

MonoClothCap: Towards Temporally Coherent Clothing Capture from Monocular RGB Video

Donglai Xiang^{1*} Fabian Prada² Chenglei Wu² Jessica Hodgins^{1,3}
¹Carnegie Mellon University ²Facebook Reality Labs Research ³Facebook AI Research
{donglaix, jkh}@cs.cmu.edu {fabianprada, chenglei}@fb.com




Figure 1: Given a monocular RGB video as input, our proposed method captures temporally coherent dynamic clothing deformation that cannot be explained by a bare human body model. From left to right, we show the input frame, body capture results using the SMPL model [25], and our clothing capture results.

Abstract

We present a method to capture temporally coherent dynamic clothing deformation from a monocular RGB video input. In contrast to the existing literature, our method does not require a pre-scanned personalized mesh template, and thus can be applied to in-the-wild videos. To constrain the output to a valid deformation space, we build statistical deformation models for three types of clothing: T-shirt, short pants and long pants. A differentiable renderer is utilized to align our captured shapes to the input frames by minimizing the difference in both silhouette, segmentation, and texture. We develop a UV texture growing method which expands the visible texture region of the clothing sequentially in order to minimize drift in deformation tracking. We also extract fine-grained wrinkle detail from the input videos by fitting the clothed surface to the normal maps estimated by a convolutional neural network. Our method produces temporally coherent reconstruction of body and clothing from monocular video. We demonstrate successful clothing capture results from a variety of challenging videos. Extensive quantitative experiments demonstrate the effectiveness of our method on metrics including body pose error and surface reconstruction error of the clothing.

1. Introduction

Dynamic capture of detailed human geometry and motion from monocular images and videos is attracting increasing attention in the computer vision and computer graphics community. High-quality human capture would enable applications in virtual and augmented reality, games, and movies. In recent years, great progress has been made on the estimation of general body shape from a single image or a monocular video [22, 50, 31, 33]. However, capturing the detailed deformation of clothing as it moves on the human body is still far from a solved problem.

Capturing a temporally coherent shape for clothing from monocular RGB imagery is an extremely challenging task, due to the fundamental ambiguity of single-view 3D reconstruction and the large deformation space of clothing. Previous work [49, 16, 17] utilizes a 3D personalized actor model as a shape prior to track the dynamic clothing deformation. This model is acquired by multi-view reconstruction on an additional video of the same actor wearing the same clothing and rotating in a T-pose. However, such a model is generally unavailable for in-the-wild videos. The need for a pre-scanned template model limits the applicability of these approaches.

With the development of deep neural networks, other efforts have been made to regress a clothed human shape directly from a single input image with supervised learning [44, 30, 39, 43, 55, 4, 40]. These methods produce

*Work partly done during internship at Facebook AI Research.

plausible results for individual input images of common human poses. However, it is difficult to extend them to capture temporally coherent dynamic clothing deformation from monocular videos for the following reasons. First, these methods are not robust to the variety of human motion due to the limited diversity of training data. They can easily produce incomplete geometry that is difficult to fix via post-processing. Second, it is non-trivial to estimate the temporal correspondence from the output of individual frames due to the data representation used (voxel [44, 55], depth map [43] or implicit function [44, 30, 39, 43, 55, 4, 40]). This limits the application of these methods in scenarios that require correspondence, such as clothing retargeting or image editing.

In this work, we present a novel method to capture dynamic clothing deformation from a monocular RGB video in a *temporally coherent* manner, as illustrated by Fig. 1. To the best of our knowledge, it is the first attempt to solve this challenging problem without the prerequisite of a pre-scanned personalized template [49, 16, 17].

Our method is based on the following observations. First, a deformation model of the clothing that provides a statistical shape prior is key to solving the problem. It not only reduces the ambiguity of single-view 3D reconstruction, but also helps to estimate temporal correspondence across frames. While clothing models have been investigated in the existing literature [51, 27] for the purpose of clothing shape generation, our work is the first study that fully demonstrates the value of a clothing model for RGB-based clothing capture¹. Second, to solve the clothing capture problem, we make use of human appearance information including silhouette, segmentation, texture and surface normal. We present a novel method to integrate all those image measurements using a differentiable renderer [24, 12]. Our method captures the realistic dynamic of clothing in a temporally coherent manner including fine-grained wrinkle details from various videos.

Our Contributions. (1) We present the first approach for temporally coherent clothing capture from a monocular RGB video without using a pre-scanned template of the subject. (2) We propose a novel method to capture clothing deformation by fitting statistical clothing models to image measurements including silhouette, segmentation, texture and surface normal with a differentiable renderer.

2. Related Work

Single-Image Human Pose and Shape Estimation. Most previous work in human pose estimation focuses on the position of body keypoints in 2D [46, 11, 10] and 3D [56, 34, 42]. Because estimating 3D pose from single

images is highly ill-posed, deformable human models including SMPL [25], SMPL-X [32] and Adam [21] are used to help with the problem by fitting the models to images [8, 32, 48]. These models not only provide a strong prior for body pose, but also enable estimation of 3D body shape from single images. Deformable human models can be further integrated in deep neural network architectures [22, 31, 33, 50]. These networks are usually trained in a weakly-supervised manner without full 3D supervision.

Because deformable human models are not able to express clothing shape, all the work above only estimates body shapes with minimal clothing. Detailed clothing shape has been largely ignored in the previous literature, except a few papers [44, 30, 39, 55, 43, 13]. These methods use deep neural networks to infer dense clothed human shapes in various data representation including voxels [44, 55], depth maps [43], point clouds [13] and implicit functions [30, 39, 40, 19], all with supervised learning. However, because the amount of available training data is very limited, these methods are not robust to human motion. In addition, it is non-trivial to estimate correspondence across frames required for clothing capture due to their data representation. Our method achieves temporally coherent body and clothing capture in terms of both geometry and correspondence with the help of a statistical clothing model.

