

The Power of Points for Modeling Humans in Clothing

Conference Paper**Author(s):**

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

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000503800>

Rights / license:

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

Originally published in:

<https://doi.org/10.1109/ICCV48922.2021.01079>

The Power of Points for Modeling Humans in Clothing

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

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

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




Figure 1: The **Power of Points (POP)** model for clothed humans. Based on a novel articulated dense point cloud representation, our cross-outfit model, named POP, produces pose-dependent shapes of clothed humans with coherent global shape and expressive local garment details. The trained model can be fitted to a *single* scan of an unseen subject wearing an unseen outfit, and can animate it with realistic pose-dependent clothing deformations. The results are color-coded with predicted point normals and rendered with a simple surfel-based renderer.

Abstract

Currently it requires an artist to create 3D human avatars with realistic clothing that can move naturally. Despite progress on 3D scanning and modeling of human bodies, there is still no technology that can easily turn a static scan into an animatable avatar. Automating the creation of such avatars would enable many applications in games, social networking, animation, and AR/VR to name a few. The key problem is one of representation. Standard 3D meshes are widely used in modeling the minimally-clothed body but do not readily capture the complex topology of clothing. Recent interest has shifted to implicit surface models for this task but they are computationally heavy and lack compatibility with existing 3D tools. What is needed is a 3D representation that can capture varied topology at high resolution and that can be learned from data. We argue that this representation has been with us all along — the point cloud. Point clouds have properties of both implicit and explicit representations that we exploit to model 3D garment geometry on a human body. We train a neural network with a novel local clothing geometric feature to represent the shape of different outfits. The network is trained from 3D point clouds of many types of clothing,

on many bodies, in many poses, and learns to model pose-dependent clothing deformations. The geometry feature can be optimized to fit a previously unseen scan of a person in clothing, enabling the scan to be reposed realistically. Our model demonstrates superior quantitative and qualitative results in both multi-outfit modeling and unseen outfit animation. The code is available for research purposes at <https://qianlim.github.io/POP>.

1. Introduction

Animatable clothed human avatars are required in many applications for 3D content generation. To create avatars with naturally-deforming clothing, existing solutions either involve heavy artist work or require 4D scans for training machine-learning models [55]. These solutions are expensive and often impractical. Instead, can we turn a single static 3D scan — which can be acquired at low cost today even with hand-held devices — into an animatable avatar? Currently, no existing technology is able to do this and produce realistic clothing deformations. Given a static scan, traditional automatic rigging-and-skinning methods [3, 18, 33] can be used to animate it, but are unable to produce pose-dependent clothing deformations. Physics-

based simulations can produce realistic deformations, but require “reverse-engineering” a simulation-ready clothing mesh from the given scan. This involves expert knowledge and is not fully automatic.

Taking a data-driven approach, the goal would be to learn a model that can produce reasonable pose-dependent clothing deformation across different outfit types and styles and can generalize to unseen outfits. However, despite the recent progress in modeling clothed human body shape deformations [22, 31, 36, 37, 44, 55], most existing models are *outfit-specific* and thus cannot generalize to unseen outfits. To date, no such cross-garment model exists, due to several technical challenges.

The first challenge lies in the choice of 3D shape representation. To handle outfits of different types and styles at once, the shape representation must handle changing topology, capture high-frequency details, be fast at inference time, and be easy to render. Classical triangle meshes excel at rendering efficiency but are fundamentally limited by their fixed topology. The implicit surface representation is topologically flexible, but is in general computationally heavy and lacks compatibility with existing graphics tools. Because point clouds are an explicit representation, they are efficient to render, but they can also be viewed as implicitly representing a surface. Thus they are flexible in topology and, as the resolution of the point cloud increases, they can capture geometric details. While point clouds are not commonly applied to representing clothing, they are widely used to represent rigid objects and many methods exist to process them efficiently with neural networks [1, 17, 32]. In this work, we show that the seemingly old-fashioned point cloud is, in fact, a powerful representation for modeling clothed humans.

In recent work, SCALE [36] demonstrates that a point cloud, grouped into local patches, can be exploited to represent clothed humans with various clothing styles, including those with thin structures and open surfaces. However, the patch-based formulation in SCALE often suffers from artifacts such as gaps between patches. In this work, we propose a new shape representation of dense point clouds. For simplicity, we avoid using patches, which have been widely used in recent point cloud shape representations [4, 15, 16, 21, 36], and show that patches are not necessary. Instead, we introduce smooth local point features on a 2D manifold that regularize the points and enable arbitrarily dense up-sampling during inference.

Another challenging aspect of cross-outfit modeling concerns how outfits of different types and styles can be encoded in a single, unified, model. In most existing outfit-specific models, the model parameters (typically the weights of a trained shape decoder network) need to represent both the intrinsic, pose-*independent* shape of a clothed person, and how this shape deforms as a function of the in-

put pose. To factor the problem, we propose to isolate the intrinsic shape from the shape decoder by explicitly conditioning it with a *geometric feature tensor*. The geometric feature tensor is learned in an auto-decoding fashion [43], with a constraint that a consistent intrinsic shape is shared across all examples of the same outfit. Consequently, the shape decoder can focus on modeling the pose-dependent effects and can leverage common deformation properties across outfits. At inference time, the geometric feature tensor can be optimized to fit to a scan of a clothed body with a previously unseen outfit, making it possible for the shape decoder to predict pose-dependent deformation of it based on the learned clothing deformation properties.

These ideas lead to POP: our dense point cloud model that produces pose-dependent clothing geometry across different outfits and demonstrating the *Power of Points* for modeling shapes of humans in clothing. POP is evaluated on both captured and synthetic datasets, showing state-of-the-art performance on clothing modeling and generalization to unseen outfits.

