

A Deeper Look into DeepCap

Marc Habermann, *MPII*, Weipeng Xu, *MPII*, Michael Zollhoefer, *Stanford University*,
Gerard Pons-Moll, *MPII*, and Christian Theobalt, *MPII*

(Invited Paper)

Abstract— Human performance capture is a highly important computer vision problem with many applications in movie production and virtual/augmented reality. Many previous performance capture approaches either required expensive multi-view setups or did not recover dense space-time coherent geometry with frame-to-frame correspondences. We propose a novel deep learning approach for monocular dense human performance capture. Our method is trained in a weakly supervised manner based on multi-view supervision completely removing the need for training data with 3D ground truth annotations. The network architecture is based on two separate networks that disentangle the task into a pose estimation and a non-rigid surface deformation step. Extensive qualitative and quantitative evaluations show that our approach outperforms the state of the art in terms of quality and robustness. This work is an extended version of [1] where we provide more detailed explanations, comparisons and results as well as applications.

Index Terms— Monocular human performance capture, 3D pose estimation, non-rigid surface deformation, human body.

1 INTRODUCTION

CAPTURING the space-time coherent geometry of entire humans including also their everyday clothing, which is also known as human performance capture, is extensively used nowadays for movie production and game development. With these tools, real humans can be ported into a virtual world and fused into the process of creating virtual content. But also other applications such as virtual try-on, telepresence and virtual and augmented reality applications can strongly benefit from dense capture methods. Especially the latter mentioned applications are mainly used by non-expert users, who do not own dense multi-capture setups. Thus, to democratize these technologies the hardware setup has to be as simple and accessible as possible, which means ideally a single color camera is sufficient to capture the entire deforming surface of a human. While there is extensive research regarding monocular skeleton-only tracking, only few works have targeted the capture of the dense deforming geometry of the entire human from a single color stream. Nonetheless, densely capturing entire humans from a single view is still far from being solved while it is indispensable for creating higher fidelity and photo-real characters.

Some works [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14] have focused on capturing entire humans from multi-camera setups that typically also include a green screen. While they achieve a superior quality, the hardware requirement makes it almost impossible to use them in other locations like an outdoor film set and restricts the usage to companies that can afford such an expensive setup.

To overcome the shortcomings of multi-view approaches and with the advent of deep learning, some recent approaches [15], [16], [17], [18], [19], [20], [21] focus on predicting the 3D clothed surface of a human from a single color image. In particular, some works leverage implicit surface representations such as 3D voxel grids [16], [22] or pixel aligned implicit surface functions [15], [23]. While such representations can handle topological surface changes as well




Fig. 1. We present the first learning-based approach for dense monocular human performance capture using weak multi-view supervision that not only predicts the pose but also the space-time coherent non-rigid deformations of the model surface.

as the capture of finer geometric details, they suffer from artifacts like missing limbs as the regressed deformations are not explicitly constrained by a discrete surface mesh. Further, these methods are designed to reconstruct a surface per image. Thus, naively applying those techniques to video streams results in individual geometries per frame which are not coherent over time. To this end, another line of work [17], [19], [24], [25] regresses surface deformation with respect to a naked human body model. Here, coherence over time can be ensured and those works do not suffer from missing limbs since the underlying body model constrains the surface deformation. However, one drawback of these approaches is that they do not capture the motion and surface deformations over time.

The state-of-the-art human performance capture methods MonoPerfCap [26] and LiveCap [27] densely track the deforming surface over time. In contrast to the proposed method, they only predict sparse 2D/3D skeletal keypoints from the images and then perform an expensive optimization based pose and surface fitting. By design, their method

suffers from the monocular setting as the deformations can only be constrained by the single view and their pose results suffer from leaning forward artifacts which originates from their biased keypoint detections. In contrast, we propose the first learning-based approach that jointly regresses the skeletal pose as well as the non-rigid surface deformation within a single inference pass resulting in a higher accuracy in terms of 3D surface and pose tracking as well as an improved robustness. Specifically, two CNN-based models predict the skeletal joint angles and embedded deformation parameters of a differentiable mesh-based character representation from a single image. By using an explicit mesh representation, our method has the advantage that the surface can be tracked over time which is key-essential for texturing and rendering in graphics. Furthermore, the coarse-to-fine modeling of articulation and surface deformation ensures that the output of our method does not suffer from missing limbs, even when those are occluded or when the actor performs out-of-plane motions.

