

Deep Physics-aware Inference of Cloth Deformation for Monocular Human Performance Capture

Yue Li^{1,2,3}, Marc Habermann^{1,2}, Bernhard Thomaszewski³, Stelian Coros³, Thabo Beeler⁴, and Christian Theobalt^{1,2}

¹Max Planck Institute for Informatics, ²Saarland Informatics Campus, ³ETH Zurich, ⁴Google Inc.

Abstract

Recent monocular human performance capture approaches have shown compelling dense tracking results of the full body from a single RGB camera. However, existing methods either do not estimate clothing at all or model cloth deformation with simple geometric priors instead of taking into account the underlying physical principles. This leads to noticeable artifacts in their reconstructions, e.g. baked-in wrinkles, implausible deformations that seemingly defy gravity, and intersections between cloth and body. To address these problems, we propose a person-specific, learning-based method that integrates a simulation layer into the training process to provide for the first time physics supervision in the context of weakly supervised deep monocular human performance capture. We show how integrating physics into the training process improves the learned cloth deformations, allows modeling clothing as a separate piece of geometry, and largely reduces cloth-body intersections. Relying only on weak 2D multi-view supervision during training, our approach leads to a significant improvement over current state-of-the-art methods and is thus a clear step towards realistic monocular capture of the entire deforming surface of a clothed human.

1. Introduction

Human performance capture plays a critical role in various computer graphics and vision applications such as virtual try-on, movies as well as video games. With rapid progress in display and capture technology, expectations on the quality of geometric reconstruction and tracking are constantly increasing. Here, not only the geometric details are of major importance but also that the deformed and posed reconstructions follow the physical behavior of real objects which includes realistic wrinkle patterns as well as coherent interaction of body and clothing. While professional content production studios can rely on involved multi-camera setups to capture high-fidelity human performances, there is an ever-growing desire to democratize per-

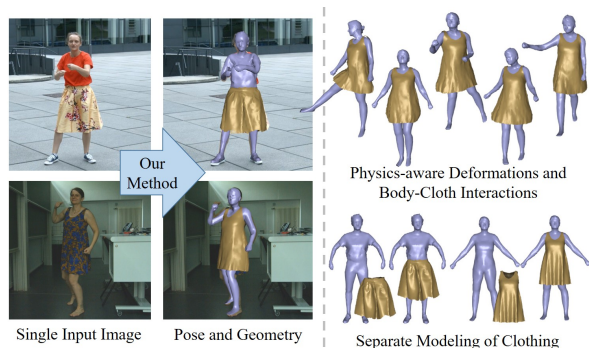


Figure 1. Our method estimates body pose and physically plausible surface deformation from a single image. Importantly, body and clothing are represented as separate meshes allowing for accurate modeling of body-cloth interactions.

formance capture for everyday applications, e.g. virtual try-on, by utilizing much simpler and cheaper capture devices.

Hence, research has shifted from expensive and complex multi-view capture setups [52, 6, 9, 7, 50, 13, 16, 38, 44, 73, 77, 50] to depth cameras [61, 21, 46, 27, 20, 35, 80, 15, 82, 81, 76, 83] over the past decade. Unfortunately, the latter are sensitive to bright sunlight and thus are not suited for outdoor use-cases. In conjunction with the advances in deep learning, the most recent research has shifted its attention onto *single RGB camera* setups, offering the most flexible and low-cost setup. Previous monocular methods have made a substantial progress in recovering the 3D unclothed body [30, 49, 31], hand pose [74, 43, 88], facial identity and expression [33, 67, 68] as well as jointly tracking all of those [48, 78, 29, 87]. However, only a few methods [23, 22, 79] coherently track the dense surface deformations with clothing included from monocular views, which is essential for a majority of applications. These person-specific methods densely deform and pose a geometry to match the body pose and the clothing deformation in the input image while assuming an initial template of the person is given. Recent learning-based monocular methods [23] only leverage image-based supervision, rendering it a challenging task to densely supervise deformations. This man-

ifests in simplified model assumptions, *e.g.*, a single geometry for both body and clothing, or the utilization of simple geometric priors disobeying physics principles. Consequently, they either fail to account for body-cloth interactions or having deformations that do not follow physical rules. Importantly, artifacts such as baked-in wrinkles from the initial scan are highly noticeable in their results.

To this end, we propose a learning-based approach for capturing the body pose and the physically plausible clothing deformation from a single RGB image (see Fig. 1). Our method comprises two networks dedicated to regress body pose in terms of joint angles and surface deformations in form of embedded deformation. Importantly, during training, we only assume weak supervision with multi-view imagery, *i.e.* 2D skeletal joint detection, and foreground masks. These supervisions alone can hardly ensure physically plausible results. Thus, at the core of our approach, we propose an efficient simulation layer that for the first time allows physically plausible self-supervision *during* the training in such a weakly supervised setting. We achieve this by integrating a physics-based simulator into a learning architecture that takes intermediate predictions of cloth and body positions and velocities to perform forward simulations. The simulation results are then used to supervise the cloth deformations during training. As cloth-body collisions are explicitly handled in the proposed layer, we can accurately model clothing as a separate piece of geometry in contrast to previous monocular methods. In summary, our contributions are:

- A monocular human performance capture approach, which outputs body pose and physically plausible cloth deformations for dressed subjects.
- A simulation network layer that allows on-the-fly simulation supervision *during* training, which also enables separate modeling of cloth and body geometry.

In contrast to prior work, our method reconstructs physically more accurate deformations without baked-in wrinkles and with correct body-cloth collision handling. Our quantitative evaluations indicate that incorporating physics-based simulation during training provides significant improvements over state-of-the-art methods.

2. Related Work

As our goal is recovering a dense surface of the human, we focus on previous works that achieve this by using parametric body models or template meshes, and works that treat body and clothing as separate mesh layers. We omit the works on 2D [10, 11, 60, 75] and 3D skeletal pose estimation [42, 41, 24, 54, 53, 86, 66, 70, 55] as they are not concerned with the problem of surface reconstruction.