In summary, our contributions are: (1) a novel dense point cloud shape representation with fine-grained local features that produces state-of-the-art detailed clothing shapes with various clothing styles; (2) a novel geometry feature tensor that enables cross-garment modeling and generalization to unseen outfits; (3) an application of animating a static scan with reasonable pose-dependent deformations. The model and code are available for research purposes at <https://qianlim.github.io/POP>.

2. Related Work

Shape Representations for Clothed Human Modeling. Surface meshes are an efficient 3D representation that is compatible with graphics pipelines, and thus are the dominant choice for modeling clothing and clothed humans. With meshes, the clothing is represented either by deforming an unclothed body template [6, 8, 37, 40, 58, 63], or using a separately defined template [22, 23, 31, 44, 56]. While recent work successfully produces detailed geometry with graph convolutions [37], multi-layer perceptrons (MLPs) [8, 44], and UV map convolutions [28, 31], meshes suffer from two fundamental limitations: the fixed topology and the requirement for template registration. This restricts their generalization to outfit styles beyond the pre-defined templates, and makes it difficult to obtain a common representation for various clothing categories. Although recent work proposes adaptable templates [42, 66] that modify mesh connectivity, the need for registering training data to the mesh template remains challenging when a complex garment geometry is involved.

Neural implicit surfaces [11, 38, 43], on the other hand, do not require any pre-defined template, are flexible with surface topology, and have recently become a promising

choice for reconstructing [7, 24, 25, 47, 53, 54, 64, 65] and modeling [10, 12, 14, 39, 41, 55] shapes of 3D humans. Despite their ability to handle varied clothing topology, it remains an open challenge to realistically represent thin structures that are often present in daily clothing. Moreover, reconstructing an explicit surface from the implicit function costs cubic time with respect to the resolution, which restricts them from many practical applications.

Point clouds are another classic 3D representation that supports arbitrary topology as well as thin structures. Going beyond prior work that generates a sparse point set [1, 17, 32], recent approaches [4, 15, 16, 21] use deep learning to generate structured point clouds with a set of surface patches. Leveraging the patch regularization, SCALE [36] proposes an articulated dense point cloud representation to model clothing deformation. However, the patch-based point clouds often suffer from overlap [21] or separation [36] between the patches, which degrades the geometric fidelity and visual quality. Our model, POP, follows the spirit of these approaches to generate a dense, structured point cloud, but, in contrast to prior work, we deprecate the concept of patches and, instead, decode a dense point cloud from fine-grained local features. The resulting clothed human model shows high geometry fidelity, is robust to topology variations, and generalizes to various outfit styles.

Modeling Outfit Shape Space. We categorize existing models for clothing or clothed bodies into three levels of increasing generalization capacity, as summarized in Tab. 1. Note that non-parametric models for clothed human reconstruction [53, 54, 65] are out of the scope for this analysis.

Outfit-Specific. Methods from this class need to train a separate model for every outfit instance [13, 23, 31, 36, 40, 55, 56, 62] or category (e.g. all short-sleeve T-shirts) [22, 44, 58, 61]. For mesh-based methods in this category [13, 22, 31, 40, 44, 58, 62], this characteristic stems from the need to manually define a mesh template: the fixed mesh topology fundamentally prohibits generalization to a different outfit category (e.g. from pants to a skirt). These methods can, however, deal with size [22, 58] or style [44, 61] within the template-defined category. Template-free methods [36, 55] require training a separate model for each outfit, i.e. the intrinsic, pose-independent shape information is stored in the model parameters; hence, the test-time generalization to an unseen outfit is restricted.

Multi-Outfit. Combining multiple pre-defined mesh templates with a multi-head network or a garment classifier, the MGN model [6], BCNet [26], and DeepFashion3D [66] can reconstruct humans in a variety of clothing from images. In a similar spirit, CAPE [37] uses a pre-labeled one-hot vector for outfit-type conditioning and can generate new clothing from four common categories with a single model. While training a single model for multiple outfits exploits the complementary information among training data, these methods

Table 1: Data-driven models for clothing / clothed humans classified by garment space generalization.

Outfit-Specific	De Aguiar [13], DRAPE [22], GarNet [23], DeepWrinkles [31], SCALE [36], Neophytou [40], TailorNet [44], SCANimate [55], Santesteban [56], Sizer [58], Wang [61], Yang [62].
Multi-Outfit	MGN [6], BCNet [26], CAPE [37], Vidaurre [59], DeepFashion3D [66].
Arbitrary Outfit	SMPLicit [12], Shen [57], POP (Ours) .

do not show the ability to handle unseen garments beyond the pre-defined categories.

Arbitrary Outfit. To overcome the limitations brought by the fixed topology of meshes, recent work opts for other representations that can unify different clothing categories and types. Shen et al. [57] represent garments using 2D sewing pattern images that are applicable to arbitrary clothing categories. However, the final 3D garment shape is represented with a single manifold mesh that is not sufficiently expressive to represent the complexity and variety of real-world clothing. Using neural implicit surfaces, SMPLicit [12] learns a topology-aware generative model for garments across multiple categories and shows continuous interpolation between them. However, the clothing geometry tends to be bulky and lacks details. In contrast, our POP model faithfully produces geometric details of various outfits, can generalize to unseen outfits, and demonstrates state-of-the-art performance on garment space modeling.