In contrast to previous work [15], [16], [17], [19], our method does not require any form of 3D supervision for training but instead leverages weak supervision in the form of multi-view videos which can be potentially also sparse. To achieve this, we propose differentiable 3D-to-2D modules which allow us to train our deep architectures in an analysis-by-synthesis manner without using 3D ground truth. At the core, our method leverages a person-specific 3D template of the actor as well as a multi-view video showing the actor while he performs a wide range of motions. Then, our dedicated network modules predict the pose and surface deformation parameters that allow to pose and deform the template. This deformed and posed template is then compared against sparse and dense multi-view observations extracted from the multi-view images in a differentiable manner. Importantly, at test time our method only takes a single image as input and predicts the posed and non-rigidly deformed surface of the entire human including also the clothing. In summary, the main technical contributions of our work are:

- A learning-based 3D human performance capture approach that jointly tracks the skeletal pose and the non-rigid surface deformations from monocular images.
- A new differentiable representation of deforming human surfaces which enables training from multi-view video footage directly.

Our new model achieves high quality dense human performance capture results on our new challenging dataset, demonstrating, qualitatively and quantitatively, the advantages of our approach over previous work. We experimentally show that our method produces reconstructions of higher accuracy and 3D stability, in particular in depth, than related work, also under difficult poses.

This work is an extended version of DeepCap [1] where additional explanations, evaluations, comparisons, applications, and limitations are provided. In particular, the character modeling and processing as well as the specific non-trivial training strategies are explained in more detail. Further, we provide additional evaluations on all 4 subjects

of the dataset of [1] and qualitative ablation results. Last, we showcase potential applications and discuss limitations.

2 RELATED WORK

In the following, we focus on related work in the field of dense 3D human performance capture and do not review work on sparse 2D pose estimation.

Capture using Parametric Models. Monocular human performance capture is an ill-posed problem due to its high dimensionality and ambiguity. Low-dimensional parametric models can be employed as shape and deformation prior. First, model-based approaches leverage a set of simple geometric primitives [28], [29], [30], [31]. Recent methods employ detailed statistical models learned from thousands of high-quality 3D scans [13], [32], [33], [34], [35], [36], [37], [38], [39], [40], [41]. Deep learning is widely used to obtain 2D and/or 3D joint detections or 3D vertex positions that can be used to inform model fitting [42], [43], [44], [45], [46]. An alternative is to regress model parameters directly [47], [48], [49]. Beyond body shape and pose, recent models also include facial expressions and hand motion [50], [51], [52], [53] leading to very expressive reconstruction results. Recently, Zhou et al. [54] also predict facial albedo and lighting parameters while even achieving real-time performance. Since parametric body models do not represent garments, variation in clothing cannot be reconstructed, and therefore many methods recover the naked body shape under clothing [55], [56], [57], [58]. The full geometry of the actor can be reconstructed by non-rigidly deforming the base parametric model to better fit the observations [18], [59], [60], [61]. But they can only model tight clothes such as T-shirts and pants, but not loose apparel which has a different topology than the body model, such as skirts. To overcome this problem, ClothCap [14] captures the body and clothing separately, but requires active multi-view setups. Physics based simulations have recently been leveraged to constrain the surface tracking [62], [63] as well as the pose estimation [64], [65], or to learn a model of clothing on top of SMPL (TailorNet [20]). Instead, our method is based on person-specific templates including clothes and employs deep learning to predict clothing deformation based on monocular video directly.

Depth-based Template-free Capture. Most approaches based on parametric models ignore clothing. The other side of the spectrum are prior-free approaches based on one or multiple depth sensors. Capturing general non-rigidly deforming scenes [66], [67], even at real-time frame rates [67], [68], [69], is feasible, but only works reliably for small, controlled, and slow motions. Higher robustness can be achieved by using higher frame rate sensors [70], [71] or multi-view setups [72], [73], [74], [75], [76]. Techniques that are specifically tailored to humans increase robustness [77], [78], [79] by integrating a skeletal motion prior [77] or a parametric model [78], [80]. HybridFusion [81] additionally incorporates a sparse set of inertial measurement units. These fusion-style volumetric capture techniques [82], [83], [84], [85], [86] achieve impressive results, but do not establish a set of dense correspondences between all frames. In addition, such depth-based methods do not directly generalize to our monocular setting, have a high power consumption, and typically do not work well under sunlight.

Monocular Template-free Capture. Quite recently, fueled by the progress in deep learning, many template-free monocular reconstruction approaches have been proposed. Due to their regular structure, many implicit reconstruction techniques [16], [87] make use of uniform voxel grids. DeepHuman [16] combines a coarse scale volumetric reconstruction with a refinement network to add high-frequency details. Multi-view CNNs can map 2D images to 3D volumetric fields enabling reconstruction of a clothed human body at arbitrary resolution [88]. SiCloPe [89] reconstructs a complete textured 3D model, including cloth, from a single image. PIFu [15], its follow-up work [23], and a real-time variant [90] regress an implicit surface representation that locally aligns pixels with the global context of the corresponding 3D object. Similar to our work, ARCH [91] proposes to evaluate the implicit surface function in pose canonical space which makes the learning task easier as the pose and surface deformation can be separated. Unlike voxel-based representations, this implicit per-pixel representation is more memory efficient. These approaches have not been demonstrated to generalize well to strong articulation. Furthermore, implicit approaches do not recover frame-to-frame correspondences which are of paramount importance for downstream applications, e.g., in augmented reality and video editing. In contrast, our method is based on a mesh representation and can explicitly obtain the per-vertex correspondences over time while being slightly less general.