Reconstruction of Parametric Body Models. The works [85, 28, 56, 18, 4, 37, 30, 71] that fall into this category use parametric body models [39]. Some works fit the model parameters to sparse 2D and 3D joint predictions [4] or regressed vertex positions [34] by minimizing corresponding energies. Others [30] directly regress these parameters from images. A set of recent works [48, 78] extended body models to account for varying hand poses and facial expressions to jointly capture hands, face, and body. While motion and shape of the undressed body are reconstructed, clothing is not considered.

Unified Reconstruction. One stream of previous work treats body and clothing as a single geometry. Volumetric representations [84, 72] use an occupancy grid to represent the body, meaning that the resolution is limited by the grid. Implicit methods [57, 26, 58] methods overcome this limitation by treating the surface as an implicit function. However, both approaches require post-processing to recover explicit surface representations. Lacking temporal consistency thus prohibits these approaches for applications such as texture replacement or motion retargeting. Closely related to our work are template-based methods [22, 23, 79, 12, 14] that track a template based on image observations. Using a mesh with fixed topology as a reference, surface correspondence over time is explicitly given. With input data originating from images only, this ill-posed problem is countered by simplified assumptions and geometric priors. Consequently, static wrinkles contained in a template remain visible across all poses, and deformations commonly appear to be physically implausible, *e.g.* defying gravity. Most importantly, all these methods treat clothing and body as a single piece of geometry ignoring dynamic body-cloth interactions. To address these limitations, we propose a simulation layer that encourages cloth deformations to not only satisfy image constraints but also exhibit physically plausible behavior.

Cloth as a Separate Part. In contrast to the above methods, there is also a line of work that reconstructs body and clothing as separate geometries. Bhatnagar *et al.* [3] recover static geometry for clothing and body from a set of RGB images. DeepWrinkles [36] enables posing a piece of cloth where their method learns to regress pose-dependent wrinkles at high resolution. ClothCap [51] uses multi-view capture to produce a clothed human body that can be used for re-targeting. Stoll *et al.* [63] recover cloth material parameters from multi-view video sequences to reproduce the observed garment deformation. SimulCap [83] performs quasi-static physics simulation with depth matching constraints to reconstruct the clothing layer. Different from the above methods, our approach relies solely on a monocular RGB camera. Also leveraging simulation, MulayCap [64]

recovers, both, the texture and the geometry of a dressed subject from monocular RGB videos by a multi-layer decomposition approach. However, simulation is only used to generate initialization for the succeeding refinement stage, whereas we consistently enforce simulation supervision.

As a potential alternative to simulation, geometric detail such as wrinkles can be added in a data-driven, pose-dependent manner [17, 59, 47, 19]. Different from these geometry-driven methods, we integrate physics-based simulation into our training framework thus encouraging physical plausibility with only a single image as input.

3. Method

Our template-based method leverages a deep neural architecture, taking a single background-segmented person image as input and regresses posed and deformed surface meshes for body and clothing which match the performance in the input image (Fig. 2). Before training, a 3D template of the person with separate cloth and body geometry and a multi-view recording of the subject performing various motions has to be acquired (Sec. 3.1). The technical core of our architecture is formed by two prediction networks, *PoseNet* and *PADefNet*, that are trained to regress body pose and physics-aware cloth deformation, respectively (Sec. 3.3). *PoseNet* [23] regresses skeleton joint angles and the root rotation from the input image using multi-view 2D joint detection as weak supervision. The proposed *PADefNet* predicts the surface deformation of the cloth template by regressing embedded graph parameters from the same input image. In addition to multi-view image data, *PADefNet* leverages our cloth simulation layer as supervision, which encourages physically plausible deformations (Sec. 3.2).

3.1. Data Processing

Template Acquisition. Similar to DeepCap[23], we acquire a single scan for body and clothing (*e.g.* using photogrammetric scanning). Surface registration against a parametric body mesh model [1, 2] is then performed to obtain an estimate for body parts occluded by clothing, *e.g.* the legs under a skirt, which are merged with the visible body parts from the scan to form a complete body mesh. The arms of the body mesh are labeled as inactive when resolving collisions. A separate cloth mesh is created manually from the scan, a task could also be automated [63, 51]. Skeleton parameters and skinning weights, required for posing the meshes, are determined automatically [23]. Two separate embedded graphs [65, 62] for body and clothing are computed by down-sampling the original meshes. These pre-processing steps only need to be done once per character. For more details, we refer to the supplemental document.

Video Capture. We capture the subject to be tracked in a multi-view green screen studio with calibrated and syn-

chronized cameras. The person is asked to perform various tasks, *e.g.* walking, and dancing, to best sample the space of possible poses. Next, we apply OpenPose [11, 10, 60, 75] on all frames and views to obtain multi-view 2D joint predictions. Color keying is used to segment the foreground from the green-screen background and compute distance transformation images D_e from the foreground masks [5].

3.2. Cloth Simulation Layer

Our simulation layer uses the publicly available cloth simulation framework ARCSim [45] as its basis, but we make several adjustments that we describe below.

Material Model and Parameter Selection. ARCSim leverages a data-driven material model defined by a total of 39 parameters. While parameter values for several real-world fabrics are provided, we found that none of them were ideally suited for the materials that we use in our examples, and manually adjusting parameters to obtain better approximation proved very difficult. For this reason, we resorted to a simpler, isotropic material model[69] defined through three parameters: Young’s modulus and Poisson’s ratio for in-plane behavior, and a single bending stiffness coefficient. We determine parameter values through best-guess initialization and a few iterations of simulation-based tuning to better approximate the qualitative behavior observed in the input video sequences. We used the same parameters for all sequences. Although this manual approach is sufficient for our examples, this task could be further automated [63].

Time Integration. During training, the initial state and velocities that are fed into the simulation layer can exhibit large deformations that, when using ARCSim’s default integration method, can lead to instabilities. To improve stability, we resort to an optimization-based formulation of fully implicit Euler [40] combined with adaptive regularization and a back-tracking line search. Finally, we make several code adaptations to enable batch operations for efficient training and integrate the simulation engine in a customized TensorFlow¹ layer. We refer to this layer as the simulation function \mathcal{S} , which takes cloth and body vertices as input and returns the cloth positions for the next time step.