3. Method

Our goal is to learn a single, unified, model of high-fidelity pose-dependent clothing deformation on human bodies across multiple outfits and subjects. We first introduce an expressive point-based representation that preserves geometric details and flexibly models varied topology (Sec. 3.1). Using this, we build a cross-outfit model enabled by a novel geometric feature tensor (Sec. 3.2).

As illustrated in Fig. 2, given an unclothed body, the model outputs the 3D clothed body by predicting a displacement field from the body surface based on local pose and geometric features. The trained model can be fitted to a scan of a person in previously unseen outfits and this scan can be animated with pose-dependent deformations (Sec. 3.3).

3.1. Representing Humans with Point Clouds

Point-based representations [20, 36] possess topological flexibility and fast inference speed, giving them an advantage over meshes or implicit functions for modeling articulated humans. In this work, we formulate a structured point cloud representation for modeling 3D clothed humans by learning a mapping from a 2D manifold to a 3D deforma-

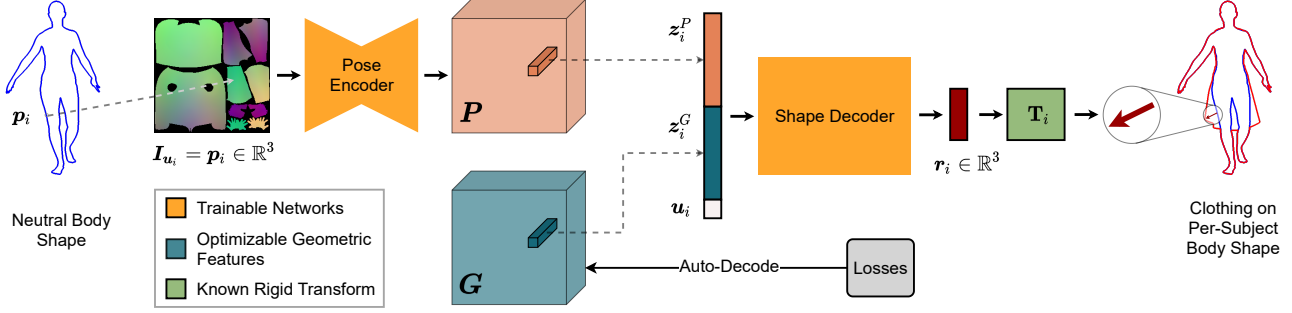


Figure 2: **Overview of POP.** Given a posed but unclothed body model (visualized as a blue 2D contour), we record the 3D positions p_i of its surface points on a UV positional map I , and encode this into a pose feature tensor P . A garment geometry feature tensor G is an optimizable variable that is pixel-aligned with P , and learned per-outfit in an auto-decoder manner. The 2D image-plane coordinate u_i describes the relative location of the points on the body surface manifold. The shape decoder queries these locations and predicts displacement vectors r_i based on the points’ local pose and geometry features.

tion field, in a similar form to AtlasNet [21]:

$$r_i = f_w(u_i; z_i) : \mathbb{R}^2 \times \mathbb{R}^Z \rightarrow \mathbb{R}^3, \quad (1)$$

where r_i is a displacement vector, $f_w(\cdot)$ is a multi-layer perceptron (MLP) with weights w , u_i is a 2D parameterization of a point i that denotes its relative location on the body surface, and z_i is the point’s local feature code containing the shape information.

We deviate from other recent point-based human representations in two key ways: 1) We use fine-grained *per-point local features* z_i as opposed to a single global feature [20, 21] or per-patch features [36] in prior work. 2) We predict the *clothing deformation field* on top of the canonically posed body, instead of absolute Euclidean point coordinates [20] or local patch deformations [36]. Both design choices lead to significant improvements in representation power, as detailed below.

Continuous Local Point Features. Recent work shows the advantage of using a local latent shape code over a global code: both for neural implicit surfaces [9, 19, 27, 46] and point clouds [36]. Decoding shape from local codes significantly improves geometry quality. In particular for modeling humans, SCALE [36] successfully leverages local features to represent pose-aware garment shape and demonstrates a significant qualitative improvement against prior work that uses a single global feature code [20].

However, the local feature in SCALE is still discretely distributed on a set of pre-defined basis points u_i on the body manifold. Each feature code is decoded into multiple points in a neighborhood (a “patch”) in the output space, but it varies discretely across patch boundaries. This is equivalent to the nearest neighbor (on the body manifold) assignment of the features, Fig. 3(a). This discrete characteristic limits the quality of the geometry. As shown in Fig. 3 of [36] and our Sec. 4.1, the patches are typically isolated from each other, leading to uneven point distributions, hence poor mesh reconstruction quality.

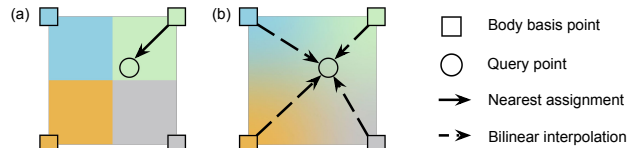


Figure 3: 2D illustration of the point feature assignment for a region on the body manifold between the body basis points. Colors represent features. (a) The nearest neighbor assignment used by SCALE [36], causes “patchy” predictions. (b) Our bilinear feature interpolation results in smoothly varying features on the manifold.

To address this problem, we make the local features more fine-grained in two ways. First, we define a denser set of basis points u_i , together with their local features, on the body manifold. In practice, this amounts to simply increasing the resolution of the body UV map (see Sec. 3.2). Second, we further diffuse the feature over the body surface: for a query coordinate on the body surface, we compute its feature by bilinearly interpolating the features from its 4 nearest basis points, Fig. 3. As a result, the network output can be evaluated at arbitrarily high resolution by querying the decoder $f_w(\cdot)$ with any point on the body surface, which we also denote as u_i from now on with a slight abuse of notation.