Template-based Capture. An interesting trade-off between being template-free and relying on parametric models are approaches that only employ a template mesh as prior. Historically, template-based human performance capture techniques exploit multi-view geometry to track the motion of a person [92]. Some systems also jointly reconstruct and obtain a foreground segmentation [2], [8], [10], [93]. Given a sufficient number of multi-view images as input, some approaches [4], [9], [94] align a personalized template model to the observations using non-rigid registration. All the aforementioned methods require expensive multi-view setups and are not practical for consumer use. Depth-based techniques enable template tracking from less cameras [72], [95] and reduced motion models [5], [6], [10], [11] increase tracking robustness. Recently, capturing 3D dense human body deformation just with a single RGB camera has been enabled [26] and real-time performance has been achieved [27]. However, their methods rely on expensive optimization leading either to very long per-frame computation times [26] or the need for two graphics cards [27]. Similar to them, our approach also employs a person-specific template mesh. But differently, our method directly learns to predict the skeletal pose and the non-rigid surface deformations. As shown by our experimental results, benefiting from our multi-view based self-supervision, our reconstruction accuracy significantly outperforms the existing methods.

3 METHOD

Given a single RGB video of a moving human in general clothing, our goal is to capture the dense deforming surface of the full body. This is achieved by training a neural network consisting of two components: As illustrated in

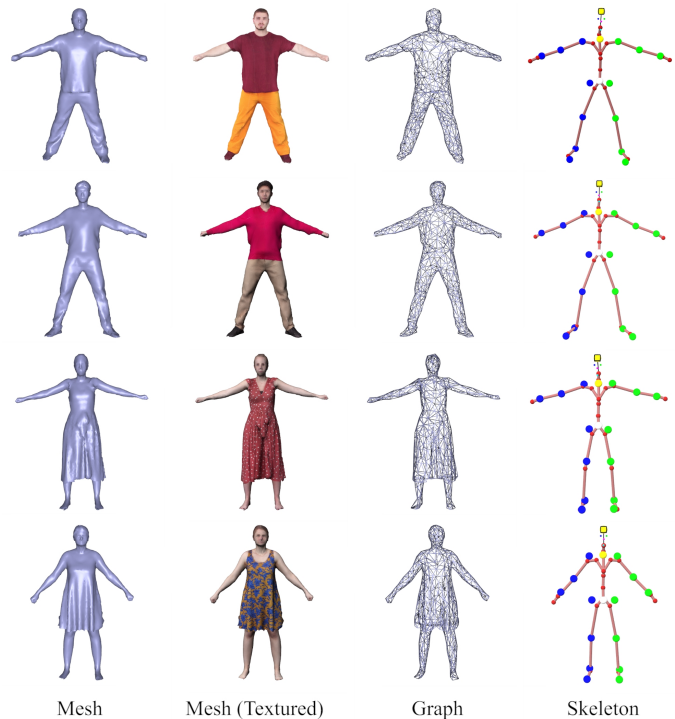


Fig. 2. Character models. Here, we show the character model of $S1$ (top to bottom) of our new dataset. It consists of the textured mesh, the underlying embedded deformation graph as well as the attached skeleton.

Fig. 3, our pose network, *PoseNet*, estimates the skeletal pose of the actor in the form of joint angles from a monocular image (Sec. 3.2). Next, our deformation network, *DefNet*, regresses the non-rigid deformation of the dense surface, which cannot be modeled by the skeletal motion, in the embedded deformation graph representation (Sec. 3.3). To avoid generating dense 3D ground truth annotation, our network is trained in a weakly supervised manner. To this end, we propose a fully differentiable human deformation and rendering model, which allows us to compare the rendering of the human body model to the 2D image evidence and back-propagate the losses. For training, we first capture a video sequence in a calibrated multi-camera green screen studio (Sec. 3.1). Note that our multi-view video is only used during training. More details about training are provided in Sec. 4. At test time we only require a single RGB video to perform dense non-rigid tracking.