Garment Modeling and Reconstruction. Human clothing, especially physically based simulation of garments [6, 9, 45, 41, 23], has been extensively studied due to its important role in animation. Recently, there is growing interest in modeling garments in a data-driven manner. Pons-Moll *et al.* [35] proposes a method to automatically segment 4D clothed human scans into different garment pieces, and track the deformation of each piece over time. The captured clothing data can be further used to train a deformable model, either a linear model [51] or a deep neural network [27]. In those methods, the clothing models are primarily used for shape generation, while we use the model to track clothing deformation from a monocular video.

Another line of work reconstructs clothing shape from images by allowing per-vertex deformation on top of the SMPL body model. Alldieck *et al.* [3, 2] builds clothed human avatar from videos of a person slowly rotating in A-pose. This is further improved to use only images of several different views [1, 7] or even a single image [4]. However, these methods reconstruct clothing as static objects without considering the temporal dynamics. By contrast, in this paper, we address the challenging problem of capturing clothing dynamics from a monocular video.

Monocular Human Performance Capture. Motion capture and performance capture refer to the capture of space-time coherent human motion sequences in the form of sparse 3D joints and surface geometry respectively. Many approaches have been developed to enable motion and per-

¹Due to the limitation in types of available clothing data to train our model, in this paper, we assume that the subject to be captured wears a T-shirt on the upper body and shorts or pants on the lower body.




Figure 2: An overview of our clothing capture pipeline.

formance capture from multi-view inputs [36, 38, 37, 18]. Here we focus on monocular-based capture methods. Mehta *et al.* [29, 28] proposes systems to capture body skeleton motion from a single RGB video in real time. Some performance capture methods [52, 53] are proposed to capture dense human body and clothing geometry from a monocular RGB-D video using a double-layer representation. Most relevant to our work are performance capture methods from monocular RGB videos [49, 16]. These methods, however, require a pre-scanned mesh template of the subject, which restricts the applications where they can be used. Habermann *et al.* [17] further proposes to train a deep neural network to deform a pre-scanned mesh template to match the surface deformation in the video. This method, requires the mesh template and multi-view images of the subject for network training. Our method relaxes the constraint to scenarios such as in-the-wild videos where neither pre-scanned templates nor multi-view images are available.

3. Method Overview

In this section, we present an overview of our approach. Our goal is to capture the dynamic deformation of three types of garments, T-shirt, shorts and pants, along with the underlying body shape from a monocular video. Our method takes as input a sequence of images, denoted as $\{\mathbf{I}_i\}_{i=1}^F$, where F is the number of frames in the sequence. The subject is assumed to be wearing a T-shirt for the upper body. The clothing for the lower body is manually identified as either short pants or long pants. Our method outputs a sequence of mesh pairs $\{\mathbf{M}_i^b, \mathbf{M}_i^c\}_{i=1}^F$, where \mathbf{M}_i^b denotes the body mesh and \mathbf{M}_i^c denotes the clothed mesh. $\{\mathbf{M}_i^b\}_{i=1}^F$ and $\{\mathbf{M}_i^c\}_{i=1}^F$ are both temporally coherent with fixed topology across time. \mathbf{M}_i^b and \mathbf{M}_i^c share the same vertex positions except for the clothing region.

Our method makes use of linear clothing deformation models defined in the canonical pose. We briefly describe our model formulation and model building procedure in Section 4. Our pipeline to capture clothing from a monocular video consists of four stages, explained in Section 5. First, we estimate the underlying body pose and shape of

subject (Section 5.1). Then, we run sequential tracking of the clothing using our linear clothing models. This step is followed by a batch optimization stage including all the frames to produce temporally coherent dynamic clothing deformation (Section 5.2). In the final stage, we add fine-grained wrinkle detail to our results (Section 5.3). A visualization of this pipeline is shown in Fig. 2.

4. Statistical Clothing Deformation Model

Statistical models of clothing have been investigated for clothing shape generation in the previous literature [51, 27], but have yet to be exploited for capturing clothing from a monocular video. In this section we give the mathematical formulation of our clothing deformation models and briefly describe the procedure to learn these models from data.

4.1. Model Formulation

Our clothing models are built on top of the SMPL body model [25]. SMPL is controlled by a set of model parameters (β, θ) , where $\beta \in \mathbb{R}^{10}$ is the shape coefficients and $\theta \in \mathbb{R}^{72}$ is the joint angles that control body pose. We denote the set of $n_v = 6890$ output vertices by $M(\beta, \theta)$. Then formally,

$$M(\beta, \theta) = W(T(\beta, \theta), J(\beta), \theta, \mathcal{W}), \quad (1)$$

where W is the Linear Blend Skinning (LBS) function; $T(\beta, \theta)$ is the rest pose body shape; $J(\beta)$ is the locations of 24 kinematic joints; \mathcal{W} is the blend weights. In particular, the unposed shape $T(\beta, \theta)$ is defined as the sum of template shape \bar{T} , shape dependent deformation $B^S(\beta)$ and pose dependent deformation $B^P(\theta)$,

$$T(\beta, \theta) = \bar{T} + B^S(\beta) + B^P(\theta) \quad (2)$$

On top of the SMPL model, an extra additive offset field D is introduced to account for clothing deformation in rest pose, i.e.,

$$T^c = T(\beta, \theta) + D. \quad (3)$$

D includes a number of n_v per-vertex offsets, each denoted by $D_j \in \mathbb{R}^3$, where $1 \leq j \leq n_v$. Here we decompose D into offsets from the upper clothing D^u and offsets from the lower clothing D^l . D^u and D^l share the same dimensionality as D . They take non-zero values if the respective garments cover body vertex j ; for an exposed skin vertex we have $D_j^u = D_j^l = 0$. Notice that some body vertices might be covered by both upper and lower clothing, for example, around the waist. To account for this phenomenon, we merge D^u and D^l into a single offset field D by

$$D_j = \begin{cases} D_j^u & \text{if } \|D_j^u\| \geq \|D_j^l\|, \\ D_j^l & \text{otherwise.} \end{cases} \quad (4)$$




Figure 3: A visualization of our clothing model formulation. From left to right, we show (1) the mean SMPL template, (2) personalized body shape with pose dependent deformation, (3) upper clothing offsets, (4) lower clothing offsets, (5) combined clothing offsets, and (6) posed clothing output.