Local Transformations. The clothing deformation vector r_i in Eq. (1) is predicted on top of the body in the *canonical pose*. Thus a large portion of shape variation is explained by body articulation so that the network can focus on modeling the residual shape. Note that unlike the mesh vertex offset representation of clothing [6, 37, 49], our formulation does not assume a constant topology, and can thus model various outfit styles such as pants and dresses.

To reconstruct the clothed body in the posed space, we transform r_i according to the transformation matrix T_i at u_i that is given by the fitted body model. The point position

from the clothed body in the posed space is then given by:

$$\mathbf{x}_i = \mathbf{T}_i \cdot \mathbf{r}_i + \mathbf{p}_i, \quad (2)$$

where \mathbf{p}_i is the 3D Euclidean coordinate of \mathbf{u}_i on the posed unclothed body. Note that we branch out the final layers of $f_w(\cdot)$ such that it also predicts the normal $\mathbf{n}(\mathbf{x}_i)$ of each point, which is transformed with the rotation part of \mathbf{T}_i .

Since our local point features are continuous over the body surface, we also perform barycentric interpolation to obtain continuously varying \mathbf{T}_i , in the same way as LoopReg [5].

3.2. Cross-garment Modeling with a Single Model

With our structured point cloud representation, we now build a system that models pose-dependent deformations of various garments, from different categories, of different shapes and topology, dressed on different body shapes, using a *single* model. This is achieved by introducing a novel geometric feature tensor. Practically, we decompose the local feature \mathbf{z}_i in Eq. (1) into pose \mathbf{z}_i^P and garment geometry \mathbf{z}_i^G , as illustrated in Fig. 2 and detailed below.

Body Shape Agnostic Pose Encoding. We first condition our network with learned local body pose features \mathbf{z}_i^P such that the output garment deformation is pose-aware. We adopt the approach based on the UV positional map of the posed body as used in [36] as it shows better pose generalization than the traditional pose parameter conditioning [37, 44, 62]. As shown in Fig. 2, a UV positional map $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ is a 2D parameterization of the body manifold, where each valid pixel corresponds to a point on the posed body surface. The 2D image-plane coordinate of a pixel describes its manifold position: $\mathbf{u}_i = (u, v)_i$. The pixel value records the point’s location in \mathbb{R}^3 : $\mathbf{p}_i = \mathbf{I}_{\mathbf{u}_i}$. A UNet [52] encodes \mathbf{I} into a pose feature tensor $\mathbf{P} \in \mathbb{R}^{H \times W \times 64}$, where H, W are the spatial dimensions of the feature tensor, and each “pixel” from \mathbf{P} is a 64-dimensional pose feature code: $\mathbf{z}_i^P = \mathbf{P}_{\mathbf{u}_i} \in \mathbb{R}^{64}$. Because of the receptive field of the UNet, the learned pose features can contain global pose information from a larger neighborhood when necessary.

The UV positional map naturally contains information about body shape. To generalize to different subjects, we use posed bodies of a *neutral shape* for the pose encoding. Still, the predicted clothing deformations are added to each subject’s body respectively.

Geometric Feature Tensor. In most learning-based, outfit/subject-specific clothed body models [36, 37, 55, 62], the clothing shape information is contained in the parameters of the trained shape decoder network, limiting generalization to unseen outfits. What is needed for cross-outfit modeling is a mechanism that decouples the intrinsic, pose-independent shape of a clothed person from the decoder,

so that it can focus on modeling *how* the shape deforms with the pose. To that end, we propose to explicitly condition the shape decoder with a *geometric feature tensor* $\mathbf{G} \in \mathbb{R}^{H \times W \times 64}$, Fig. 2.

The geometric feature tensor follows the spirit of being local and is pixel-aligned with the pose feature tensor. Each of its “pixels” represents a local shape feature on a body point: $\mathbf{z}_i^G = \mathbf{G}_{\mathbf{u}_i} \in \mathbb{R}^{64}$. Unlike the pose features, the geometry features are learned in an auto-decoding [43] fashion; i.e. they are updated during training such that the optimal representation for the garment geometry is discovered by the network itself. Importantly, we use a consistent \mathbf{G} for each outfit across all of its training examples (in different poses). In this way, \mathbf{G} is enforced to be a pose-agnostic canonical representation of each outfit’s geometry.

Our geometry feature tensor plays a similar role as the pre-defined clothing templates in many mesh-based garment models [23, 44, 56, 59]. However, by auto-decoding the geometric features, we get rid of the reliance on manual template definition. More importantly, our decoder network and the neural geometry feature are fully differentiable. This enables generalization to *unseen* garments at test time, which is done by optimizing \mathbf{G} to fit the target clothed body scan. See Sec. 3.3 for details.

Shared Garment Shape Decoder. With the introduced local pose and geometry features, we can re-write Eq. (1) in a more concrete form: $\mathbf{r}_i = f_w([\mathbf{u}_i, \mathbf{z}_i^P, \mathbf{z}_i^G])$, where $[\cdot, \cdot, \cdot]$ denotes concatenation. While \mathbf{z}_i^G is optimized for each garment and \mathbf{z}_i^P is acquired from each pose, $f_w(\cdot)$ is shared for the entire set of all garments and poses. By training on many outfits and poses, the decoder learns common properties of clothing deformation, with which it can animate scans in unseen outfits at test-time, as described below.