3.1 Template and Data Acquisition

Character Model. Our method relies on a person-specific 3D template model. To this end, we first scan the actor with a 3D scanner [96] to obtain the textured mesh (see Fig. 2). Next, we rig a skeleton onto the template which consists of 23 joints and 21 attached landmarks (17 body and 4 face landmarks) and it is parameterized with 27 joint angles $\theta \in \mathbb{R}^{27}$, the camera relative rotation $\alpha \in \mathbb{R}^3$ and translation $\mathbf{t} \in \mathbb{R}^3$. The landmark placement follows the convention of OpenPose [97], [98], [99], [100]. To model the non-rigid surface deformation, we automatically build an embedded deformation graph \mathcal{G} with K nodes. The connections of a

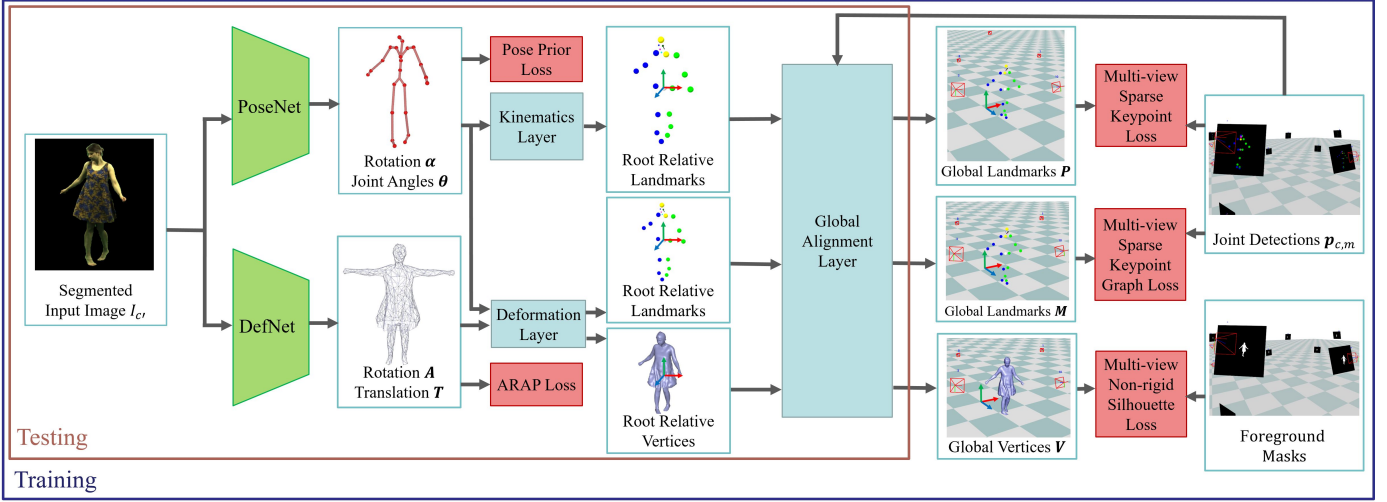


Fig. 3. Overview of our approach. Our method takes a single segmented image as input. First, our pose network, *PoseNet*, is trained to predict the joint angles and the camera relative rotation using sparse multi-view 2D joint detections as weak supervision. Second, the deformation network, *DefNet*, is trained to regress embedded graph rotation and translation parameters to account for non-rigid deformations. To train *DefNet*, multi-view 2D joint detections and silhouettes are used for supervision.

node k to neighboring nodes are denoted as the set $\mathcal{N}_n(k)$. The position of the graph nodes is denoted as $\mathbf{G} \in \mathbb{R}^{K \times 3}$ where \mathbf{G}_k is the position of node k . Finally, the vertex-to-node weights between the graph node k and the template vertex i are defined as $w_{i,k}$ and $\mathcal{N}_{vn}(i)$ denotes the set of nodes that influence vertex i . The nodes are parameterized with Euler angles $\mathbf{A} \in \mathbb{R}^{K \times 3}$ and translations $\mathbf{T} \in \mathbb{R}^{K \times 3}$. Similar to [27], we segment the mesh into different non-rigidity classes resulting in per-vertex rigidity weights s_i . This allows us to model varying deformation behaviors of different surface materials, e.g. skin deforms less than clothing (see Eq. 13).

Training Data. To acquire the training data, we record a multi-view video using C calibrated cameras of the actor doing various actions in a calibrated multi-camera studio with green screen. To provide weak supervision for the training, we first perform 2D pose detection on the sequences using OpenPose [97], [98], [99], [100] and apply temporal filtering. Then, we generate the foreground mask using color keying and compute the corresponding distance transform image $D_{f,c}$ [101], where $f \in [0, F]$ and $c \in [0, C]$ denote the frame index and camera index, respectively. During training, we randomly sample one camera view c' and frame f' for which we crop the recorded image with a bounding box, based on the 2D joint detections. The final training input image $I_{f',c'} \in \mathbb{R}^{256 \times 256 \times 3}$ is obtained by removing the background and augmenting the foreground with random brightness, hue, contrast and saturation changes. For simplicity, we describe the operation on frame f' and omit the subscript f' in following equations.