Silhouette Constraint. While our simulation model captures the characteristic behavior of clothing, it is still an approximation, and deviations from the input images must be expected due to external forces such as air drag, viscous damping, and friction that are not modeled. To better track the real-world behavior, we add a multi-view silhouette constraint term. Specifically, this constraint ensures that the vertices $\tilde{V}_{\text{cloth}}^t$ of the simulated cloth geometry matches

¹<https://www.tensorflow.org/>

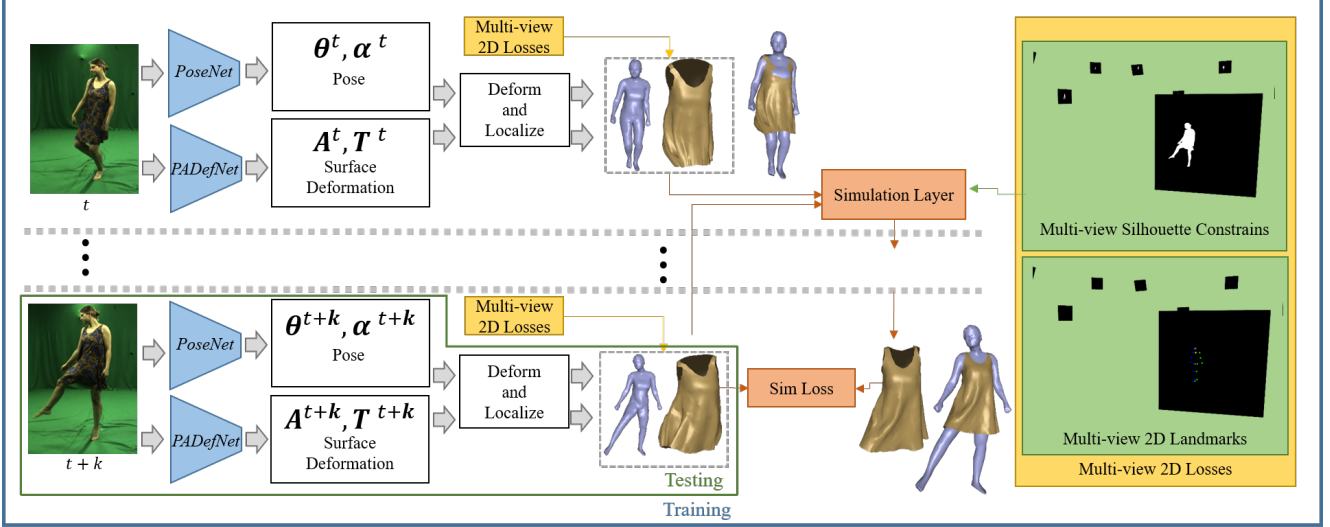


Figure 2. Our method takes a single image as input and two networks, *PoseNet* and *PADefNet*, regress the skeletal pose as well as embedded deformation parameters for the clothing. Combining the outputs of the two networks allows posing and deforming the body and clothing geometry. During training, we use multi-view image losses for *PoseNet* and *PADefNet* is additionally supervised by our proposed simulation loss to encourage physically plausible deformations. To evaluate the simulation loss, we run on-the-fly cloth simulation on small windows of subsequent frames from the training sequence and penalize the difference between regressed deformations and simulation outputs.

the image silhouettes from all camera views for frame t . We construct a 3D ray going through the camera origin and the silhouette pixel p and search for the boundary vertex $\tilde{V}_{\text{cloth},p}^t$ that minimizes the distance to this ray. The closest point on the ray is used as 3D point correspondence for the boundary vertex, enforced via soft constraints

$$E_{\text{cons}} = \sum_p \|\tilde{V}_{\text{cloth},p}^t - V_{\text{ray},p}^t\|^2. \quad (1)$$

3.3. Pose and Deformation Regression

We separate the task of regressing the full surface deformation into predicting pose and surface deformation independently. Therefore, our method consists of two ResNet50 based CNNs [25], *PoseNet* and *PADefNet*, which regress skeleton pose and embedded deformation parameters from a segmented input image, respectively.

3.3.1 Pose Regression and Deformation Model

To pose and deform template vertices as well as sparse body markers, a deformation layer [23] denoted as

$$V_{\text{loc}}, K_{\text{loc}} = f(\theta, \alpha, \mathbf{A}, \mathbf{T}) \quad (2)$$

is used, which is a combination of dual quaternion skinning [32] and embedded deformation [65, 62]. It takes the pose in terms of skeleton joint angles $\theta \in \mathbb{R}^3$ and camera-relative root joint rotation $\alpha \in \mathbb{R}^3$ as well as the embedded graph node rotation $\mathbf{A} \in \mathbb{R}^{K \times 3}$ and translation $\mathbf{T} \in \mathbb{R}^{K \times 3}$ where each row encodes rotations in terms of Euler angles

and translation vectors for each of the K nodes. The output is the posed and deformed vertices V_{loc} and markers K_{loc} in camera and root-relative space. The body pose parameters θ , \mathbf{T} and α are obtained from *PoseNet* [23].

3.3.2 Physics-aware Deformation Regression

To not only pose the template but also account for surface deformation, a dedicated network *PADefNet* predicts the translation vectors \mathbf{T} and rotation angles \mathbf{A} of the embedded graph (EG) from the segmented input image. *PADefNet* is supervised using a combination of both image-based and physics-based metrics, which ensure that deformations match image-based observations while minimizing violations of physical equilibrium conditions. In the remainder of this chapter, we assume *PoseNet* is fixed and provides the posed and deformed vertices \mathbf{V} and markers \mathbf{K} in global space. As \mathbf{V} and \mathbf{K} are a function of the *PADefNet* outputs (\mathbf{T}, \mathbf{A}) , we can then supervise *PADefNet* on \mathbf{V} and \mathbf{K} .

Warm Start. To jump-start our training including the simulation layer, we first pre-train *PADefNet* without running simulation but use a geometric regularizer (ARAP [62]). This adds robustness to the training as geometric regularizers are more stable than simulation and significantly reduces overall training time. Once the network predicts reasonable shapes, we add the simulation loss to supervise the physical deformation. The loss is defined as

$$L_{\text{warm}} = L_{\text{sil}} + L_{\text{lm}} + L_{\text{reg}} + L_{\text{att}} \quad (3)$$

which comprises multi-view losses as well as geometric priors. The individual loss terms are defined as follows.