3.3. Training and Inference

Loss Functions. We train POP with multi-subject and outfit data. During training, the parameters of the UNet pose encoder, the garment shape decoder, and the geometric feature tensor \mathbf{G} are optimized, by minimizing the loss function:

$$\mathcal{L}_{\text{total}} = \lambda_d \mathcal{L}_d + \lambda_n \mathcal{L}_n + \lambda_{\text{rd}} \mathcal{L}_{\text{rd}} + \lambda_{\text{rg}} \mathcal{L}_{\text{rg}}, \quad (3)$$

where the λ ’s are weights that balance the loss terms, and the \mathcal{L} ’s are the following loss terms.

First, the normalized Chamfer Distance \mathcal{L}_d is employed to penalize the average bi-directional point-to-point L_2 distances between the generated point cloud \mathbf{X} and a sampled point set \mathbf{Y} from the ground truth surface: $\mathcal{L}_d = d(\mathbf{x}, \mathbf{y}) =$

$$\frac{1}{M} \sum_{i=1}^M \min_j \|\mathbf{x}_i - \mathbf{y}_j\|_2^2 + \frac{1}{N} \sum_{j=1}^N \min_i \|\mathbf{x}_i - \mathbf{y}_j\|_2^2, \quad (4)$$

where M, N are the number of points from the generated point cloud and the ground truth surface, respectively.

The normal loss \mathcal{L}_n is the averaged $L1$ discrepancy between the normal prediction on each generated point and its nearest neighbor from the ground truth point set:

$$\mathcal{L}_n = \frac{1}{M} \sum_{i=1}^M \left\| \mathbf{n}(\mathbf{x}_i) - \mathbf{n}(\operatorname{argmin}_{\mathbf{y}_j \in \mathbf{Y}} d(\mathbf{x}_i, \mathbf{y}_j)) \right\|_1, \quad (5)$$

where $\mathbf{n}(\cdot)$ denotes the unit normal of the given point.

An $L2$ regularizer \mathcal{L}_{rd} discourages the predicted point displacements from being extremely large. Similarly, the term \mathcal{L}_{rg} penalizes the $L2$ -norm of the vectorized geometric feature tensor to regularize the garment shape space:

$$\mathcal{L}_{rd} = \frac{1}{M} \sum_{i=1}^M \|\mathbf{r}_i\|_2^2, \quad \mathcal{L}_{rg} = \frac{1}{C} \sum_{m=1}^C \|\mathbf{G}_m\|_2^2, \quad (6)$$

where C is the number of garments seen in training.

The detailed model architecture, hyper-parameters and training procedure are provided in the SupMat.

Inference: Scan Animation. At test-time, POP can generalize to unseen poses of both the previously seen and unseen outfits. For a seen outfit, we use its geometric feature tensor optimized from training, and infer the clothing deformation on unseen poses with a simple forward pass of the network.

To test on a scan $\hat{\mathbf{Y}}$ of a human wearing unseen clothing, we first fix the weights of the UNet pose encoder and the shape decoder $g_w(\cdot)$, and optimize the geometric feature tensor such that the total loss against the scan is minimized:

$$\hat{\mathbf{G}} = \operatorname{argmin} \mathcal{L}_{\text{total}}(\hat{\mathbf{Y}}). \quad (7)$$

Afterwards, the estimated $\hat{\mathbf{G}}$ is fixed, and is then treated as in the case of a seen garment.

As with other point-based human models, the point cloud generated by POP can either be meshed using classical tools such as the Poisson Surface Reconstruction (PSR) [29], or directly rendered into realistic images using recent point-based neural rendering techniques [2, 30, 50]. However, in this work, we do not rely on neural rendering to inpaint the gaps between the points. Instead, we show qualitative results using a simple surfel-based renderer to more directly highlight the geometric properties of the smooth, high-resolution, human point cloud generated by POP.

4. Experiments

Due to the lack of comparable existing work on POP’s two key features, namely cross-outfit modeling and single scan animation, we first evaluate its representation power on a simpler but related task: outfit-specific shape modeling, and compare with two state-of-the-art methods (Sec. 4.1). We then discuss the efficacy of our cross-outfit learning formulation (Sec. 4.2) and demonstrate single scan animation (Sec. 4.3).

Datasets. We train and evaluate our method and baselines on both a captured clothed human dataset, CAPE [37], and our new synthetic dataset called *ReSynth*. From the CAPE dataset, we use the three subjects (00096, 00215, 03375) that contain the most abundant outfit variations (14 outfits in total) to compare the representation capacity of different methods. The synthetic ReSynth dataset is created with a larger variation in outfit shapes, styles, and poses. We worked with a professional clothing designer to create 3D outfit designs that faithfully reflect those in a set of commercial 3D clothed human scans (Renderpeople [51]), resulting in 24 outfits including challenging cases such as skirts and jackets. We then use physics simulation to drape the clothing on the 3D bodies from the AGORA dataset [45], which we animate to generate many poses. Details of the datasets are provided in the SupMat., and we will release ReSynth for research purposes.

Baselines. To evaluate the representation power of our model, we first compare with two recent methods for pose-dependent human shape modeling (Sec. 4.1): NASA [14] and SCALE [36]. To evaluate the effectiveness of our cross-outfit modeling formulation (Sec. 4.2), we compare two versions of our model: per-outfit trained and a cross-outfit model trained with data from all outfits, both using the same architecture. For animating unseen scans (Sec. 4.3), we qualitatively compare with classical Linear Blend Skinning using the SMPL [34] body model.