The dimensions of D^u and D^l are very high ($3n_v$), so we use PCA dimension reduction to enable control with low-dimensional parameters $\mathbf{z}^u, \mathbf{z}^l \in \mathbb{R}^{n_z}$. Formally,

$$D^k(\mathbf{z}^k) = \mathbf{A}^k \mathbf{z}^k + \bar{\mathbf{d}}^k, k \in \{u, l\} \quad (5)$$

where $\mathbf{A}^k \in \mathbb{R}^{3n_v \times n_z}$ is the matrix of PCA bases and $\bar{\mathbf{d}}^k$ is the mean value vector. We use the skinning function W of SMPL to transform the clothed shape from rest pose to target pose. Finally, our clothing model is formulated as

$$M^c(\beta, \theta, \mathbf{z}) = W(T^c(\beta, \theta, \mathbf{z}), J(\beta), \theta, \mathcal{W}), \quad (6)$$

$$T^c(\beta, \theta, \mathbf{z}) = T(\beta, \theta) + D(\mathbf{z}), \quad (7)$$

where $\mathbf{z} = \{\mathbf{z}^u, \mathbf{z}^l\}$ is the collection of clothing parameters. A visual illustration of our clothing model formulation is shown in Figure 3.

4.2. Model Building

We build our models from the BUFF dataset [54], a collection of high-resolution 4D people scan. We build a model for each of the three garment types in the dataset, T-shirts, shorts and pants. For each garment type k , we need to train the model parameters $\{\mathbf{A}^k, \bar{\mathbf{d}}^k\}$ from a collection of clothing offsets, denoted by $\mathbf{X}^k \in \mathbb{R}^{3n_v \times n_k}$, where n_k is the number of samples for garment type k in the dataset. To obtain each sample in the collection, we follow [54, 35] to register the raw scan with SMPL model. This operation not only brings the raw scan data into the same topology, but also “unposes” the human shape with clothing into the rest pose, denoted by \mathbf{X}^c . We also follow [54] to estimate the underlying body shape of the subject in rest pose, denoted by \mathbf{X}^b . In addition, we obtain a per-vertex binary mask σ^k that has value 1 for the region of garment type k and 0 for any other regions (skin and other clothing types) by rendering the meshes to images and applying a state-of-the-art clothing segmentation algorithm [14]. Then we obtain the clothing offset data \mathbf{X}^k for garment type k by

$$\mathbf{X}^k = (\mathbf{X}^c - \mathbf{X}^b) \odot \sigma^k, \quad (8)$$

where \odot denotes the element-wise multiplication. We use a standard PCA training algorithm based on Singular Value Decomposition (SVD), leaving $n_z = 50$ bases in our model. We refer readers to the original papers [54, 35] for details on scan registration and underlying body shape estimation.

5. Monocular Clothing Capture

Given the pre-trained clothing models, we now present our approach for temporally coherent clothing capture from only a monocular video.

5.1. Body Motion Estimation

In the first stage, we estimate underlying body motion in 3D with the SMPL body model. We estimate per-frame SMPL pose parameters θ_i and global translation $\mathbf{t}_i \in \mathbb{R}^3$, together with SMPL shape parameters β across the whole sequence. Meanwhile, we estimate the camera intrinsics \mathbf{K} of a full perspective projection model for all the frames. In order to achieve good robustness under different in-the-wild scenarios, we integrate a variety of different image measurements into an energy optimization problem. Formally, we solve the following minimization problem:

$$\min_{\beta, \{\theta_i, \mathbf{t}_i\}_{i=1}^F, \mathbf{K}} E^b = E_{2d}^b + E_{dp}^b + E_{sil}^b + E_{pof}^b + E_{reg}^b. \quad (9)$$

In particular, E_{2d}^b is the squared L_2 error between projected SMPL joints and 2D keypoint detection from OpenPose [46, 11, 10]. E_{dp}^b , also used in Guler *et al.* [15], is an energy term for dense correspondence estimation from DensePose [5]. Specifically, for any pixel \mathbf{p} in the image with DensePose prediction, we identify the corresponding SMPL vertex index $j(\mathbf{p})$ and optimize an energy term defined as

$$E_{dp}^b = \frac{1}{F} \sum_{i=1}^F \sum_{\mathbf{p}} \|\Pi(M_{j(\mathbf{p})}(\beta, \theta_i) + \mathbf{t}_i; \mathbf{K}) - \mathbf{p}\|^2, \quad (10)$$

where Π denotes the projection function determined by the camera intrinsics \mathbf{K} . E_{sil}^b is the silhouette matching term.

We extract silhouettes \mathbf{S}_i of our SMPL body mesh with a differentiable renderer [24], and obtain the target silhouette $\hat{\mathbf{S}}_i$ from a clothing segmentation algorithm [14]. We use an Intersection-over-Union error [24]

$$E_{\text{sil}}^b = \frac{1}{F} \sum_{i=1}^F \left(1 - \frac{\|\mathbf{S}_i \odot \hat{\mathbf{S}}_i\|_1}{\|\mathbf{S}_i + \hat{\mathbf{S}}_i - \mathbf{S}_i \odot \hat{\mathbf{S}}_i\|_1} \right). \quad (11)$$

E_{por}^b is an error term based on 3D orientation between adjacent joints in the body skeleton hierarchy. We match the spatial orientation of SMPL body joints to the prediction of Part Orientation Field (POF) similar to [48]. We refer readers to the original papers [48, 26] for details. We also apply regularization on our estimation, denoted by E_{reg}^b , which consists of a Mixture of Gaussian prior for body pose $\{\theta_i\}_{i=1}^F$ [8], L_2 regularization on the shape parameters β , and temporal smoothness terms to reduce motion jitters.