Multi-view Losses. Our multi-view 2D landmark loss

$$L_{\text{lm}} = \beta_{\text{lm}} \sum_c \sum_m \|\Pi_c(\mathbf{K}_m) - \mathbf{p}_{c,m}\|^2 \quad (4)$$

ensures that the projected landmark matches the 2D detection $\mathbf{p}_{c,m}$ for all views c and landmarks m . Here, Π_c denotes the projection function of the camera c . To densely supervise the surface, we also introduce a silhouette loss

$$L_{\text{sil}} = \beta_{\text{sil}} \sum_c \sum_{b \in \mathcal{B}_c} \rho_{c,b} \|\Pi_c(\mathbf{V}_b) - \mathbf{D}_c\|^2, \quad (5)$$

which ensures that the set of mesh boundary vertices \mathcal{B}_c matches the zero contour line in the distance transformation image \mathbf{D}_c for all views. $\rho_{c,b}$ is a weighting term ensuring that silhouettes are only matched if the normal of the surface aligns with the gradient of the distance transformation [22].

Regularization Loss. To regularize deformations and to avoid drifting of the surface, we employ the as-rigid-as-possible prior [62] to ensure smooth local embedded deformations. We further adopt the rigidity weights formulation [23] to model material-dependent deformation behaviors, *e.g.* the skirt can deform more freely than the skin.

Attachment Loss. Note that our entire mesh \mathbf{V} can be split into body and garment meshes, denoted as $\mathbf{V}_{\text{cloth}}$ and \mathbf{V}_{body} in the remainder of this section. To ensure a coherent movement of these two, an attachment loss

$$L_{\text{att}} = \beta_{\text{att}} \sum_{i \in \mathcal{A}} \|\mathbf{V}_{\text{cloth},i} - \sum_{j=0}^2 \gamma_{i,j} \mathcal{C}(\mathbf{V}_{\text{cloth},i}, \mathbf{V}_{\text{body}})_j\|^2 \quad (6)$$

is included to ensure that the cloth is attached to the body at some anchor positions, *e.g.* the waistband of a skirt has to be attached to the hip of the body mesh. Here, \mathcal{A} are the selected vertices on the garment that act as anchor points, \mathcal{C} is a function that takes the cloth vertex id i and returns the 3 vertices of its closest triangle on the undeformed body mesh, and γ_{ij} are barycentric weights computed from the closest point on this triangle and its three vertices.

Physics-aware Training. While the previous training stage constrains the surface mesh to match the image evidence, it can neither account for the collision of body and clothing nor ensure physically plausible cloth deformations. To this end, we introduce a dedicated simulation-based loss as a better substitution for the ARAP term to explicitly penalize collision behavior and physically implausible deformations. Our final loss is then defined as

$$L = L_{\text{sil}} + L_{\text{lm}} + L_{\text{sim}} + L_{\text{att}}. \quad (7)$$

As our simulation layer \mathcal{S} is directly integrated into a learning framework, we can perform on-the-fly simulation during training. While our method takes a single image as input, the simulation-based loss term is designed to be a multi-frame function to better leverage the sequential training data available. More concretely, this term penalizes the accumulated error on a set of consecutive frames, *i.e.* the mismatching between the per-frame predictions and the on-the-fly simulation results within a frame window (see Eqn. 8). We found that in practice performing simulation over long sequences is extremely challenging when using shapes and poses predicted by a network, since even visually unnoticeable errors, *e.g.* cloth getting trapped in body self-intersections, can lead to catastrophic failures. Hence, we designed our framework specifically to rely only on small simulation windows \mathcal{F} starting at random frames t' to have shorter but successful simulations for training. Additionally, the chosen design is well suited for machine learning, as it allows to access data randomly and in parallel for training. In the following, we refer to a specific frame in this window using the superscript \cdot^t , where $t \in \{t', \dots, t' + \mathcal{F}\}$. Our physics loss then reads

$$L_{\text{sim}} = \beta_{\text{sim}} \sum_i \sum_{t=t'+1}^{t'+\mathcal{F}} \|\mathbf{V}_{\text{cloth},i}^t - \tilde{\mathbf{V}}_{\text{cloth},i}^t\|^2. \quad (8)$$

Here, $\tilde{\mathbf{V}}_{\text{cloth},i}^t$ denote the post-simulation cloth vertex positions, defined as

$$\tilde{\mathbf{V}}_{\text{cloth}}^t = \begin{cases} \mathbf{V}_{\text{cloth}}^t, & t = t' \\ \mathcal{S}(\mathbf{V}_{\text{cloth}}^{t-1}, \mathbf{V}_{\text{cloth}}^t, \mathbf{V}_{\text{body}}^{t-1}, \mathbf{V}_{\text{body}}^t), & t = t' + 1 \\ \mathcal{S}(\tilde{\mathbf{V}}_{\text{cloth}}^{t-2}, \tilde{\mathbf{V}}_{\text{cloth}}^{t-1}, \mathbf{V}_{\text{body}}^{t-1}, \mathbf{V}_{\text{body}}^t), & t > t' + 1 \end{cases}$$

where \mathcal{S} is the aforementioned simulation operation. We initialize the cloth position with *PADefNet* outputs at $t = t'$, where no history is available, and the velocity is initialized using finite difference with cloth position between the succeeding frame. The body vertices positions come from the network predictions and velocities are computed always with finite differences. The loss is then evaluated on the $\mathcal{F} - 1$ frames (excluding the first frame). Even though the first frame in a training sample sequence does not receive this supervision, that frame is supervised by our multi-view supervision, such that in practice all frames are supervised. We opted to not backpropagate gradients through the simulation inputs with respect to the EG parameters as guaranteeing convergence during training would be harder.