Metrics. We quantitatively evaluate each method using the Chamfer Distance (in m^2 , Eq. (4)) and the $L1$ normal discrepancy (Eq. (5)), computed over the 50K points generated by our method and SCALE. For the implicit surface baseline NASA, the points for evaluation are sampled from surface extracted using Marching Cubes [35]. To evaluate the visual quality of the generated results, we perform a large-scale user study on the Amazon Mechanical Turk (AMT) and report the percentage of users that favor the results from our method over the baseline. Details of the user study are provided in the SupMat.

4.1. Representation Power

Tab. 2 summarizes the numerical results of reconstructing pose-dependent garment shape from different methods, tested with seen outfits on *unseen* motion sequences. As the difficulty of shape modeling varies greatly across different outfit types (e.g. how a loose jacket deforms is much more complex than that of a tight T-shirt), we report three types of statistics to holistically reflect the performance of each model: the mean error averaged across all test examples from all outfits, the median of the per-outfit calculated mean error (denoted as “Outfit Median” in Tab. 2), and the per-outfit maximum error (“Outfit Max”).

Comparison with SoTA. The upper section of Tab. 2 shows a comparison with NASA [14] and SCALE [36]. NASA

Table 2: Results of pose-dependent deformation prediction on unseen test sequences from the captured CAPE dataset and our ReSynth data. Best results are in **boldface**.

Methods	CAPE Data						ReSynth Data					
	Chamfer- L_2 ($\times 10^{-4} m^2$) \downarrow			Normal diff. ($\times 10^{-1}$) \downarrow			Chamfer- L_2 ($\times 10^{-4} m^2$) \downarrow			Normal diff. ($\times 10^{-1}$) \downarrow		
	Mean	Outfit Median	Outfit Max	Mean	Outfit Median	Outfit Max	Mean	Outfit Median	Outfit Max	Mean	Outfit Median	Outfit Max
NASA [14]	6.087	1.190	32.35	1.275	1.277	1.497	—	—	—	—	—	—
SCALE [36]	0.721	0.689	0.971	1.168	1.170	1.335	1.491	0.680	8.451	1.041	1.054	1.321
Ours, per-outfit	0.639	0.607	0.831	1.146	1.150	1.293	1.356	0.651	7.339	1.013	1.006	1.289
Ours, multi	0.592	0.550	0.757	1.115	1.116	1.256	1.366	0.635	7.386	1.022	1.037	1.280
1/2 Data	0.598	0.560	0.765	1.122	1.127	1.257	1.405	0.665	7.414	1.032	1.042	1.299
1/4 Data	0.621	0.586	0.841	1.134	1.142	1.271	1.406	0.674	7.469	1.032	1.043	1.296
1/8 Data	0.662	0.623	0.992	1.165	1.176	1.310	1.490	0.720	7.859	1.050	1.056	1.326

represents the body shape with an ensemble of articulated occupancy functions defined per body part. As it requires pre-computing occupancy values for training, it is in general not applicable to our synthetic data, which typically does not contain water-tight meshes. SCALE produces a point cloud of clothed bodies, where the points are grouped into local patches. Notably, both methods need to train a separate model per outfit. Under the same outfit-specific training setting, POP outperforms both baselines on both datasets under all metrics. We further conduct an AMT user study to evaluate the perceptual quality, comparing POP side-by-side against the strongest baseline, SCALE. On the CAPE data, 89.8% participants rate POP’s results as having “higher visual quality” than SCALE (10.2%); while on ReSynth, 84.8% users favor POP over SCALE (15.2%).

The differences in perceptual quality can be seen in Fig. 4. All approaches provide clear pose-dependent effects. However, NASA suffers from non-smooth transitions between separately-modeled body parts and is unable to handle thin structures and open surfaces such as the skirt in ReSynth. SCALE, on the other hand, produces a smooth global shape with local details, but the isolation between the patches leads to sub-optimal visual quality. In contrast, the dense point clouds generated by POP manifest a coherent overall shape and expressive local details. This illustrates the advantage of our continuous local point features as opposed to the discretely defined patch features.

4.2. Cross-Outfit Modeling

A key feature of POP is the ability to handle multiple outfits of varied types using a single model, without pre-defining garment templates. As shown in Tab. 2, our cross-outfit model trained with all subjects and outfits (“Ours, multi”) reaches overall comparable performance to the outfit-specific models (“Ours, per-outfit”) on the synthetic ReSynth data, and has significantly lower error on the CAPE dataset. Still, on both datasets, the cross-outfit model consistently outperforms the two baselines analyzed above.

The result on the CAPE data reveals the advantage of cross-outfit modeling: the information between different