After solving the energy optimization, we obtain a temporally consistent body mesh for every frame by $\mathbf{M}_i^b = M(\beta, \theta_i)$. We fix the SMPL parameters $\beta, \{\theta_i, \mathbf{t}_i\}_{i=1}^F$ and camera parameters \mathbf{K} during later stages of our pipeline. The estimated body meshes provide a strong guidance for the subsequent estimation of clothing deformation.

5.2. Clothing Deformation Capture

We now illustrate our proposed method to capture clothing deformation. Compared to previous work [49, 16, 17] where a pre-scanned template of the subject is assumed, this problem is significantly more challenging due to the lack of strong shape prior to resolve the single-view 3D ambiguity, and the lack of a pre-defined personalized texture that provides correspondence for surface tracking. To solve this problem, we (1) exploit the deformation space learned in our clothing models and (2) progressively extract a personalized texture from the input image sequence to enable surface tracking across time and reduce drifting.

We perform clothing capture in a sequential manner. For each frame i , we estimate per-frame clothing parameters for clothing on the upper and lower body $\mathbf{z}_i = \{\mathbf{z}_i^u, \mathbf{z}_i^l\}$ given the input image \mathbf{I}_i , initializing from the result of the previous frame \mathbf{z}_{i-1} . We formulate the task as solving an energy optimization problem, formally,

$$\min_{\mathbf{z}_i} E^c = E_{\text{sil}}^c + E_{\text{seg}}^c + E_{\text{photo}}^c + E_{\text{reg}}^c. \quad (12)$$

Now we explain each cost term individually. An illustration of the different cost terms is shown in Fig. 4.

Silhouette matching term E_{sil}^b : Similar to E_{sil}^b in the first stage, we use a differentiable renderer to match the silhouette of our rendering output with a target silhouette extracted from the original image. Differently, here what we compare with target silhouette is the silhouette of human shape with clothing $M^c(\beta, \theta_i, \mathbf{z}_i)$, instead of bare body shape $M(\beta, \theta_i)$ in the previous stage.

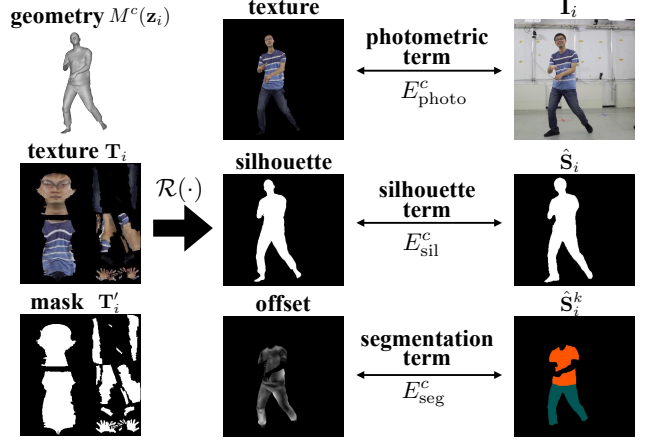


Figure 4: Explanation of the energy terms used for clothing capture. We obtain the rendered texture, silhouette and clothing offset with a differentiable renderer, which are compared with target images using different energy terms.

Clothing segmentation term E_{seg}^c : Clothing segmentation [14] provides not only the overall silhouette of the person, but also the boundary between different garment regions and exposed skin in the image. We utilize this information by penalizing the clothing offset on vertices whose projection falls outside the segmentation region. Concretely, the differentiable renderer is used to render the per-vertex offset fields $D^u(\mathbf{z}_i^u)$, $D^l(\mathbf{z}_i^l)$ on the clothed mesh $M^c(\mathbf{z}_i)$ ²; we denote the output by $\mathcal{R}(D^k(\mathbf{z}_i^k))$, where k represents either u for upper clothing or l for lower clothing. In addition, we denote the segmentation masks for the clothing region $k \in \{u, l\}$ by $\hat{\mathbf{S}}_i^k$. Then we have

$$E_{\text{seg}}^c = \sum_{k \in \{u, l\}} \sum_{\mathbf{p}} (1 - \hat{\mathbf{S}}_i^k) \odot \|\mathcal{R}(D^k(\mathbf{z}_i^k))\|^2, \quad (13)$$

where \mathbf{p} iterates over all pixels in the image. Effectively, for each clothing type we penalize the clothing offset outside the corresponding clothing region, where $\hat{\mathbf{S}}_i^k$ is 0. The gradient in the image domain is propagated to the mesh by the differentiable renderer \mathcal{R} .

Photometric tracking term E_{photo}^c : This term is introduced to estimate temporal correspondence more accurately, especially when the garments we capture have high-contrast texture. We progressively build a personalized RGB texture image \mathbf{T}_i in a pre-defined UV space of SMPL model, along with a binary mask \mathbf{T}'_i that indicates texels where RGB values in \mathbf{T}_i have been identified. The photometric tracking term is defined to compare the rendered output of our clothing models using \mathbf{T}_i with the input image. For this purpose

²SMPL body parameters β, θ_i and \mathbf{t}_i are fixed in this stage thus omitted here.