3.2 Pose Network

In our *PoseNet*, we use ResNet50 [102] pretrained on ImageNet [103] as backbone and modify the last fully connected layer to output a vector containing the joint angles θ and the camera relative root rotation α , given the input image $I_{c'}$. Since generating the ground truth for θ and α is a non-trivial task, we propose weakly supervised training based on fitting the skeleton to multi-view 2D joint detections.

Kinematics Layer. To this end, we introduce a kinematics layer as the differentiable function that takes the joint angles θ and the camera relative rotation α and computes the positions $\mathbf{P}_{c'} \in \mathbb{R}^{M \times 3}$ of the M 3D landmarks attached to the skeleton (17 body joints and 4 face landmarks). Note that $\mathbf{P}_{c'}$ lives in a camera-root-relative coordinate system. In order to project the landmarks to other camera views, we need to transform $\mathbf{P}_{c'}$ to the world coordinate system:

$$\mathbf{P}_m = \mathbf{R}_{c'}^T \mathbf{P}_{c',m} + \mathbf{t}, \quad (1)$$

where $\mathbf{R}_{c'}$ is the rotation matrix of the input camera c' and \mathbf{t} is the global translation of the skeleton.

Global Alignment Layer. To obtain the global translation \mathbf{t} , we propose a global alignment layer that is attached to the kinematics layer. It localizes our skeleton model in the world space, such that the globally rotated landmarks $\mathbf{R}_{c'}^T \mathbf{P}_{c',m}$ project onto the corresponding detections in all camera views. This is done by minimizing the distance between the rotated landmarks $\mathbf{R}_{c'}^T \mathbf{P}_{c',m}$ and the corresponding rays cast from the camera origin \mathbf{o}_c to the 2D joint detections:

$$\sum_c \sum_m \sigma_{c,m} \|(\mathbf{R}_{c'}^T \mathbf{P}_{c',m} + \mathbf{t} - \mathbf{o}_c) \times \mathbf{d}_{c,m}\|^2, \quad (2)$$

where $\mathbf{d}_{c,m}$ is the direction of a ray from camera c to the 2D joint detection $\mathbf{p}_{c,m}$ corresponding to landmark m :

$$\mathbf{d}_{c,m} = \frac{(\mathbf{E}_c^{-1} \tilde{\mathbf{p}}_{c,m})_{xyz} - \mathbf{o}_c}{\|(\mathbf{E}_c^{-1} \tilde{\mathbf{p}}_{c,m})_{xyz} - \mathbf{o}_c\|}. \quad (3)$$

Here, $\mathbf{E}_c \in \mathbb{R}^{4 \times 4}$ is the projection matrix of camera c and $\tilde{\mathbf{p}}_{c,m} = (\mathbf{p}_{c,m}, 1, 1)^T$. Each point-to-line distance is weighted by the joint detection confidence $\sigma_{c,m}$, which is set to zero if below 0.4. The minimization problem of Eq. 2 can be solved in closed form:

$$\mathbf{t} = \mathbf{W}^{-1} \sum_{c,m} \mathbf{D}_{c,m} (\mathbf{R}_{c'}^T \mathbf{P}_{c',m} - \mathbf{o}_c) + \mathbf{o}_c - \mathbf{R}_{c'}^T \mathbf{P}_{c',m}, \quad (4)$$

where

$$\mathbf{W} = \sum_c \sum_m \mathbf{I} - \mathbf{D}_{c,m}. \quad (5)$$

Here, \mathbf{I} is the 3×3 identity matrix and $\mathbf{D}_{c,m} = \mathbf{d}_{c,m} \mathbf{d}_{c,m}^T$. Note that the operation in Eq. 4 is differentiable with respect to the landmark position $\mathbf{P}_{c'}$.

Sparse Keypoint Loss. Our 2D sparse keypoint loss for the *PoseNet* can be expressed as

$$\mathcal{L}_{\text{kp}}(\mathbf{P}) = \sum_c \sum_m \lambda_m \sigma_{c,m} \|\pi_c(\mathbf{P}_m) - \mathbf{p}_{c,m}\|^2, \quad (6)$$

which ensures that each landmark projects onto the corresponding 2D joint detections $\mathbf{p}_{c,m}$ in all camera views. Here, π_c is the projection function of camera c and $\sigma_{c,m}$ is the same as in Eq. 2. λ_m is a hierarchical re-weighting factor that varies during training for better convergence. More precisely, for the first one third of the training iterations per training stage (see Sec. 4) for *PoseNet*, we multiply the keypoint loss with a factor of $\lambda_m = 3$ for all torso markers and with a factor of $\lambda_m = 2$ for elbow and knee markers. For all other markers, we set $\lambda_m = 1$. For the remaining iterations, we set $\lambda_m = 3$ for all markers. This re-weighting allows us to let the model first focus on the global rotation (by weighting torso markers higher than others). We found that this gives better convergence during training and joint angles overshoot less often, especially at the beginning of training.