4. Results

We evaluate our approach on various outdoor and indoor environment settings with three subject-cloth combinations under a wide range of motions (see Fig. 3). To bridge the

domain gap between training data recorded in the capture studio and in-the-wild testing sequences, *e.g.* different light conditions, we apply a domain adaptation step. *PoseNet* and *PADefNet* are refined for 300 iterations on the testing sequence leveraging the losses introduced before but *using only a single camera*. For in-the-wild captures with varying and dynamic backgrounds, we segment the input images using OSVOS [8]. While the result is almost collision-free thanks to *PADefNet*, minute intersections can remain, which is why we run a final collision resolution step (see also supplemental video). This optional step takes 2s per frame on an Intel i7-9700 CPU. Following DeepCap [23], we apply a temporal Gaussian filter of size 5 frames.

Dataset. Our training dataset contains 3 green screen studio capture sequences with actors performing a large range of motions. For testing, we recorded an additional multi-view green screen sequence to evaluate our reconstruction on reference views and multiple in-the-wild captures using a single camera with a resolution of 1920×1080 for every subject. Apart from a public available sequence *S4* [23], we additionally acquired two training sequences and templates, *F1* and *F2*, with 18 cameras at a resolution of 1285×940 , where each sequence contains around 20,000 frames. We will release the dataset for future research.

Qualitative Results. In Fig. 3, we test our method on various in-the-wild environments while the subjects perform a wide range of motions. Our method does not only provide accurate image overlays and plausible 3D body and cloth geometries but our reconstruction also show physics-aware cloth deformations and plausible body-cloth interactions. *PADefNet* predicts different physically plausible wrinkle patterns related to the character motion as shown in Fig. 4. This is due to our separate modeling of body and cloth geometry and the fact that body-cloth interactions are taken into account by our simulation supervision. We further visualize the underlying body geometry without clothing where also the occluded parts are predicted accurately.

4.1. Comparisons

We compare our approach to the state-of-the-art template-based monocular human performance capture methods [22, 23]. LiveCap [22] optimizes the pose via inverse kinematics to match predicted 2D and 3D joint positions and computes surface deformations via analysis by synthesis. DeepCap [23] uses weak supervision from multi-view images during training to predict pose and embedded deformation parameters from a single segmented image.

4.1.1 Qualitative Comparisons

In Fig. 5, we compare our method with state-of-the-art template-based methods [22, 23]. Unlike our approach,

Methods	Avg	Max
LiveCap[22]	79.85	1140
DeepCap[23]	2.119	29.84
Ours	1.017	7.063

Table 1. Out-of-balance force evaluation. We compare our method to LiveCap [22] and DeepCap [23] with respect to out-of-balance force magnitude. It can be seen that our physics-aware method outperforms state-of-the-art geometry-based methods.

both LiveCap and DeepCap only use geometric priors on the deformations during optimization and training, respectively. Consequently, the resulting cloth deformation contains static wrinkles from the initial template (see top left corners) that persist across all poses. By using simulation supervision and separate modeling of cloth and body geometries, the wrinkles generated by our method are less constrained by the template and, consequently, exhibit more variety and better physical plausibility.

4.1.2 Quantitative Comparisons

We evaluate our results using the green screen testing sequence of *S4* for all metrics below. Note that obtaining accurate ground truth 3D geometry from such a sparse camera system is impossible and thus we resort to image-based and physics-based metrics. For a fair comparison, we use the same cloth-body geometry, obtained through manual clean-up of the input scans, for all approaches.

Out-of-balance Force Evaluation. In Tab. 1, we list the magnitude of the out-of-balance forces, which are defined as the difference between inertial forces and the sum over internal, external, and collision forces. This physical measure indicates to what extent the results deviate from Newton’s second law of motion and vanishes for physically correct motion. The acceleration of the body and the garment for a given frame is determined using a centered difference approximation based on network predictions for three consecutive frames. To reduce the global translation error irrespective of our network predictions, we apply the ground truth global translation for all methods as described by Habermann *et al.* [23]. Our method performs not only better on average compared to other approaches but also significantly reduces the peak value. LiveCap [22] performs significantly worse due to the inherent ambiguity of the single-image setting combined with the inability of geometric priors to capture physical behavior. As a result, the cloth geometry returned by their method exhibits large distortions, in particular in regions occluded from view, resulting in large internal forces. DeepCap [23] leverages neural network models trained with multi-view supervision. Despite substantial improvements compared to LiveCap, our physics-aware method leads to a 50% decrease in error.



Figure 3. Reconstructions obtained with our method for various in-the-wild environments and challenging motion combinations. Our results show good overlay quality throughout, attesting to the pose and clothing estimation accuracy of our method. Furthermore, diverse and physically plausible cloth deformations are observed for a wider range of poses.

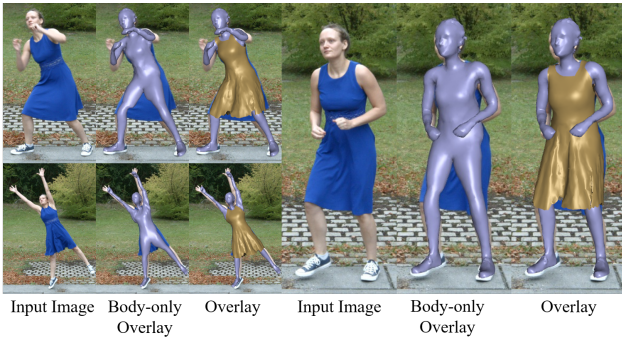


Figure 4. Reconstruction results. Despite minor distortions in occluded regions, our reconstructed body geometry matches the image evidence. Thanks to the separate modeling, our cloth reconstruction is able to reproduce the folding and unfolding behavior of the dress driven by the underlying body motions.

Garment-body Intersection Distance. In Tab. 2, we further compare the garment-body intersection distance to LiveCap [22] and DeepCap [23]. Their simplified one-piece templates sidestep this issue, however, when we evaluate with our more accurate mesh model, *e.g.* separate geometry for clothing and body, collisions severely affect the reconstruction quality. We show that our network predictions significantly reduce cloth penetration.

IoU Percentage. To measure reconstruction quality from different camera views, we compare our results with previous methods using the intersection over union (IoU) met-




Figure 5. Comparison with state-of-the-art template-based monocular methods [22, 23]. Using simulation supervision during training, our method produces physically more realistic results without baked-in wrinkles from the initial template mesh (shown in the top left corners).