we use a differential renderer that works with UV texture [12, 20], and denotes the rendered output as $\mathcal{R}(\mathbf{T}_i)$. We also render the mesh with \mathbf{T}'_i to indicate pixels where texture from \mathbf{T}_i is available. Formally, we have

$$E_{\text{photo}}^c = \sum_{\mathbf{p}} \|\mathcal{R}(\mathbf{T}_i) - \mathbf{I}_i\|^2 \odot \mathcal{R}(\mathbf{T}'_i), \quad (14)$$

where the summation is taken over the pixels in the image. After the optimization in Eq. 12 is solved for frame i , we update $\mathbf{T}_i, \mathbf{T}'_i$ to obtain $\mathbf{T}_{i+1}, \mathbf{T}'_{i+1}$, which will be used for solving optimization for frame $i+1$. To achieve this, we project \mathbf{I}_i to the mesh surface and fill in new RGB values to UV texels in \mathbf{T}_i where no previous values have been identified, indicated by 0s in \mathbf{T}'_i . Corresponding texels in \mathbf{T}'_i are also set to 1 to obtain \mathbf{T}'_{i+1} . This process is initialized by setting \mathbf{T}_1 and \mathbf{T}'_1 to 0; in other words, the photometric tracking term takes no effect for the first frame in the sequence, since no texture has been extracted.

Regularization term E_{reg}^c : Our clothing deformation models are PCA-based linear models. They may produce unreasonable shapes when the parameters \mathbf{z} are large. Therefore we apply regularization on the cloth parameters using an adaptive cost function ρ that penalizes large input values:

$$E_{\text{reg}}^c = \rho(\|\mathbf{z}_i\|). \quad (15)$$

The sequential tracking stage is then followed by a batch optimization stage that optimize for all F frames in the sequence together. The energy function we use is the same as the previous stage (Eq. 12) with an additional term that penalizes too drastic temporal change of clothing parameters, which helps to produce temporally stable results. This term is defined as

$$E_{\text{temp}}^c = \frac{1}{F-1} \sum_{i=1}^{F-1} \|\mathbf{z}_{i+1} - \mathbf{z}_i\|^2. \quad (16)$$

The output of batch optimization stage is a sequence of capture results with clothing $\{\mathbf{M}_i^c = M^c(\beta, \theta_i, \mathbf{z}_i)\}_{i=1}^F$.

5.3. Wrinkle Detail Extraction

Up to the batch optimization stage, we can capture large clothing deformation. However, the results are limited by low mesh resolution and unable to capture the fine-grained wrinkles on the clothing. Therefore, the last stage of our approach is to extract wrinkle details from the original images and apply them to our coarsely tracked meshes.

Traditionally, such wrinkle details are captured with Shape from Shading (SfS) [47, 2]. For in-the-wild monocular clothing capture, we empirically find it difficult to extract wrinkles reliably by SfS due to complex garment albedo, large variation of lighting conditions and self-shadowing. Recently, we observed the success of learning-based approaches in estimating accurate surface normal for human

appearance using neural networks [43, 40]. The estimated surface normal provides strong and direct clues on how the wrinkles should be added to our clothing capture results in order to match the original images.

Formally, let us denote the output of a surface normal estimation network for frame i to be \mathbf{I}_i^n , a 3-channel normal map for each pixel in the original image. We first subdivide the mesh \mathbf{M}_i^c with Loop subdivision to increase the spatial resolution, with the subdivided mesh denoted by \mathbf{M}_i^s . Then, we solve for a deformed mesh \mathbf{O}_i whose rendered normal map matches the estimated normal map \mathbf{I}_i^n in the garment region. We denote the rendered normal output by $\mathcal{R}^n(\mathbf{O}_i)$, where \mathcal{R}^n is the differential renderer. We solve the following optimization problem for each frame i individually:

$$\min_{\mathbf{O}_i} E^w = E_{\text{normal}}^w + E_{\text{reg}}^w + E_{\text{tpl}}^w, \quad (17)$$

$$E_{\text{normal}}^w = \sum_{\mathbf{p} \in \hat{\mathbf{S}}_i^c} \|\nabla \mathcal{R}^n(\mathbf{O}_i) - \nabla \mathbf{I}_i^n\|^2, \quad (18)$$

$$E_{\text{reg}}^w = \|\mathbf{O}_i - \mathbf{M}_i^s\|^2, \quad E_{\text{tpl}}^w = \|\mathbf{L}\mathbf{O}_i\|^2. \quad (19)$$

where ∇ denotes the image gradient operator, and \mathbf{L} denotes the mesh Laplacian operator, $\hat{\mathbf{S}}_i^c = \hat{\mathbf{S}}_i^u \cup \hat{\mathbf{S}}_i^l$ is the union of all pixels in the clothing segmentation masks. Here we penalize the difference between normal maps in the image gradient domain to be more tolerant to error in absolute normal direction from the neural network. We also restrict the deformation of \mathbf{O}_i from \mathbf{M}_i^s be in the direction of the camera rays. Here, we use the implementation of differentiable renderer in [12, 20] and surface normal network in [40]. The final results are the deformed meshes $\{\mathbf{O}_i\}_{i=1}^F$.

6. Quantitative Evaluation

In this section, we present the results of quantitative evaluation. We use a benchmark sequence from the MonoPerfCap dataset [49] and video sequences rendered from BUFF dataset [54] to test the performance of our method.

6.1. Evaluation on MonoPerfCap Dataset

Experiment Setting. We follow previous work [49] to use the *Pablo* sequence in their dataset to perform quantitative comparison. Surface meshes and 3D joints obtained by a multi-view performance capture method [38] are provided as the ground truth in the dataset. We compare our method with a state-of-the-art template-based performance capture method [49]³ and single-image human reconstruction methods [55, 39, 40, 4, 57]. Body pose is not estimated in [4], so we apply our estimated body pose to their T-pose results.

Evaluation of Clothing Surface Reconstruction. We first evaluate our method using a surface reconstruction metric.

³Monocular capture results are provided by the author.