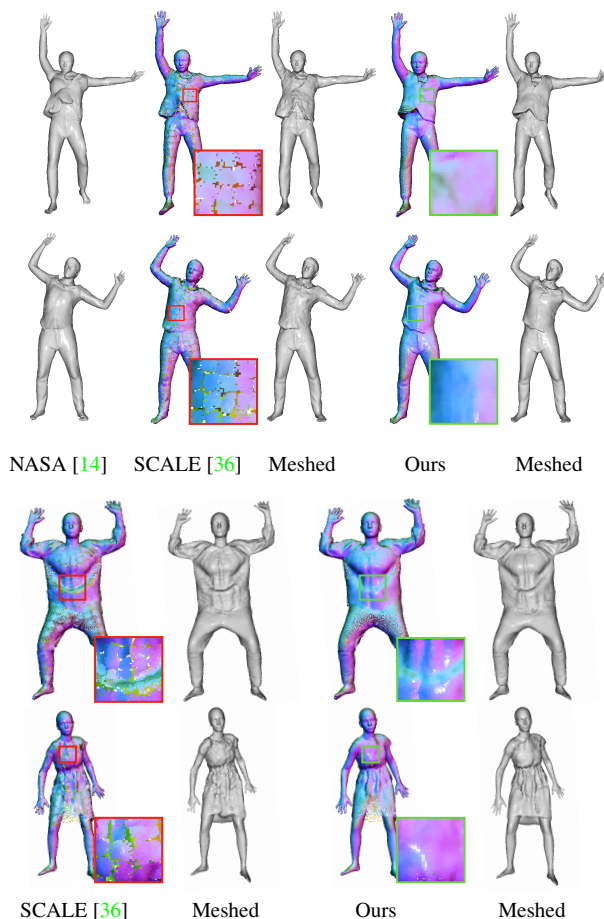


Figure 4: Comparison with SoTAs on the CAPE (upper 2 rows) and ReSynth (lower 2 rows) data. The dense point clouds from POP (cross-outfit model) are clean and preserve local details, while baseline methods suffer from artifacts. Note the point clouds are colored according to predicted normal and rendered with surfels [48]. Best viewed zoomed-in on a color screen.

but similar outfit styles can be shared. The CAPE dataset mostly consists of simple and similar garments such as long/short T-shirts and pants, but the motions are performed differently by different subjects. Intuitively, training a

cross-outfit model leads to mutual data augmentation in the pose space among different outfits, hence a better generalization to unseen poses as seen in the numerical results. In contrast, the synthetic data are simulated with a consistent set of motion sequences for all outfit types with largely varied geometry. As a result, the inter-outfit knowledge sharing is limited, leading to similar performance between per-outfit and cross-outfit models.

Robustness Against Limited Data. In the lower section of Tab. 2, we evaluate the performance of our cross-outfit model trained with subsets sampled from the full training set. Even with only 1/4 of data from each outfit, our cross-outfit model is comparable to the outfit-specific models of SCALE that are trained with full data.

4.3. Single Scan Animation

Once trained, our cross-outfit POP model can be fit to a single scan of an unseen subject with an unseen outfit, and then animated with pose-dependent clothing variation, as described in Sec. 3.3. Notably, this is a very challenging task since it requires generalization in both the outfit shape space and pose space. Figure 5 qualitatively compares our model (trained on CAPE data) and the classical Linear Blend Skinning (LBS) technique that uses the SMPL model [34] to animate the given unseen scan with an unseen motion. Here we use sampled points from the mesh provided in the CAPE dataset as the target scan. As LBS uses simple rigid transformations of the body parts only, it cannot produce complex pose-dependent shape variation. In contrast, POP produces reasonable and vivid clothing deformation such as the lifted hems. In Fig. 6 we deploy POP (trained on the ReSynth data) to animate an unseen scan from ReSynth and one from real-world captures [51], respectively. Note that the latter test is much more challenging due to the domain gap between our synthetic training data and the captured test data. POP produces high-quality clothing geometry on the ReSynth test example and generates reasonable animation for the challenging real-world scan, which opens the promising new direction of automatic 3D avatar creation from a “one-shot” observation.

5. Conclusion

We have introduced POP, a dense, structured point cloud shape representation for articulated humans in clothing. Being template-free, geometrically expressive, and topologically flexible, POP models clothed humans in various outfit styles with a single model, producing high-quality details and realistic pose-dependent deformation. The learned POP model can also be used to animate a single human scan from an unseen subject and clothing. Our evaluations on both captured data and our synthetic data demonstrate the efficacy of the continuous local features and the advantage of a cross-outfit model over traditional subject-specific ones.

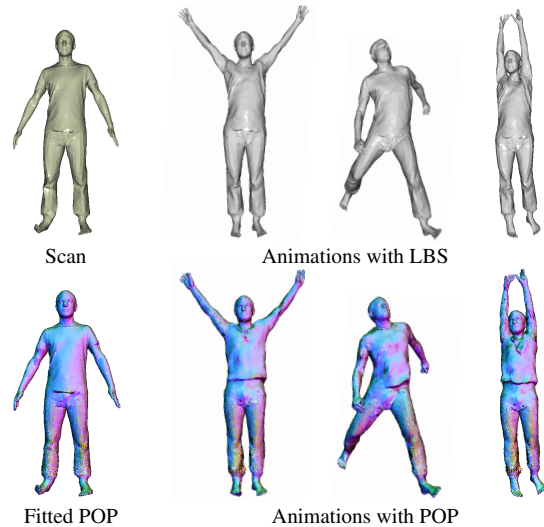


Figure 5: Comparison of animation with POP and an LBS-based method. The unseen scan from the CAPE dataset is animated on unseen motions.




Figure 6: Animation of an unseen example from ReSynth (upper row) and a real-world capture (lower row) with unseen motions.

POP assumes that the minimally closed body under the clothing is given and the training scans have no noise or missing data. Dealing with partial, noisy, scans and combining it with automatic body shape and pose estimation [5, 60] are promising future directions for fully-automatic scan animation. Our use of the UV map for body pose conditioning can sometimes lead to “seams” in the outputs (see SupMat.). Future work should explore more continuous parameterizations of the human body manifold. Additionally, here we factor out dependencies of clothing deformation on body shape for simplicity. Given sufficient training data, this would be easy to learn by replacing the neutral shape with the true shape in the UV positional map.