Average Penetration Depth (cm)	
Methods	Distance
LiveCap[22]	25.58
DeepCap[23]	23.83
Ours	4.165

Table 2. Penetration Depth. We compute average penetration depths for cloth-body intersections across a 10,000 frame testing sequence. Not taking collisions into account, both LiveCap [22] and DeepCap [23] produce severe penetrations. Our method handles collisions during training, which leads to substantially reduced penetration depths.

Methods	AMVIoU (%)	RVIoU(%)	SVIoU(%)
HMR[30]	65.10	64.66	70.84
LiveCap[22]	59.96	59.02	72.16
DeepCap[23]	82.53	82.22	86.66
Ours	80.83	80.53	84.83

Table 3. IoU percentage comparison. Average multi view (*AMVIoU*), reference view (*RVIoU*) and single view (*SVIoU*) values correspond to IoU evaluation on all views, all views except input view, and input view, respectively. Our reconstruction provides comparable accuracy with state-of-the-art methods while delivering more physically plausible results.

ric (see Tab. 3). The IoU metric indicates the overlapping percentage of the camera projection images of our reconstruction and the foreground segmentation of input images (ground truth). To be consistent with DeepCap [23], the evaluation is performed for every 100th frame of the testing sequence from *S4*, and we apply the same ground truth global translation and a temporal filter. Comparing to body-only reconstruction methods, our method achieves significantly better performance. It should be noted that, compared to the other approaches, the cloth geometry in our method is more constrained due to physics. For example, the strap of a dress cannot detach from the body to match the image silhouettes. Nonetheless, we achieve comparable IoU accuracy while maintaining better physical plausibility.

4.2. Ablation Study

Simulation during Training. Here, we verify that, with simulation supervision in the training process, physically unrealistic cloth deformations and other artifacts resulting from merely image-based supervision can be reduced. In Fig. 6, the strap of the dress remains on the body, and the bottom of the dress does not distort to match the silhouette.




Figure 6. Simulation during training. We evaluate our deformation network with and without simulation loss on a testing sequence. It can be seen that penetrations and deformation artifacts are largely reduced when using simulation.

Simulation during Testing. We compare our learned cloth deformation with traditional cloth simulation (TCS) [45, 69, 40] performing sequentially on a test sequence using *PoseNet* to drive the body mesh. As shown




Figure 7. Simulation during testing. Although running simulation directly on the test sequences can lead to stronger wrinkle patterns, these deformations neither match the wrinkles in the image nor the image silhouettes. In contrast, our predicted cloth closely matches the silhouettes while also showing plausible deformations.

in Fig. 7, our results faithfully match the image evidence due to the silhouette constraint. In contrast, TCS solely provides plausible cloth animation irrespective of the image observation. Moreover, a single failure in the simulation process is fatal for TCS methods, since they cannot easily recover from such failure due to their sequential nature. As such, we refrained from reporting quantitative numbers for TCS methods as they failed already after several frames for the green screen evaluation sequences. In contrast, the presented frame-based approach is robust to failure cases and can recover from bad frames by design.

5. Conclusion

We propose a physics-aware deep learning-based method for monocular human performance capture. With physics-based simulation running on the fly as a network layer, we enforce physics plausibility to compensate for the shorthand of using only multi-view images. We show more visually pleasing results and much-improved physics metrics over state-of-the-art methods.

Limitations & Future Work. The simulation layer is not differentiable in our current implementation. Nevertheless, a fully differentiable physics solver would improve data efficiency and it would open the door to automatic material parameter estimation from video input. Our method is able to faithfully track body pose and cloth deformations for dynamic input motion, but it cannot produce dynamic effects from a single input image—an inherently ill-posed problem. To further improve physical fidelity and reconstruction quality, we would like to extend our method to regress dynamically consistent cloth motion by leveraging deep temporal architectures, which take short videos as input instead of single frames.

6. Acknowledgments

The authors would like to thank the anonymous reviewers for their valuable feedback, and Gereon Fox for the video narration. The authors from MPII were supported by the ERC Consolidator Grant 4DRepLy (770784).

References

- [1] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to re-construct people in clothing from a single RGB camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, jun 2019. 3
- [2] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Combining implicit function learning and parametric models for 3d human reconstruction. In *European Conference on Computer Vision (ECCV)*. Springer, August 2020. 3
- [3] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *IEEE International Conference on Computer Vision (ICCV)*. IEEE, oct 2019. 2
- [4] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [5] Gunilla Borgefors. Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing*, 34(3):344 – 371, 1986. 3
- [6] Matthieu Bray, Pushmeet Kohli, and Philip HS Torr. Posecut: Simultaneous segmentation and 3d pose estimation of humans using dynamic graph-cuts. In *European conference on computer vision*, pages 642–655. Springer, 2006. 1
- [7] Thomas Brox, Bodo Rosenhahn, Juergen Gall, and Daniel Cremers. Combined region and motion-based 3d tracking of rigid and articulated objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(3):402–415, 2010. 1
- [8] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [9] Cedric Cagniart, Edmond Boyer, and Slobodan Ilic. Free-form mesh tracking: a patch-based approach. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1339–1346. IEEE, 2010. 1
- [10] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. In *arXiv preprint arXiv:1812.08008*, 2018. 2, 3
- [11] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017. 2, 3
- [12] Joel Carranza, Christian Theobalt, Marcus A Magnor, and Hans-Peter Seidel. Free-viewpoint video of human actors. *ACM transactions on graphics (TOG)*, 22(3):569–577, 2003. 2
- [13] Edilson De Aguiar, Carsten Stoll, Christian Theobalt, Naveed Ahmed, Hans-Peter Seidel, and Sebastian Thrun. Performance capture from sparse multi-view video. In *ACM Transactions on Graphics (TOG)*, volume 27, page 98. ACM, 2008. 1
- [14] Edilson De Aguiar, Carsten Stoll, Christian Theobalt, Naveed Ahmed, Hans-Peter Seidel, and Sebastian Thrun. Performance capture from sparse multi-view video. In *ACM SIGGRAPH 2008 papers*, pages 1–10. 2008. 2
- [15] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d: Real-time performance capture of challenging scenes. *ACM Transactions on Graphics (TOG)*, 35(4):114, 2016. 1
- [16] Juergen Gall, Carsten Stoll, Edilson De Aguiar, Christian Theobalt, Bodo Rosenhahn, and Hans-Peter Seidel. Motion capture using joint skeleton tracking and surface estimation. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1746–1753. IEEE, 2009. 1
- [17] Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. Drape: Dressing any person. *ACM Transactions on Graphics (TOG)*, 31(4):1–10, 2012. 3
- [18] Peng Guan, Alexander Weiss, Alexandru O Bălan, and Michael J Black. Estimating human shape and pose from a single image. In *ICCV*, pages 1381–1388, 2009. 2
- [19] 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 of the IEEE International Conference on Computer Vision*, pages 8739–8748, 2019. 3
- [20] Kaiwen Guo, Jonathan Taylor, Sean Fanello, Andrea Tagliasacchi, Mingsong Dou, Philip Davidson, Adarsh Kowdle, and Shahram Izadi. Twinfusion: High framerate non-rigid fusion through fast correspondence tracking. 09 2018. 1
- [21] Kaiwen Guo, Feng Xu, Tao Yu, Xiaoyang Liu, Qionghai Dai, and Yebin Liu. Real-time geometry, albedo, and motion reconstruction using a single rgb-d camera. *ACM Transactions on Graphics (TOG)*, 36(3):32, 2017. 1
- [22] Marc Habermann, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Livecap: Real-time human performance capture from monocular video. *ACM Transactions on Graphics (TOG)*, 38(2):1–17, 2019. 1, 2, 5, 6, 7, 8
- [23] Marc Habermann, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, and Christian 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, 4, 5, 6, 7, 8
- [24] Ikhsanul Habibie, Weipeng Xu, Dushyant Mehta, Gerard Pons-Moll, and Christian Theobalt. In the wild human pose estimation using explicit 2d features and intermediate 3d representations, 2019. 2
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [26] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony 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

