics simulation. In THuman-CloSET, there are more than 2,000 scans of 15 outfits with a large variation in clothing style, including T-shirts, pants, skirts, dresses, jackets, and coats, to name a few. For each outfit, the subject is guided to perform different poses by imitating the poses in CAPE. Moreover, each subject has a scan with minimal clothing in A-pose. THuman-CloSET contains well-fitted body models in the form of SMPL-X [50]. Note that

Table 1. Quantitative comparison with previous point-based methods on ReSynth. † denotes the methods using 1/8 training data.

Method	Chamfer- $L_2$ ↓		Normal diff. ↓	
	Mean	Max	Mean	Max
<b>outfit-specific</b>				
SCALE [36]	1.491	8.451	1.041	1.321
POP [39]	1.356	7.339	<b>1.013</b>	<b>1.289</b>
<b>multi-outfit</b>				
Baseline (POP [39]) †	1.490	7.859	1.050	1.326
Baseline w. PN++ †	1.290	5.940	1.028	1.330
CloSET (Ours) †	<b>1.240</b>	<b>5.543</b>	1.019	1.315

Table 2. Quantitative comparison of different methods on the proposed THuman-CloSET dataset in the outfit-specific setting.

Subject ID	SCANimate		SNARF		POP		CloSET	
	CD	NML	CD	NML	CD	NML	CD	NML
sweater-000	1.06	1.64	7.11	2.09	0.76	1.55	<b>0.68</b>	<b>1.48</b>
longshirt-001	1.42	1.85	6.66	2.21	1.54	1.83	<b>1.39</b>	<b>1.71</b>
skirt-005	1.93	1.74	9.39	2.31	1.66	1.43	<b>1.49</b>	<b>1.36</b>

the loose clothing makes the fitting of the underlying body models quite challenging. For more accurate fitting of the body models, we first fit a SMPL-X model on the scan of the subject in minimal clothing and then adopt its shape parameters for fitting the outfit scans in different poses. More details can be found in the Supp.Mat. In our experiments, we use the outfit scans in 100 different poses for training and use the remaining poses for evaluation. We hope our new dataset can open a promising direction for clothed human modeling and animation from real-world scans.

**Metrics.** Following previous work [39], we generate 50K points from our method and point-based baselines and adopt the Chamfer Distance (see Eq. (4)) and the  $L_1$  normal discrepancy (see Eq. (5)) for quantitative evaluation. By default, the Chamfer distance (CD) and normal discrepancy (NML) are reported in the unit of  $\times 10^{-4}m^2$  and  $\times 10^{-1}$ , respectively. To evaluate the implicit modeling methods, the points are sampled from the surface extracted using Marching Cubes [35].

#### 4.1. Comparison with the State-of-the-art Methods

We compare results with recent state-of-the-art methods, including point-based approaches SCALE [36], POP [39], and SkiRT [37], and implicit approaches SCANimate [59] and SNARF [10].

**ReSynth.** Tab. 1 reports the results of the pose-dependent clothing predictions on unseen motion sequences from the ReSynth [39] dataset, where all 12 outfits are used for eval-

Table 3. Quantitative comparison of different methods on the ReSynth dataset in the outfit-specific setting. The garment styles are non-skirt, skirt, and dress for carla-004, christine-027, and felice-004, respectively.

Subject ID	SCANimate		POP		SkiRT		CloSET	
	CD	NML	CD	NML	CD	NML	CD	NML
carla-004	0.90	1.52	0.51	<b>1.02</b>	<b>0.48</b>	1.06	0.49	1.04
christine-027	3.21	1.66	1.72	0.97	1.54	0.99	<b>1.49</b>	<b>0.97</b>
felice-004	20.79	2.94	7.34	1.24	6.45	1.25	<b>6.01</b>	<b>1.16</b>

uation. As can be seen, the proposed approach has the lowest mean and max errors, which outperforms all other approaches including POP [39]. Note that our approach needs fewer data for the pose-dependent deformation modeling. By using only 1/8 data, our approach achieves a performance comparable to or even better than other models trained with full data. Tab. 3 also reports outfit-specific performances on 3 selected subject-outfit types, including jackets, skirts, and dresses. In comparison with the recent state-of-the-art method SkiRT [37], our method achieves better results on challenging skirt/dress outfits and comparable results on non-skirt clothing.