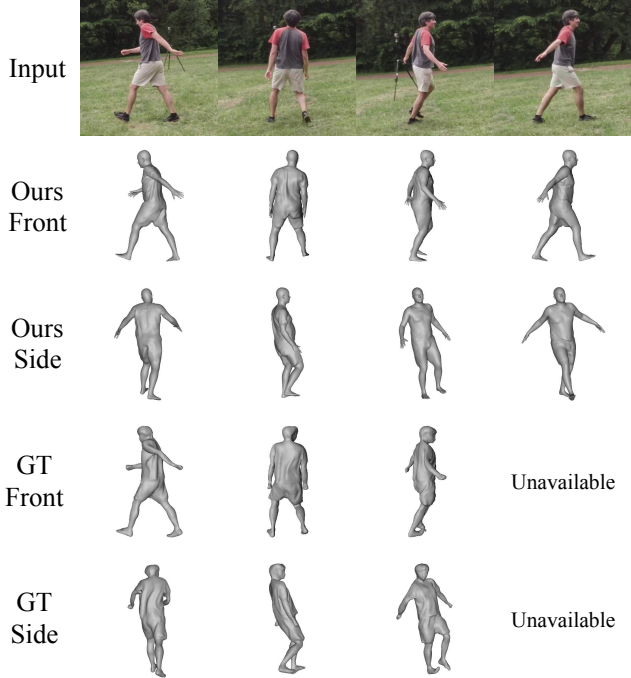


Figure 5: Visualization of experiment results on the *Pablo* sequence. From top to bottom we show original images, our results from the front and side views, and ground truth from the front and side views. The ground truth mesh for the last frame is not provided in the dataset.

Due to the intrinsic depth-scale ambiguity of single-view reconstruction, we compute a global scaling factor from our result to the ground truth, which is applied to our result before comparison. Following [49] we align our results to the ground truth with a translation to eliminate the global depth offsets. We compute the average point-to-surface distance from all the ground truth vertices in the clothing region to the output mesh as the evaluation metric. The clothing region (the T-shirt and shorts) is obtained by manual segmentation on the ground truth surface mesh. The same procedure is applied to all the methods under evaluation. A visualization of our results is shown in Fig. 5.

We report the mean surface error averaged across all frames in the middle column of Table 1. For the visualization of per-frame error curves please refer to our supplementary material. Our method achieves a significantly lower surface error compared to all previous single-image surface reconstruction methods. Our performance even comes close to the template-based tracking method [49] which requires a pre-scanned personalized template that provides strong prior information about the body and clothing shape of the subject. By contrast, our method does not require a pre-processed template, and therefore can be applied to a wider range of videos.

Methods	Surface Error	Joint Error
MonoPerfCap* [49]	14.6	118.7
HMD [57]	31.9	-
Tex2Shape [4]	27.7	-
DeepHuman [55]	24.2	-
PIFu [39]	30.5	-
PIFuHD [40]	26.5	-
Ours	17.9	77.3

Table 1: Quantitative comparison with previous work on *Pablo* sequence using mean point-to-surface error and mean joint error across frames. All numbers are in mm. The method annotated with ‘*’ uses a pre-scanned personalized template that provides a strong shape prior. Please see our supplementary material for per-frame error.



Figure 6: Visualization of three viewpoints in the BUFF evaluation. In each view we show the input image, body estimation result, and clothing capture result.

Evaluation of 3D Pose Estimation. Although body pose estimation is not a focus of this paper, we follow [49] to validate our method on the metric of 3D joint error on the *Pablo* sequence. Average per-joint 3D position error after alignment with translation is reported in Table 1 (right). Our method achieves an error of 77.3 mm, significantly lower than 118.7 mm in [49]. This verifies the effectiveness of our body pose initialization that utilizes various image measurements including 2D joints, dense correspondences, silhouette, etc.

6.2. Evaluation on BUFF Dataset

Experiment Setting. BUFF [54] is a dataset of high-resolution 4D textured scan sequences of five people. In this experiment, we sample a test sequence from the BUFF dataset (00096-shortlong_hips, first 200 frames) and train a pair of upper and lower clothing models with the data of four other people. We render the sequence from three views: front, left and front-left, as visualized in Fig. 6. We evaluate our method in four stages: body initialization, sequential tracking, batch optimization and wrinkle extraction. The evaluation protocol is the same as Section 6.1: we rigidly align the estimated and ground truth meshes with a global scaling and translation, and compute average distance from ground truth clothing vertices to our results.

Results. The quantitative results are shown in Table 2. First, from all three viewpoints, results with clothing con-



Figure 7: Examples of our clothing capture results. In each example, we show the input image, capture results from the front view and the side view. Please see our supplementary video for complete results.

	Front	Front-left	Left
Body only	29.4	30.3	29.2
Clothed w/o batch	26.8	25.5	24.7
Clothed w/ batch	26.7	25.3	24.7
Clothed w/ wrinkle	26.8	25.5	24.9

Table 2: Quantitative ablation study of different stages of our method on the rendered BUFF dataset using mean point-to-surface error. All numbers are in mm. Please see our supplementary material for per-frame error.

sistently achieve lower reconstruction error than body only. This verifies that our method captures clothing shape that cannot be explained by the SMPL body shape space. Second, we can see that temporal smoothing and wrinkle extraction, which improve the visual quality as shown in qualitative results, have little influence on the reconstruction error. Third, our results show similar range of error in the clothing region across different views, implying that our method is not very sensitive to the viewpoint variation.

7. Qualitative Evaluation

We qualitatively evaluate our method on various videos including public benchmark and in-the-wild videos where no pre-scanned template is available. Example results are shown in Figure 7. Please see our supplementary video for full results and qualitative comparison with other work.

As shown in the supplementary video, our result not only demonstrates better temporal robustness than the single-image 3D human reconstruction methods in terms of reconstructed surfaces, but also provides 3D temporal correspondences effectively shown by the re-rendering of our output mesh with a consistent texture map. This is hard to obtain by methods that regress 3D shape in voxels [44, 55], depth

maps [43] or implicit functions [30, 39, 40]. Template-based monocular performance capture methods [49, 16] rely heavily on non-rigid surface regularization such as As-Rigid-As-Possible (ARAP), which often prevents those methods from capturing natural dynamics of the clothing deformation. In comparison, our method is able to capture more realistic dynamics of the garment with regularization provided by the clothing models.