- [27] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. VolumeDeform: Real-time Volumetric Non-rigid Reconstruction. October 2016. **1**
- [28] Arjun Jain, Thorsten Thormählen, Hans-Peter Seidel, and Christian Theobalt. MovieReshape: Tracking and reshaping of humans in videos. *ACM Transactions on Graphics*, 29(5), 2010. **2**
- [29] Hanbyul Joo, Tomas Simon, and Yaser 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*, 2018. **1**
- [30] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. **1, 2, 8**
- [31] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5614–5623, 2019. **1**
- [32] Ladislav Kavan, Steven Collins, Jiří Žára, and Carol O’Sullivan. Skinning with dual quaternions. In *Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games*, I3D ’07, 2007. **4**
- [33] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Nießner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *ACM Transactions on Graphics (TOG)*, 37(4):163, 2018. **1**
- [34] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction, 2019. **2**
- [35] Adarsh Kowdle, Christoph Rhemann, Sean Fanello, Andrea Tagliasacchi, Jonathan Taylor, Philip Davidson, Mingsong Dou, Kaiwen Guo, Cem Keskin, Sameh Khamis, David Kim, Danhang Tang, Vladimir Tankovich, Julien Valentin, and Shahram Izadi. The need 4 speed in real-time dense visual tracking. In *SIGGRAPH Asia 2018 Technical Papers*, SIGGRAPH Asia ’18, pages 220:1–220:14, New York, NY, USA, 2018. ACM. **1**
- [36] Zorah Lahner, Daniel Cremers, and Tony Tung. Deepwrinkles: Accurate and realistic clothing modeling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 667–684, 2018. **2**
- [37] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, and Peter V.Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *Proc. CVPR*, 2017. **2**
- [38] Yebin Liu, Carsten Stoll, Juergen Gall, Hans-Peter Seidel, and Christian Theobalt. Markerless motion capture of interacting characters using multi-view image segmentation. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1249–1256. IEEE, 2011. **1**
- [39] 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):1–16, 2015. **2**
- [40] Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. Example-based elastic materials. *ACM Trans. Graph.*, 30(4), July 2011. **3, 8**
- [41] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Mohamed Elgharib, Pascal Fua, Hans-Peter Seidel, Helge Rhodin, Gerard Pons-Moll, and Christian Theobalt. XNect: Real-time multi-person 3D motion capture with a single RGB camera. volume 39, 2020. **2**
- [42] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. volume 36, 2017. **2**
- [43] Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, and Christian Theobalt. Generated hands for real-time 3d hand tracking from monocular rgb. In *Proceedings of Computer Vision and Pattern Recognition (CVPR)*, June 2018. **1**
- [44] Armin Mustafa, Hansung Kim, Jean-Yves Guillemaut, and Adrian Hilton. General dynamic scene reconstruction from multiple view video. In *ICCV*, 2015. **1**
- [45] Rahul Narain, Armin Samii, and James F O’Brien. Adaptive anisotropic remeshing for cloth simulation. *ACM transactions on graphics (TOG)*, 31(6):1–10, 2012. **3, 8**
- [46] Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. **1**
- [47] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7365–7375, 2020. **3**
- [48] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image, 2019. **1, 2**
- [49] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3d human pose and shape from a single color image, 2018. **1**
- [50] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael Black. ClothCap: Seamless 4D clothing capture and retargeting. *ACM Transactions on Graphics, (Proc. SIGGRAPH)*, 36(4), 2017. **1**
- [51] 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):1–15, 2017. **2, 3**
- [52] Gerard Pons-Moll, Javier Romero, Naureen Mahmood, and Michael J Black. Dyna: a model of dynamic human shape in motion. *ACM Transactions on Graphics (TOG)*, 34(4):120, 2015. **1**
- [53] Alin-Ionut Popa, Mihai Zanfir, and Cristian Sminchisescu. Deep multitask architecture for integrated 2d and 3d human sensing. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. **2**