Acknowledgements and disclosure: Q. Ma is partially funded by DFG (German Research Foundation)-276693517 SFB 1233 and supported by the Max Planck ETH Center for Learning Systems. The conflict of interest of MJB can be found at <https://files.is.tue.mpg.de/black/CoI/ICCV2021.txt>.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3D point clouds. In *Proceedings of Machine Learning and Systems (ICML)*, pages 40–49. PMLR, 2018. [2, 3](#)
- [2] 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. [6](#)
- [3] Ilya Baran and Jovan Popović. Automatic rigging and animation of 3D characters. *ACM Transactions on Graphics (TOG)*, 26(3):72–es, 2007. [1](#)
- [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. [2, 3](#)
- [5] 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. [5, 8](#)
- [6] 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, 3, 4](#)
- [7] Aljaz Bozic, Pablo Palafox, Michael Zollhöfer, Justus Thies, Angela Dai, and Matthias Nießner. Neural deformation graphs for globally-consistent non-rigid reconstruction. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. [3](#)
- [8] Andrei Burov, Matthias Nießner, and Justus Thies. Dynamic surface function networks for clothed human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [9] 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. [4](#)
- [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 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [3](#)
- [11] 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. [2](#)
- [12] 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)*, pages 11875–11885, 2021. [3](#)
- [13] Edilson De Aguiar, Leonid Sigal, Adrien Treuille, and Jessica K Hodgins. Stable spaces for real-time clothing. In *ACM Transactions on Graphics (TOG)*, volume 29, page 106. ACM, 2010. [3](#)
- [14] 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](#)
- [15] 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. [2, 3](#)
- [16] 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](#)
- [17] 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, 3](#)
- [18] Andrew Feng, Dan Casas, and Ari Shapiro. Avatar reshaping and automatic rigging using a deformable model. In *Proceedings of the 8th ACM SIGGRAPH Conference on Motion in Games*, pages 57–64, 2015. [1](#)
- [19] Kyle Genova, Forrester Cole, Avneesh 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. [4](#)
- [20] 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. [3, 4](#)
- [21] 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](#)
- [22] 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](#)
- [23] 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, 3, 5](#)
- [24] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. ARCH++: Animation-ready clothed human reconstruction revisited. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [3](#)
- [25] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. ARCH: Animatable reconstruction of clothed humans. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3093–3102, 2020. [3](#)
- [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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 18–35. Springer, 2020. 3
- [27] 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. 4
- [28] 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. 2
- [29] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics Symposium on Geometry Processing*, volume 7, 2006. 6
- [30] Maria Kolos, Artem Sevastopolsky, and Victor Lempitsky. TRANSPR: Transparency ray-accumulating neural 3D scene point renderer. In *International Conference on 3D Vision (3DV)*, pages 1167–1175, 2020. 6
- [31] 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
- [32] 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, 3
- [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 Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 1
- [34] 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. 6, 8
- [35] 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
- [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 *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. 2, 3, 4, 5, 6, 7
- [37] 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, 4, 5, 6
- [38] 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
- [39] 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
- [40] 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
- [41] Pablo Palafox, Aljaz Bozic, Justus Thies, Matthias Nießner, and Angela Dai. Neural parametric models for 3D deformable shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 3
- [42] 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. 2
- [43] 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, 5
- [44] Chaitanya Patel, Zhouyuecheng 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
- [45] Priyanka Patel, Chun-Hao Huang Paul, Joachim Tesch, David Hoffmann, Shashank Tripathi, and Michael J. Black. AGORA: Avatars in geography optimized for regression analysis. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 13468–13478, June 2021. 6
- [46] Songyu 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. 4
- [47] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 9054–9063, June 2021. 3
- [48] Hanspeter Pfister, Matthias Zwicker, Jeroen Van Baar, and Markus Gross. Surfels: Surface elements as rendering primitives. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 335–342, 2000. 7
- [49] 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. 4
- [50] Sergey Prokudin, Michael J. Black, and Javier Romero. SM-PLpix: Neural avatars from 3D human models. In *Winter Conference on Applications of Computer Vision (WACV)*, 2021. 6
- [51] Renderpeople, 2020. <https://renderpeople.com>. 6, 8
- [52] 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
- [53] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization.

- tion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2304–2314, 2019. [3](#)
- [54] 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. [3](#)
- [55] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. [1](#), [2](#), [3](#), [5](#)
- [56] 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](#), [3](#), [5](#)
- [57] Yu Shen, Junbang Liang, and Ming C Lin. Gan-based garment generation using sewing pattern images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 1, page 3, 2020. [3](#)
- [58] 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](#), [3](#)
- [59] Raquel Vidas, Igor Santesteban, Elena Garces, and Dan Casas. Fully convolutional graph neural networks for parametric virtual try-on. In *Computer Graphics Forum*, volume 39, pages 145–156. Wiley Online Library, 2020. [3](#), [5](#)
- [60] Shaofei Wang, Andreas Geiger, and Siyu Tang. Locally aware piecewise transformation fields for 3D human mesh registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7639–7648, June 2021. [8](#)
- [61] Tuanfeng Y Wang, Duygu Ceylan, Jovan Popović, and Niloy J Mitra. Learning a shared shape space for multimodal garment design. In *ACM Transactions on Graphics (TOG)*, pages 203–216, 2018. [3](#)
- [62] Jinlong Yang, Jean-Sébastien Franco, Franck Hétroy-Wheeler, and Stefanie Wuhler. Analyzing clothing layer deformation statistics of 3D human motions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 237–253, 2018. [3](#), [5](#)
- [63] 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](#)
- [64] Ze Yang, Shenlong Wang, Siva Manivasagam, Zeng Huang, Wei-Chiu Ma, Xinchun Yan, Ersin Yumer, and Raquel Urtasun. S3: Neural shape, skeleton, and skinning fields for 3D human modeling. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 13284–13293, 2021. [3](#)
- [65] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. PaMIR: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [3](#)
- [66] 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. [2](#), [3](#)