Pose Prior Loss. To avoid unnatural poses, we impose a pose prior loss on the joint angles

$$\mathcal{L}_{\text{limit}}(\boldsymbol{\theta}) = \sum_{i=1}^{27} \Psi(\theta_i) \quad (7)$$

$$\Psi(x) = \begin{cases} (x - \theta_{\max,i})^2, & \text{if } x > \theta_{\max,i} \\ (\theta_{\min,i} - x)^2, & \text{if } x < \theta_{\min,i} \\ 0, & \text{otherwise} \end{cases}, \quad (8)$$

that encourages that each joint angle θ_i stays in a range $[\theta_{\min,i}, \theta_{\max,i}]$ depending on the anatomic constraints.

3.3 Deformation Network

With the skeletal pose from *PoseNet* alone, the non-rigid deformation of the skin and clothes cannot be fully explained. Therefore, we disentangle the non-rigid deformation and the articulated skeletal motion. *DefNet* takes the input image $I_{c'}$ and regresses the non-rigid deformation parameterized with rotation angles \mathbf{A} and translation vectors \mathbf{T} of the nodes of the embedded deformation graph. *DefNet* uses the same backbone architecture as *PoseNet*, while the last fully connected layer outputs a $6K$ -dimensional vector reshaped to match the dimensions of \mathbf{A} and \mathbf{T} . The weights of *PoseNet* are fixed while training *DefNet*. Again, we do not use direct supervision on \mathbf{A} and \mathbf{T} . Instead, we propose a deformation layer with differentiable rendering and use multi-view silhouette-based weak supervision.

Deformation Layer. The deformation layer takes \mathbf{A} and \mathbf{T} from *DefNet* as input to non-rigidly deform the surface

$$\mathbf{Y}_i = \sum_{k \in \mathcal{N}_{\text{vn}}(i)} w_{i,k} (R(\mathbf{A}_k)(\hat{\mathbf{V}}_i - \mathbf{G}_k) + \mathbf{G}_k + \mathbf{T}_k). \quad (9)$$

Here, $\mathbf{Y}, \hat{\mathbf{V}} \in \mathbb{R}^{N \times 3}$ are the vertex positions of the deformed and undeformed template mesh, respectively. $w_{i,k}$ are vertex-to-node weights, but in contrast to [104] we

compute them based on geodesic distances. $\mathbf{G} \in \mathbb{R}^{K \times 3}$ are the node positions of the undeformed graph, $\mathcal{N}_{\text{vn}}(i)$ is the set of nodes that influence vertex i , and $R(\cdot)$ is a function that converts the Euler angles to rotation matrices. We further apply the skeletal pose on the deformed mesh vertices to obtain the vertex positions in the input camera space

$$\mathbf{V}_{c',i} = \sum_{k \in \mathcal{N}_{\text{vn}}(i)} w_{i,k} (R_{\text{sk},k}(\boldsymbol{\theta}, \boldsymbol{\alpha}) \mathbf{Y}_i + t_{\text{sk},k}(\boldsymbol{\theta}, \boldsymbol{\alpha})), \quad (10)$$

where the node rotation $R_{\text{sk},k}$ and translation $t_{\text{sk},k}$ are derived from the pose parameters using dual quaternion skinning [105]. Eq. 9 and Eq. 10 are differentiable with respect to pose and graph parameters. Thus, our layer can be integrated in the learning framework and gradients can be propagated to *DefNet*. So far, $\mathbf{V}_{c',i}$ is still rotated relative to the camera c' and located around the origin. To bring them to global space, we apply the inverse camera rotation and the global translation, defined in Eq. 4, $\mathbf{V}_i = \mathbf{R}_{c'}^T \mathbf{V}_{c',i} + \mathbf{t}$. **Non-rigid Silhouette Loss.** This loss encourages that the non-rigidly deformed mesh matches the multi-view silhouettes in all camera views. It can be formulated using the distance transform representation [101]

$$\mathcal{L}_{\text{sil}}(\mathbf{V}) = \sum_c \sum_{i \in \mathcal{B}_c} \rho_{c,i} \|D_c(\pi_c(\mathbf{V}_i))\|^2. \quad (11)$$

Here, \mathcal{B}_c is the set of vertices that lie on the boundary when the deformed 3D mesh is projected onto the distance transform image D_c of camera c . Those vertices are computed by rendering a depth map using a custom CUDA-based rasterizer that can be easily integrated into deep learning architectures as a separate layer. The vertices that project onto a depth discontinuity (background vs. foreground) in the depth map are treated as boundary vertices. $\rho_{c,i}$ is a directional weighting [27] that guides the gradient in D_c . The silhouette loss ensures that the boundary vertices project onto the zero-set of the distance transform, *i.e.*, the foreground silhouette.