**THuman-CloSET.** The effectiveness of our method is also validated on our real-world THuman-CloSET dataset. The sparse training poses and loose clothing make this dataset very challenging for clothed human modeling. Tab. 2 reports the quantitative comparisons of different methods on three representative outfits. Fig. 4 also shows example results of different methods, where we follow previous work [39] to obtain meshed results via Poisson surface reconstruction. We can see that our method generalizes better to unseen poses and produces more natural pose-dependent wrinkles than other methods. In our experiments, we found that SNARF [10] fails to learn correct skinning weights due to loose clothing and limited training poses. As discussed in FITE [32], there is an ill-posed issue of jointly optimizing the canonical shape and the skinning fields, which becomes more severe in our dataset.

#### 4.2. Ablation Study

**Evaluation of Continuous Surface Features.** The point features in our method are learned on the body surface, which provides a continuous and compact feature learning space. To validate this, Tab. 4 summarizes the feature learning space of different approaches and their performances on two representative outfits from CAPE [38]. Here, we only include the proposed Continuous Surface Features (CSF) in Tab. 4 by applying the continuous features on POP [39] for fair comparisons with SCALE [36] and POP [39]. As discussed previously, existing solutions learn features either in a discontinuous space (*e.g.*, CAPE [38]) on the fixed resolu-

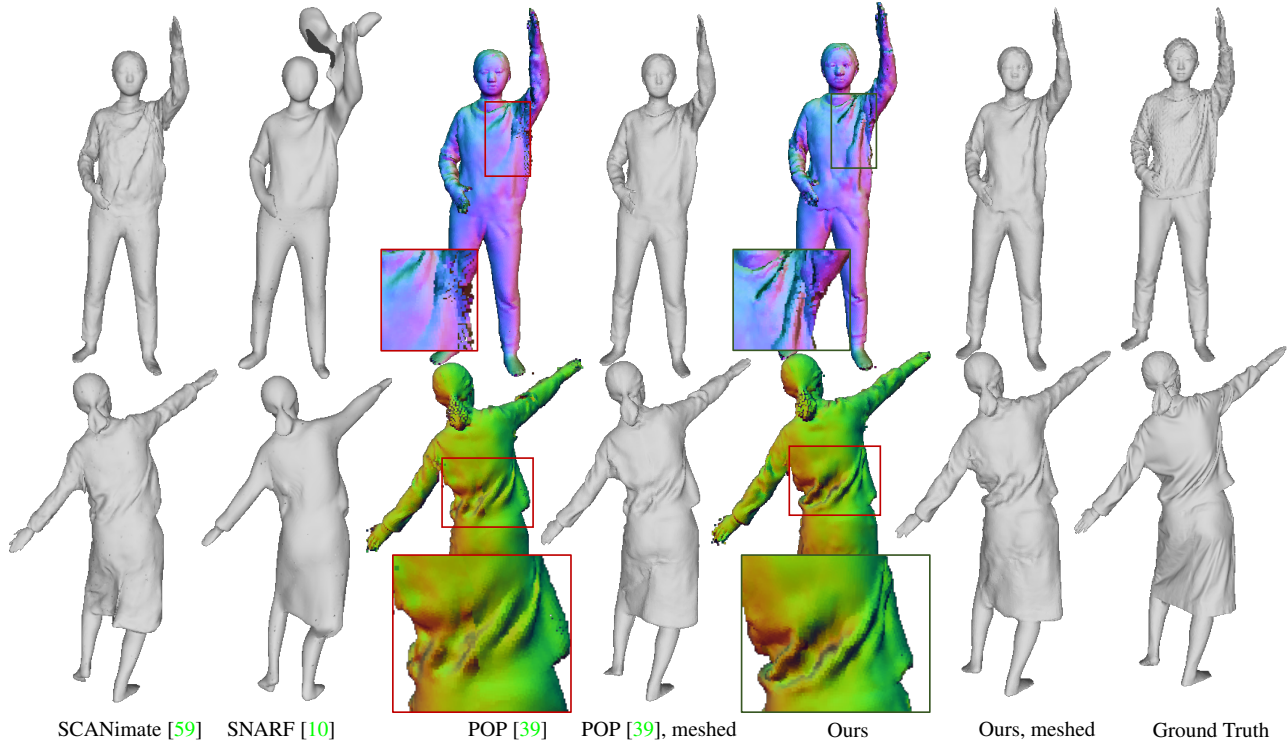


Figure 4. Comparison of different clothed human modeling methods on the proposed real-world scan dataset.