In addition, we perform extensive ablation studies on various loss terms used in our pipeline. Please refer to our supplementary document for the results.

8. Conclusion and Future Work

In this paper, we have presented a method to capture temporally coherent dynamic deformation of clothing from a monocular video. To the best of our knowledge, we have shown the first result of temporally coherent clothing capture from a monocular RGB video without using a pre-scanned template. Our results on various in-the-wild videos endorse the effectiveness and robustness of our method.

Our method is limited by the types of garments in the available training data. We have demonstrated results on several types of tight clothing. Treatment of free-flowing garments like skirts requires collection of more data and additional design of the clothing model. Our method is constrained in the ability to capture drastically changing deformations due to the limited expressiveness of our models, which may be addressed by using higher-capacity models like a deep neural network. We also would like to further incorporate physics into the clothing models to enable more physically realistic clothing capture.

Acknowledgements. We would like to thank Eric Yu for his help with the rendering of our results using Blender.

References

- [1] T. Alldieck, M. Magnor, B. L. Bhatnagar, C. Theobalt, and G. Pons-Moll. Learning to reconstruct people in clothing from a single rgb camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1175–1186, 2019. [2](#)
- [2] T. Alldieck, M. Magnor, W. Xu, C. Theobalt, and G. Pons-Moll. Detailed human avatars from monocular video. In *2018 International Conference on 3D Vision (3DV)*, pages 98–109. IEEE, 2018. [2](#), [6](#)
- [3] T. Alldieck, M. Magnor, W. Xu, C. Theobalt, and G. Pons-Moll. Video based reconstruction of 3d people models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8387–8397, 2018. [2](#)
- [4] T. Alldieck, G. Pons-Moll, C. Theobalt, and M. Magnor. Tex2shape: Detailed full human body geometry from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2293–2303, 2019. [1](#), [2](#), [6](#), [7](#)
- [5] R. Alp Güler, N. Neverova, and I. Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018. [4](#)
- [6] D. Baraff and A. Witkin. Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 43–54, 1998. [2](#)
- [7] B. L. Bhatnagar, G. Tiwari, C. Theobalt, and G. Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *IEEE International Conference on Computer Vision (ICCV)*. IEEE, oct 2019. [2](#)
- [8] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578. Springer, 2016. [2](#), [5](#)
- [9] R. Bridson, S. Marino, and R. Fedkiw. Simulation of clothing with folds and wrinkles. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 28–36, 2003. [2](#)
- [10] Z. Cao, G. H. Martinez, T. Simon, S.-E. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. [2](#), [4](#)
- [11] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7291–7299, 2017. [2](#), [4](#)
- [12] W. Chen, H. Ling, J. Gao, E. Smith, J. Lehtinen, A. Jacobson, and S. Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. In *Advances in Neural Information Processing Systems*, pages 9605–9616, 2019. [2](#), [6](#)
- [13] V. Gabeur, J.-S. Franco, X. Martin, C. Schmid, and G. Rogez. Moulding humans: Non-parametric 3d human shape estimation from single images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2232–2241, 2019. [2](#)
- [14] K. Gong, Y. Gao, X. Liang, X. Shen, M. Wang, and L. Lin. Graphonomy: Universal human parsing via graph transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7450–7459, 2019. [4](#), [5](#)
- [15] R. A. Guler and I. Kokkinos. Holopose: Holistic 3d human reconstruction in-the-wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10884–10894, 2019. [4](#)
- [16] M. Habermann, W. Xu, M. Zollhofer, G. Pons-Moll, and C. Theobalt. Livecap: Real-time human performance capture from monocular video. *ACM Transactions on Graphics (TOG)*, 38(2):1–17, 2019. [1](#), [2](#), [3](#), [5](#), [8](#)
- [17] M. Habermann, W. Xu, M. Zollhofer, G. Pons-Moll, and C. Theobalt. Deepcap: Monocular human performance capture using weak supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5052–5063, 2020. [1](#), [2](#), [3](#), [5](#)
- [18] Z. Huang, T. Li, W. Chen, Y. Zhao, J. Xing, C. LeGendre, L. Luo, C. Ma, and H. Li. Deep volumetric video from very sparse multi-view performance capture. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 336–354, 2018. [3](#)
- [19] Z. Huang, Y. Xu, C. Lassner, H. Li, and T. Tung. Arch: Animatable reconstruction of clothed humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3093–3102, 2020. [2](#)
- [20] K. M. Jatavallabhula, E. Smith, J.-F. Lafleche, C. F. Tsang, A. Rozantsev, W. Chen, and T. Xiang. Kaolin: A pytorch library for accelerating 3d deep learning research. *arXiv preprint arXiv:1911.05063*, 2019. [6](#)
- [21] H. Joo, T. Simon, and Y. Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8320–8329, 2018. [2](#)
- [22] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7122–7131, 2018. [1](#), [2](#)
- [23] J. Li, G. Daviet, R. Narain, F. Bertails-Descoubes, M. Overby, G. E. Brown, and L. Boissieux. An implicit frictional contact solver for adaptive cloth simulation. *ACM Transactions on Graphics (TOG)*, 37(4):1–15, 2018. [2](#)
- [24] S. Liu, T. Li, W. Chen, and H. Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7708–7717, 2019. [2](#), [5](#)
- [25] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. [1](#), [2](#), [3](#)
- [26] C. Luo, X. Chu, and A. Yuille. Orinet: A fully convolutional network for 3d human pose estimation. In *BMVC*, 2018. [5](#)
- [27] Q. Ma, J. Yang, A. Ranjan, S. Pujades, G. Pons-Moll, S. Tang, and M. J. Black. Learning to dress 3d people in generative clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6469–6478, 2020. [2](#), [3](#)