Sparse Keypoint Graph Loss. Only using the silhouette loss can lead to wrong mesh-to-image assignments, especially for highly articulated motions. To this end, we use a sparse keypoint loss to constrain the mesh deformation, which is similar to the keypoint loss for *PoseNet* in Eq. 6

$$\mathcal{L}_{\text{kpG}}(\mathbf{M}) = \sum_c \sum_m \sigma_{c,m} \|\pi_c(\mathbf{M}_m) - \mathbf{p}_{c,m}\|^2. \quad (12)$$

Differently from Eq. 6, the deformed and posed landmarks \mathbf{M} are derived from the embedded deformation graph. To this end, we can deform and pose the canonical landmark positions by attaching them to its closest graph node g in canonical pose with weight $w_{m,g} = 1.0$. Landmarks can then be deformed according to Eq. 9, 10, resulting in $\mathbf{M}_{c'}$ which is brought to global space via $\mathbf{M}_m = \mathbf{R}_{c'}^T \mathbf{M}_{c',m} + \mathbf{t}$. **As-rigid-as-possible Loss.** To enforce local smoothness of the surface, we impose an as-rigid-as-possible loss [106]

$$\mathcal{L}_{\text{arap}}(\mathbf{A}, \mathbf{T}) = \sum_k \sum_{l \in \mathcal{N}_n(k)} u_{k,l} \|d_{k,l}(\mathbf{A}, \mathbf{T})\|_1, \quad (13)$$

where

$$d_{k,l}(\mathbf{A}, \mathbf{T}) = R(\mathbf{A}_k)(\mathbf{G}_l - \mathbf{G}_k) + \mathbf{T}_k + \mathbf{G}_k - (\mathbf{G}_l + \mathbf{T}_l).$$

$\mathcal{N}_n(k)$ is the set of indices of the neighbors of node k . In contrast to [106], we propose weighting factors $u_{k,l}$ that influence the rigidity of respective parts of the graph. We derive $u_{k,l}$ by averaging all per-vertex rigidity weights s_i [27] for all vertices (see Sec. 3.1), which are connected to node k or l . Thus, the mesh can deform either less or more depending on the surface material. For example, graph nodes that are mostly connected to vertices on a skirt can deform more freely than nodes that are mainly connected to vertices on the skin. Without this loss, the deformations can strongly drift along the visual hull carved by the silhouette images without receiving any penalty leading to strong visual artifacts.

3.4 In-the-wild Domain Adaptation

Since our training set is captured in a green screen studio and our test set is captured in the wild, there is a significant domain gap between them, due to different lighting conditions and camera response functions. To improve the performance of our method on in-the-wild images, we fine-tune our networks on the monocular test images for a small number of iterations using the same 2D keypoint and silhouette losses as before, *but only on a single view*. This drastically improves the performance at test time as shown in the supplemental material.

4 IMPLEMENTATION DETAILS

Both network architectures as well as the GPU-based custom layers are implemented in the Tensorflow framework [107]. We use the Adam optimizer [108] in all our experiments.

Template Acquisition and Rigging. To create the textured mesh (see Fig. 2), we capture the person in a static T-pose with an RGB-based scanner¹ which has 134 RGB cameras resulting in 134 images $\mathcal{I}_{\text{rec}} = \{I_{\text{rec}_1}, \dots, I_{\text{rec}_{134}}\}$. The textured 3D geometry is obtained by leveraging a commercial 3D reconstruction software, called Agisoft Metashape², that takes as input the images \mathcal{I}_{rec} and reconstructs a textured 3D mesh of the person (see Fig. 2). We apply Metashape’s mesh simplification to reduce the number of vertices N and Meshmixer’s³ remeshing to obtain roughly uniform shaped triangular surfaces. Next, we automatically fit the skeleton (see Fig. 2) to the 3D mesh by fitting the SMPL model [35]. To this end, we first optimize the pose by performing a sparse non-rigid ICP where we use the head, hands and feet as feature points since they can be easily detected in a T-pose. Then, we perform a dense non-rigid ICP on vertex level to obtain the final pose and shape parameters. For clothing types that roughly follow the human body shape, e.g., pants and shirt, we propagate the per-vertex skinning weights of the naked SMPL model to the template vertices. For other types of clothing, like skirts and dresses, we leverage Blender’s⁴ automated skinning weight computation.

Embedded Graph Construction. To build the embedded deformation graph \mathcal{G} with K the template mesh is decimated to around 500 vertices (see Fig. 2). The connections of a node

k to neighboring nodes are given by the vertex connections of the decimated mesh. For each vertex of the decimated mesh we search for the closest vertex on the template mesh in terms of Euclidean distance. These points then define the position of the graph nodes. To compute the vertex-to-node weights $w_{i,k}$, we measure the geodesic distance between the graph node k and the template vertex i .