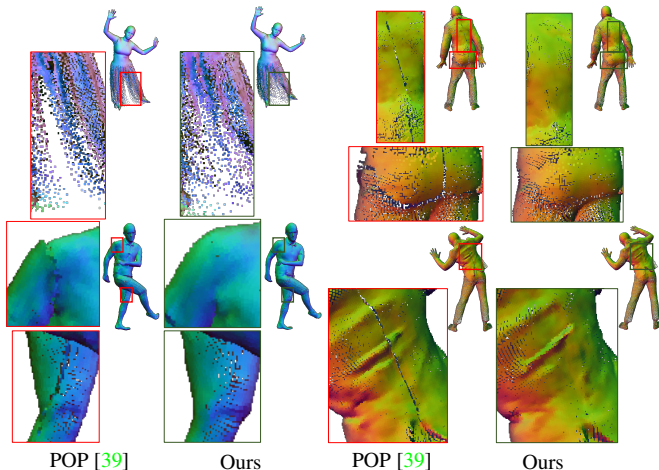


Figure 5. Comparison of the approach learned on UV planes (POP [39]) and the approach learned on continuous surfaces (Ours). Our solution alleviates the seam artifacts of POP.

tion mesh, POP [39] on the 2D UV plane) or in a space that is too flexible (e.g., NASA [16] in the implicit 3D space), while our approach learns features on continuous and compact surface space. Though SCALE [36] has also investigated using the point-based encoder (PointNet [52]) for pose-dependent feature extraction, it only uses the global features which lack fine-grained information. In contrast, we adopt PointNet++ [53] (PN++) to learn hierarchical surface features, so that the pose-dependent features can be learned more effectively. Fig. 5 shows the qualitative re-

Table 4. Comparison of the modeling ability of different approaches and their feature learning space on the CAPE dataset.

Methods	Features	Chamfer- $L_2$ ↓		Normal diff. ↓	
		<i>blazerlong</i>	<i>shortlong</i>	<i>blazerlong</i>	<i>shortlong</i>
CAPE [38]	Mesh	1.96	1.37	1.28	1.15
NASA [16]	3D space	1.37	0.95	1.29	1.17
SCALE [36]	Global	1.46	1.03	1.34	1.16
SCALE [36]	UV plane	1.07	0.89	1.22	1.12
POP [39]	UV plane	0.78	0.57	1.29	1.24
CSF	Surface	<b>0.71</b>	<b>0.54</b>	<b>1.15</b>	<b>1.09</b>

Table 5. Ablation study on the effectiveness of continuous surface features (CSF) and Explicit Template Decomposition (ETD) on a dress outfit (felice-004 from ReSynth [39]).

Method	POP	POP + ETD	CSF	CSF + ETD
CD	7.34	7.05	6.53	<b>6.01</b>
NML	1.24	1.17	1.16	<b>1.16</b>

sults of the ablation approaches learned on UV planes and continuous surfaces. We can see that our solution clearly alleviates the seam artifacts of POP.

**Evaluation of Explicit Template Decomposition.** The decomposed templates help to capture more accurate pose-dependent deformations and produce more natural wrinkles in unseen poses, especially for outfits that differ largely from the body template. To validate the effectiveness of our decomposition strategy, Tab. 5 reports the ablation experiments on a dress outfit (felice-004) of the ReSynth [39]