- [54] Helge Rhodin, Jörg Spörri, Isinsu Katircioglu, Victor Constantin, Frédéric Meyer, Erich Müller, Mathieu Salzmann, and Pascal Fua. Learning monocular 3d human pose estimation from multi-view images, 2018. [2](#)
- [55] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net: Localization-classification-regression for human pose. In *CVPR 2017-IEEE Conference on Computer Vision & Pattern Recognition*, 2017. [2](#)
- [56] Lorenz Rogge, Felix Klose, Michael Stengel, Martin Eiseermann, and Marcus Magnor. Garment replacement in monocular video sequences. *ACM Transactions on Graphics (TOG)*, 34(1):6, 2014. [2](#)
- [57] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2304–2314, 2019. [2](#)
- [58] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul 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. [2](#)
- [59] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Learning-based animation of clothing for virtual try-on. In *Computer Graphics Forum*, volume 38, pages 355–366. Wiley Online Library, 2019. [3](#)
- [60] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multi-view bootstrapping. In *CVPR*, 2017. [2](#), [3](#)
- [61] Miroslava Slavcheva, Maximilian Baust, Daniel Cremers, and Slobodan Ilic. Killingfusion: Non-rigid 3d reconstruction without correspondences. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 3, page 7, 2017. [1](#)
- [62] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, pages 109–116, 2007. [3](#), [4](#), [5](#)
- [63] Carsten Stoll, Juergen Gall, Edilson De Aguiar, Sebastian Thrun, and Christian Theobalt. Video-based reconstruction of animatable human characters. *ACM Transactions on Graphics (TOG)*, 29(6):1–10, 2010. [2](#), [3](#)
- [64] Z Su, W Wan, T Yu, L Liu, L Fang, W Wang, and Y Liu. Mulaycap: Multi-layer human performance capture using a monocular video camera. *IEEE Transactions on Visualization and Computer Graphics*, 2020. [2](#)
- [65] Robert W Sumner, Johannes Schmid, and Mark Pauly. Embedded deformation for shape manipulation. In *ACM SIGGRAPH 2007 papers*, pages 80–es. 2007. [3](#), [4](#)
- [66] Xiao Sun, Jiaxiang Shang, Shuang Liang, and Yichen Wei. Compositional human pose regression. *ICCV*, 2017. [2](#)
- [67] Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeonwoo Kim, Patrick Pérez, and Christian Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 Hz. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [1](#)
- [68] Ayush Tewari, Michael Zollöfer, Hyeonwoo Kim, Pablo Garrido, Florian Bernard, Patrick Perez, and Theobalt Christian. MoFA: Model-based Deep Convolutional Face Autoencoder for Unsupervised Monocular Reconstruction. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. [1](#)
- [69] Bernhard Thomaszewski, Simon Pabst, and Wolfgang Straßer. Asynchronous cloth simulation. In *Computer Graphics International*, volume 2, page 2, 2008. [3](#), [8](#)
- [70] Denis Tome, Chris Russell, and Lourdes Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. *IEEE Conf. on Computer Vision and Pattern Recognition. Proceedings*, 2017. [2](#)
- [71] Gül Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. BodyNet: Volumetric inference of 3D human body shapes. In *ECCV*, 2018. [2](#)
- [72] Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 20–36, 2018. [2](#)
- [73] Daniel Vlasic, Pieter Peers, Ilya Baran, Paul Debevec, Jovan Popović, Szymon Rusinkiewicz, and Wojciech Matusik. Dynamic shape capture using multi-view photometric stereo. *ACM Transactions on Graphics (TOG)*, 28(5):174, 2009. [1](#)
- [74] Jiayi Wang, Franziska Mueller, Florian Bernard, Suzanne Sorli, Oleksandr Sotnychenko, Neng Qian, Miguel A. Otaduy, Dan Casas, and Christian Theobalt. RGB2Hands: Real-Time Tracking of 3D Hand Interactions from Monocular RGB Video. *ACM Transactions on Graphics (TOG)*, 39(6), 12 2020. [1](#)
- [75] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional Pose Machines. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [2](#), [3](#)
- [76] X. Wei, P. Zhang, and J. Chai. Accurate realtime full-body motion capture using a single depth camera. *ACM TOG (Proc. SIGGRAPH Asia)*, 31(6):188:1–188:12, 2012. [1](#)
- [77] Chenglei Wu, Carsten Stoll, Levi Valgaerts, and Christian Theobalt. On-set Performance Capture of Multiple Actors With A Stereo Camera. In *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2013)*, volume 32, pages 161:1–161:11, November 2013. [1](#)
- [78] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [1](#), [2](#)
- [79] Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. Monoperfcap: Human performance capture from monocular video. *ACM Transactions on Graphics (ToG)*, 37(2):1–15, 2018. [1](#), [2](#)
- [80] Genzhi Ye, Yebin Liu, Nils Hasler, Xiangyang Ji, Qionghai Dai, and Christian Theobalt. Performance capture of interacting characters with handheld kinects. In *ECCV*, volume 7573 LNCS, pages 828–841, 2012. [1](#)

- [81] Mao Ye and Ruigang Yang. Real-time simultaneous pose and shape estimation for articulated objects using a single depth camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2345–2352, 2014. [1](#)
- [82] Tao Yu, Zerong Zheng, Kaiwen Guo, Jianhui Zhao, Qionghai Dai, Hao Li, Gerard Pons-Moll, and Yebin Liu. Doublefusion: Real-time capture of human performances with inner body shapes from a single depth sensor. In *The IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2018. [1](#)
- [83] Tao Yu, Zerong Zheng, Yuan Zhong, Jianhui Zhao, Qionghai Dai, Gerard Pons-Moll, and Yebin 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. [1](#), [2](#)
- [84] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7739–7749, 2019. [2](#)
- [85] Shizhe Zhou, Hongbo Fu, Ligang Liu, Daniel Cohen-Or, and Xiaoguang Han. Parametric reshaping of human bodies in images. *ACM Transactions on Graphics (TOG)*, 29(4):126, 2010. [2](#)
- [86] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: A weakly-supervised approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 398–407, 2017. [2](#)
- [87] Yuxiao Zhou, Marc Habermann, Ikhsanul Habibie, Ayush Tewari, Christian Theobalt, and Feng Xu. Monocular real-time full body capture with inter-part correlations, 2020. [1](#)
- [88] Yuxiao Zhou, Marc Habermann, Weipeng Xu, Ikhsanul Habibie, Christian Theobalt, and Feng Xu. Monocular real-time hand shape and motion capture using multi-modal data. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 0–0, 2020. [1](#)