- [28] D. Mehta, O. Sotnychenko, F. Mueller, W. Xu, M. Elgharib, P. Fua, H.-P. Seidel, H. Rhodin, G. Pons-Moll, and C. Theobalt. Xnect: Real-time multi-person 3d motion capture with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 39(4):82–1, 2020. 3
- [29] D. Mehta, S. Sridhar, O. Sotnychenko, H. Rhodin, M. Shafiei, H.-P. Seidel, W. Xu, D. Casas, and C. Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 36(4):1–14, 2017. 3
- [30] R. Natsume, S. Saito, Z. Huang, W. Chen, C. Ma, H. Li, and S. Morishima. Siclope: Silhouette-based clothed people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4480–4490, 2019. 1, 2, 8
- [31] M. Omran, C. Lassner, G. Pons-Moll, P. Gehler, and B. Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *2018 international conference on 3D vision (3DV)*, pages 484–494. IEEE, 2018. 1, 2
- [32] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10975–10985, 2019. 2
- [33] G. Pavlakos, N. Kolotouros, and K. Daniilidis. Texturepose: Supervising human mesh estimation with texture consistency. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 803–812, 2019. 1, 2
- [34] G. Pavlakos, X. Zhou, K. G. Derpanis, and K. Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7025–7034, 2017. 2
- [35] G. Pons-Moll, S. Pujades, S. Hu, and M. J. Black. Clothcap: Seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics (TOG)*, 36(4):1–15, 2017. 2, 4
- [36] H. Rhodin, N. Robertini, D. Casas, C. Richardt, H.-P. Seidel, and C. Theobalt. General automatic human shape and motion capture using volumetric contour cues. In *European conference on computer vision*, pages 509–526. Springer, 2016. 3
- [37] N. Robertini, F. Bernard, W. Xu, and C. Theobalt. Illumination-invariant robust multiview 3d human motion capture. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1661–1670. IEEE, 2018. 3
- [38] N. Robertini, D. Casas, H. Rhodin, H.-P. Seidel, and C. Theobalt. Model-based outdoor performance capture. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 166–175. IEEE, 2016. 3, 6
- [39] S. Saito, Z. Huang, R. Natsume, S. Morishima, A. Kanazawa, and H. Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2304–2314, 2019. 1, 2, 6, 7, 8
- [40] S. Saito, T. Simon, J. Saragih, and H. Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020. 1, 2, 6, 7, 8
- [41] L. Sigal, M. Mahler, S. Diaz, K. McIntosh, E. Carter, T. Richards, and J. Hodgins. A perceptual control space for garment simulation. *ACM Transactions on Graphics (TOG)*, 34(4):1–10, 2015. 2
- [42] X. Sun, B. Xiao, F. Wei, S. Liang, and Y. Wei. Integral human pose regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 529–545, 2018. 2
- [43] S. Tang, F. Tan, K. Cheng, Z. Li, S. Zhu, and P. Tan. A neural network for detailed human depth estimation from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7750–7759, 2019. 1, 2, 6, 8
- [44] G. Varol, D. Ceylan, B. Russell, J. Yang, E. Yumer, I. Laptev, and C. Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 20–36, 2018. 1, 2, 8
- [45] P. Volino, N. Magnenat-Thalmann, and F. Faure. A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Transactions on Graphics (TOG)*, 28(4):1–16, 2009. 2
- [46] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4724–4732, 2016. 2, 4
- [47] C. Wu, C. Stoll, L. Valgaerts, and C. Theobalt. On-set performance capture of multiple actors with a stereo camera. *ACM Transactions on Graphics (TOG)*, 32(6):1–11, 2013. 6
- [48] D. Xiang, H. Joo, and Y. Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10965–10974, 2019. 2, 5
- [49] W. Xu, A. Chatterjee, M. Zollhöfer, H. Rhodin, D. Mehta, H.-P. Seidel, and C. Theobalt. Monoperfcap: Human performance capture from monocular video. *ACM Transactions on Graphics (TOG)*, 37(2):1–15, 2018. 1, 2, 3, 5, 6, 7, 8
- [50] Y. Xu, S.-C. Zhu, and T. Tung. Denserac: Joint 3d pose and shape estimation by dense render-and-compare. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7760–7770, 2019. 1, 2
- [51] J. Yang, J.-S. Franco, F. Hétroy-Wheeler, and S. Wuhler. Analyzing clothing layer deformation statistics of 3d human motions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 237–253, 2018. 2, 3
- [52] T. Yu, Z. Zheng, K. Guo, J. Zhao, Q. Dai, H. Li, G. Pons-Moll, and Y. Liu. Doublefusion: Real-time capture of human performances with inner body shapes from a single depth sensor. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7287–7296, 2018. 3
- [53] T. Yu, Z. Zheng, Y. Zhong, J. Zhao, Q. Dai, G. Pons-Moll, and Y. Liu. Simulcap: Single-view human performance capture with cloth simulation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5499–5509. IEEE, 2019. 3
- [54] C. Zhang, S. Pujades, M. J. Black, and G. Pons-Moll. Detailed, accurate, human shape estimation from clothed 3d

- scan sequences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4191–4200, 2017. 4, 6, 7
- [55] Z. Zheng, T. Yu, Y. Wei, Q. Dai, and Y. Liu. Deephuman: 3d human reconstruction from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7739–7749, 2019. 1, 2, 6, 7, 8
- [56] X. Zhou, Q. Huang, X. Sun, X. Xue, and Y. Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 398–407, 2017. 2
- [57] H. Zhu, X. Zuo, S. Wang, X. Cao, and R. Yang. Detailed human shape estimation from a single image by hierarchical mesh deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4491–4500, 2019. 6, 7