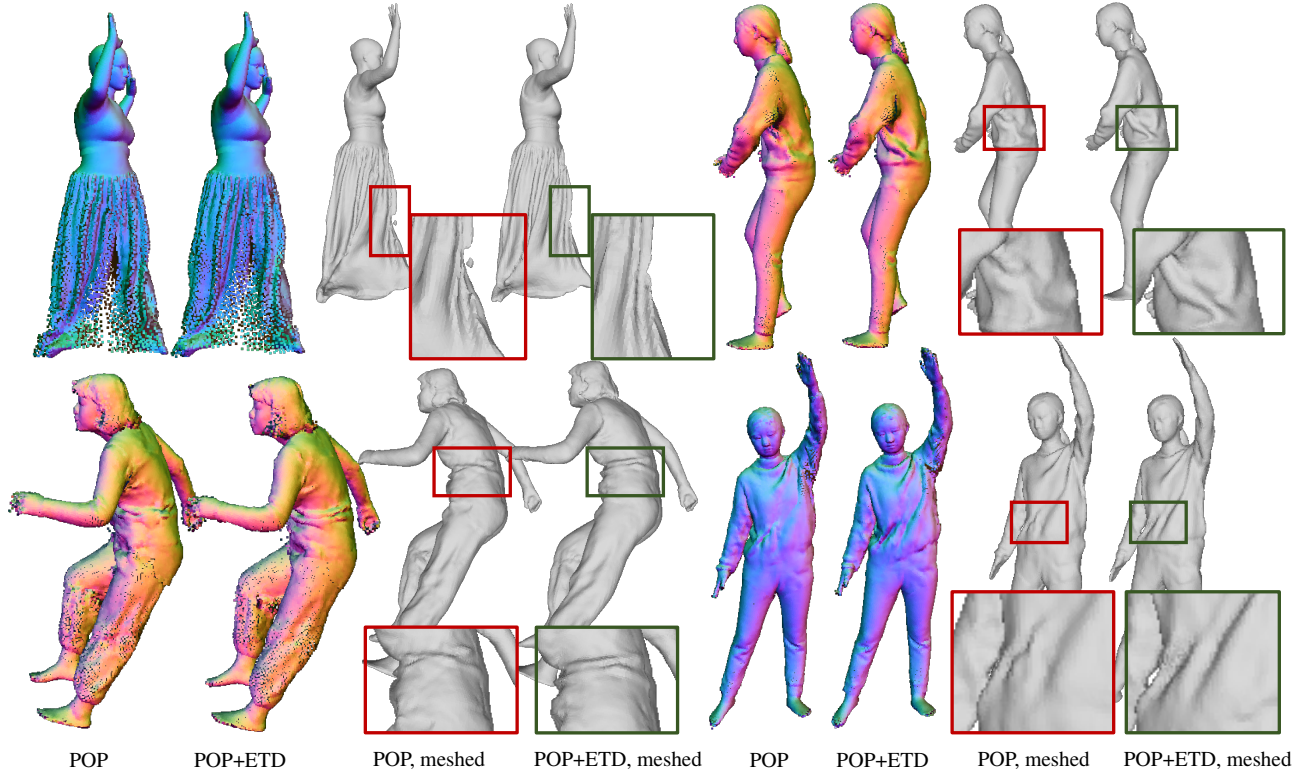


Figure 6. Comparison of the clothing deformation in unseen poses. Explicit template decomposition (ETD) helps to capture more natural pose-dependent wrinkle details than POP [39].

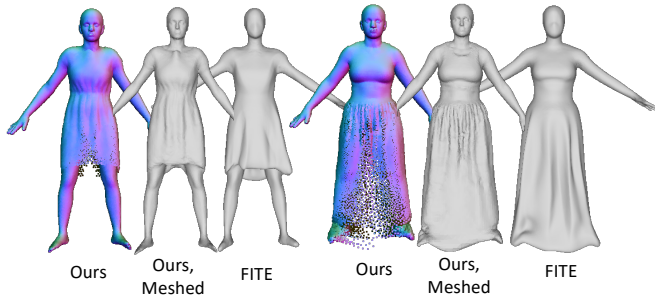


Figure 7. Comparison of the learned templates with FITE [32].

dataset. We can see that the proposed Explicit Template Decomposition (ETD) brings clear performance gains over baseline methods. Fig. 6 shows the visual improvements of pose-dependent wrinkles when applying the explicit template decomposition. Note that the templates are decomposed in an end-to-end manner in our method. Compared with the implicit template learned in the recent approach FITE [32], the explicit templates in our method contain more details, as shown in Fig. 7.

## 5. Conclusions and Future Work

In this work, we present CloSET, a point-based clothed human modeling method that is built upon a continuous surface and learns to decompose explicit garment templates for better learning of pose-dependent deformations. By learn-

ing features on a continuous surface, our solution gets rid of the seam artifacts in previous state-of-the-art point-based methods [36, 39]. Moreover, the explicit template decomposition helps to capture more accurate and natural pose-dependent wrinkles. To facilitate the research in this direction, we also introduce a high-quality real-world scan dataset with diverse outfit styles and accurate body model fitting.

**Limitations and Future Work.** Due to the incorrect skinning weight used in our template, the issue of the non-uniform point distribution remains for the skirt and dress outfits. Combining our method with recent learnable skinning solutions [37, 59] could alleviate this issue and further improve the results. Currently, our method does not leverage information from adjacent poses. Enforcing temporal consistency and correspondences between adjacent frames would be interesting for future work. Moreover, incorporating physics-based losses into the learning process like SNUG [61] would also be a promising solution to address the artifacts like self-intersections.

**Acknowledgements.** This work was supported by the National Key R&D Program of China (2022YFF0902200), the National Natural Science Foundation of China (No.62125107 and No.61827805), and the China Postdoctoral Science Foundation (No.2022M721844).



## References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3D point clouds. In *ICML*, pages 40–49, 2018. 2
- [2] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *ICCV*, pages 2293–2303, 2019. 1, 3
- [3] Ziqian Bai, Timur Bagautdinov, Javier Romero, Michael Zollhöfer, Ping Tan, and Shunsuke Saito. AutoAvatar: Autoregressive neural fields for dynamic avatar modeling. *ECCV*, 2022. 2
- [4] Ilya Baran and Jovan Popović. Automatic rigging and animation of 3D characters. *ACM TOG*, 26(3):72–es, 2007. 1, 3
- [5] Jan Bednařík, Shaifali Parashar, Erhan Gundogdu, Mathieu Salzmann, and Pascal Fua. Shape reconstruction by learning differentiable surface representations. In *CVPR*, pages 4715–4724, 2020. 2
- [6] Hugo Bertiche, Meysam Madadi, Emilio Tylson, and Sergio Escalera. DeepPSD: Automatic deep skinning and pose space deformation for 3D garment animation. In *ICCV*, pages 5471–5480, 2021. 2
- [7] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-Garment Net: Learning to dress 3D people from images. In *ICCV*, pages 5420–5430, 2019. 2
- [8] Andrei Burov, Matthias Nießner, and Justus Thies. Dynamic surface function networks for clothed human bodies. In *ICCV*, pages 10754–10764, October 2021. 1, 2
- [9] Xu Chen, Tianjian Jiang, Jie Song, Jinlong Yang, Michael J Black, Andreas Geiger, and Otmar Hilliges. gDNA: Towards generative detailed neural avatars. In *CVPR*, pages 20427–20437, 2022. 2
- [10] Xu Chen, Yufeng Zheng, Michael J. Black, Otmar Hilliges, and Andreas Geiger. SNARF: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, pages 11594–11604, October 2021. 1, 2, 6, 7
- [11] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. In *NeurIPS*, pages 21638–21652, 2020. 2
- [12] Enric Corona, Albert Pumarola, Guillem Alenya, Gerard Pons-Moll, and Francesc Moreno-Noguer. SMPLicit: Topology-aware generative model for clothed people. In *CVPR*, pages 11875–11885, 2021. 1, 2
- [13] Edilson De Aguiar, Leonid Sigal, Adrien Treuille, and Jessica K Hodgins. Stable spaces for real-time clothing. In *ACM TOG*, volume 29, page 106. ACM, 2010. 2
- [14] Deform Dynamics. <https://deformdynamics.com/>. 1, 3
- [15] Boyang Deng, JP Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Neural articulated shape approximation. In *ECCV*, pages 612–628, 2020. 2
- [16] Boyang Deng, John P Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Nasa neural articulated shape approximation. In *ECCV*, pages 612–628. Springer, 2020. 2, 7
- [17] Zhantao Deng, Jan Bednařík, Mathieu Salzmann, and Pascal Fua. Better patch stitching for parametric surface reconstruction. In *3DV*, pages 593–602, 2020. 2
- [18] Theo Deprelle, Thibault Groueix, Matthew Fisher, Vladimir Kim, Bryan Russell, and Mathieu Aubry. Learning elementary structures for 3D shape generation and matching. In *NeurIPS*, pages 7433–7443, 2019. 2
- [19] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3D object reconstruction from a single image. In *CVPR*, pages 2463–2471, 2017. 2
- [20] Andrew Feng, Dan Casas, and Ari Shapiro. Avatar reshaping and automatic rigging using a deformable model. In *Proceedings of the ACM SIGGRAPH Conference on Motion in Games*, pages 57–64, 2015. 1, 3
- [21] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *ICML*, pages 3569–3579, 2020. 2
- [22] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. 3D-CODED: 3D correspondences by deep deformation. In *ECCV*, pages 230–246, 2018. 2
- [23] Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. DRAPE: DRessing Any PErson. *ACM TOG*, 31(4):35–1, 2012. 1, 2, 3
- [24] 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 *CVPR*, pages 8739–8748, 2019. 1, 2, 3
- [25] Boyan Jiang, Xinlin Ren, Mingsong Dou, Xiangyang Xue, Yanwei Fu, and Yinda Zhang. LoRD: Local 4d implicit representation for high-fidelity dynamic human modeling. In *ECCV*, pages 307–326. Springer, 2022. 2
- [26] Boyi Jiang, Juyong Zhang, Yang Hong, Jinhao Luo, Ligang Liu, and Hujun Bao. BCNet: Learning body and cloth shape from a single image. In *ECCV*, pages 18–35. Springer, 2020. 2
- [27] Hyomin Kim, Hyeonseo Nam, Jungeon Kim, Jaesik Park, and Seungyong Lee. LaplacianFusion: Detailed 3D clothed-human body reconstruction. *ACM TOG*, 41(6):1–14, 2022. 2
- [28] Zorah Lahner, Daniel Cremers, and Tony Tung. Deepwrinkles: Accurate and realistic clothing modeling. In *ECCV*, pages 667–684, 2018. 2, 3
- [29] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhöfer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. TAVA: Template-free animatable volumetric actors. In *ECCV*, pages 419–436. Springer, 2022. 2
- [30] Zhe Li, Zerong Zheng, Hongwen Zhang, Chaonan Ji, and Yebin Liu. Avatacap: Animatable avatar conditioned monocular human volumetric capture. In *ECCV*, pages 322–341. Springer, 2022. 2, 5
- [31] Chen-Hsuan Lin, Chen Kong, and Simon Lucey. Learning efficient point cloud generation for dense 3D object reconstruction. In *AAAI*, pages 7114–7121, 2018. 2

- [32] Siyou Lin, Hongwen Zhang, Zerong Zheng, Ruizhi Shao, and Yebin Liu. Learning implicit templates for point-based clothed human modeling. *ECCV*, 2022. 2, 3, 6, 8
- [33] 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 TOG*, 38(4):1–12, 2019. 1, 3
- [34] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 34(6):248, 2015. 2, 3
- [35] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *SIG-GRAPH*, volume 21, pages 163–169, 1987. 6
- [36] 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 *CVPR*, pages 16082–16093, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [37] Qianli Ma, Jinlong Yang, Michael J Black, and Siyu Tang. Neural point-based shape modeling of humans in challenging clothing. *3DV*, 2022. 3, 6, 8
- [38] 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 *CVPR*, pages 6469–6478, 2020. 1, 2, 3, 5, 6, 7
- [39] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J Black. The power of points for modeling humans in clothing. In *ICCV*, pages 10974–10984, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [40] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point cloud: A simple residual mlp framework. In *ICLR*, 2022. 5
- [41] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *CVPR*, pages 4460–4470, 2019. 2
- [42] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. LEAP: Learning articulated occupancy of people. In *CVPR*, pages 10461–10471, 2021. 2, 4
- [43] Alexandros Neophytou and Adrian Hilton. A layered model of human body and garment deformation. In *3DV*, pages 171–178, 2014. 2
- [44] Hayato Onizuka, Zehra Hayirci, Diego Thomas, Akihiro Sugimoto, Hideaki Uchiyama, and Rin-ichiro Taniguchi. TetraTSDF: 3D human reconstruction from a single image with a tetrahedral outer shell. In *CVPR*, pages 6011–6020, 2020. 2
- [45] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. NPMs: Neural parametric models for 3D deformable shapes. In *ICCV*, pages 12695–12705, October 2021. 2
- [46] Pablo Palafox, Nikolaos Sarafianos, Tony Tung, and Angela Dai. SPAMs: Structured implicit parametric models. In *CVPR*, pages 12851–12860, 2022. 2
- [47] Junyi Pan, Xiaoguang Han, Weikai Chen, Jiapeng Tang, and Kui Jia. Deep mesh reconstruction from single RGB images via topology modification networks. In *ICCV*, pages 9963–9972, 2019. 2
- [48] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, pages 165–174, 2019. 2
- [49] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. TailorNet: Predicting clothing in 3D as a function of human pose, shape and garment style. In *CVPR*, pages 7363–7373, 2020. 1, 2, 3
- [50] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, pages 10975–10985, 2019. 3, 5
- [51] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael Black. ClothCap: Seamless 4d clothing capture and retargeting. *ACM TOG*, 36(4), 2017. Two first authors contributed equally. 3
- [52] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, pages 652–660, 2017. 7
- [53] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 30, 2017. 3, 5, 7
- [54] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. *NeurIPS*, 35:23192–23204, 2022. 5
- [55] Shenhan Qian, Jiale Xu, Ziwei Liu, Liqian Ma, and Shenghua Gao. UNIF: United neural implicit functions for clothed human reconstruction and animation. In *ECCV*, 2022. 2
- [56] Renderpeople, 2020. <https://renderpeople.com>. 3
- [57] 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 *ICCV*, pages 2304–2314, 2019. 2, 4
- [58] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *CVPR*, pages 84–93, 2020. 2
- [59] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, pages 2886–2897, 2021. 1, 2, 3, 4, 6, 7, 8
- [60] Igor Santesteban, Miguel A. Otaduy, and Dan Casas. Learning-Based Animation of Clothing for Virtual Try-On. *Comput. Graph. Forum*, 38(2):355–366, 2019. 2
- [61] Igor Santesteban, Miguel A Otaduy, and Dan Casas. SNUG: Self-supervised neural dynamic garments. In *CVPR*, pages 8140–8150, 2022. 8
- [62] Igor Santesteban, Nils Thuerey, Miguel A Otaduy, and Dan Casas. Self-supervised collision handling via generative 3D garment models for virtual try-on. In *CVPR*, pages 11763–11773, 2021. 2
- [63] Yu Shen, Junbang Liang, and Ming C Lin. Gan-based garment generation using sewing pattern images. In *ECCV*, volume 1, page 3, 2020. 2

- [64] 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 *ECCV*, volume 12348, pages 1–18, 2020. 2
- [65] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-GIF: Neural generalized implicit functions for animating people in clothing. In *ICCV*, pages 11708–11718, 2021. 2
- [66] Raquel Vidas, Igor Santesteban, Elena Garces, and Dan Casas. Fully convolutional graph neural networks for parametric virtual try-on. In *Comput. Graph. Forum*, volume 39, pages 145–156. Wiley Online Library, 2020. 2
- [67] Shaofei Wang, Marko Mihajlovic, Qianli Ma, Andreas Geiger, and Siyu Tang. Metaavatar: Learning animatable clothed human models from few depth images. In *NeurIPS*, 2021. 2, 3
- [68] Jane Wu, Zhenglin Geng, Hui Zhou, and Ronald Fedkiw. Skinning a parameterization of three-dimensional space for neural network cloth. *arXiv preprint arXiv:2006.04874*, 2020. 2
- [69] Donglai Xiang, Fabian Prada, Timur Bagautdinov, Weipeng Xu, Yuan Dong, He Wen, Jessica Hodgins, and Chenglei Wu. Modeling clothing as a separate layer for an animatable human avatar. *ACM TOG*, dec 2021. 2
- [70] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. ICON: implicit clothed humans obtained from normals. In *CVPR*, pages 13286–13296. IEEE, 2022. 2
- [71] Jinlong Yang, Jean-Sebastien Franco, Franck Hetroy-Wheeler, and Stefanie Wuhler. Analyzing clothing layer deformation statistics of 3D human motions. In *ECCV*, September 2018. 3
- [72] 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 TOG*, 37(5):1–14, 2018. 2
- [73] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4D: Real-time human volumetric capture from very sparse consumer rgbd sensors. In *CVPR*, June 2021. 3
- [74] Ilya Zakharkin, Kirill Mazur, Artur Grigorev, and Victor Lempitsky. Point-based modeling of human clothing. In *ICCV*, pages 14718–14727, October 2021. 2
- [75] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *CVPR*, pages 15893–15903, 2022. 2
- [76] Zerong Zheng, Tao Yu, Qionghai Dai, and Yebin Liu. Deep implicit templates for 3D shape representation. In *CVPR*, pages 1429–1439, 2021. 3, 5
- [77] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. PaMIR: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE TPAMI*, 2021. 2
- [78] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. DeepHuman: 3D human reconstruction from a single image. In *ICCV*, pages 7739–7749, 2019. 3
- [79] 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 *ECCV*, volume 12346, pages 512–530, 2020. 2