Multi-view video. The number of frames per subject varies between 26000 and 38000 depending on how fast the person performed all the motions. We used C calibrated and synchronized cameras with a resolution of 1024×1024 for capturing where for all subjects we used between 11 and 14 cameras. The original image resolution is too large to transfer all the distance transform images to the GPU during training. Fortunately, most of the image information is anyways redundant since we are only interested in the image region where the person is. Therefore, we crop the distance transform images using the bounding box that contains the segmentation mask with a conservative margin. Finally, we resize it to a resolution of 350×350 without losing important information.

Training Strategy for PoseNet. As we are interested in joint angle regression, one has to note that multiple solutions for the joint angles exist due the fact that every correct solution can be multiplied by 2π leading to the same loss value. To this end, training has to be carefully designed. In general, our strategy first focuses on the torso markers by giving them more weight (see Sec. 3.2). Using this strategy, the global rotation will be roughly correct and joint angles are slowly trained to avoid overshooting of angular values. This is further ensured by our limits term. After several epochs, when the network already learned to fit the poses roughly, we turn off the regularization and let it refine the angles further. More precisely, the training of *PoseNet* proceeds in three stages. First, we train *PoseNet* for $120k$ iterations with a learning rate of 10^{-5} and weight \mathcal{L}_{kp} with 0.01. $\mathcal{L}_{\text{limit}}$ has a weight of 1.0 for the first $40k$ iterations. Between $40k$ and $60k$ iterations we re-weight $\mathcal{L}_{\text{limit}}$ with a factor of 0.1. Finally, we set $\mathcal{L}_{\text{limit}}$ to zero for the remaining training steps. Second, we train *PoseNet* for another $120k$ iterations with a learning rate of 10^{-6} and \mathcal{L}_{kp} is weighted with a factor of 10^{-4} . Third, we train *PoseNet* again $120k$ iterations with a learning rate of 10^{-6} and \mathcal{L}_{kp} is weighted with a factor of 10^{-5} . We always use a batch size of 90.

Training Strategy for DefNet. We train *DefNet* for $120k$ iterations with a batch size of 50. We used a learning rate of 10^{-5} and weight \mathcal{L}_{sil} , \mathcal{L}_{kp} , and $\mathcal{L}_{\text{arap}}$ with $1k$, 0.05 , and $1.5k$, respectively.

Training Strategy for the Domain Adaptation. To fine-tune the network for in-the-wild monocular test sequences, we train the pre-trained *PoseNet* and *DefNet* for 250 iterations, respectively. To this end, we replace the multi-view losses with a single view loss which can be trivially achieved. For *PoseNet*, we disable $\mathcal{L}_{\text{limit}}$ and weight \mathcal{L}_{kp} with 10^{-6} . For *DefNet*, we weight \mathcal{L}_{sil} , \mathcal{L}_{kp} , and $\mathcal{L}_{\text{arap}}$ with $1k$, 0.05 , and $1.5k$ respectively. Further, we use a learning rate of 10^{-6} and use the same batch sizes as before. This fine-tuning in total takes around 5 minutes. All hyperparameters are empirically determined and fixed across different subjects.

1. <https://www.treedys.com/>

2. <http://www.agisoft.com>

3. <http://www.meshmixer.com/>

4. <https://www.blender.org/>

5 RESULTS

All our experiments were performed on a machine with an NVIDIA Tesla V100 GPU. A forward pass of our method takes around 50ms, which breaks down to 25ms for *PoseNet* and 25ms for *DefNet*. During testing, we use the off-the-shelf video segmentation method of [109] to remove the background in the input image. Our method requires OpenPose’s 2D joint detections [97], [98], [99], [100] as input during testing to crop the frames and to obtain the 3D global translation with our global alignment layer. Finally, we temporally smooth the output mesh vertices with a Gaussian kernel of size 5 frames.

Dataset. We evaluate our approach on 4 subjects (*S1* to *S4*) with varying types of apparel. For qualitative evaluation, we recorded 13 in-the-wild sequences in different indoor and outdoor environments shown in Fig. 4. For quantitative evaluation, we captured 4 sequences in a calibrated multi-camera green screen studio (see Fig. 5), for which we computed the ground truth 3D joint locations using the multi-view motion capture software, The Captury [110], and we use a color keying algorithm for ground truth foreground segmentation. All sequences contain a large variety of motions, ranging from simple ones like walking up to more difficult ones like fast dancing or baseball pitching. We will release the dataset for future research.

Qualitative Comparisons. Fig. 4 shows our qualitative results on in-the-wild test sequences with various clothing styles, poses and environments. Our reconstructions not only precisely overlay with the input images, but also look plausible from arbitrary 3D view points. In Fig. 6 and 7, we qualitatively compare our approach to the related human capture and reconstruction methods [15], [16], [27], [47] on our green screen and the in-the-wild sequences, respectively. In terms of the shape representation, our method is most closely related to LiveCap [27] that also uses a person-specific template. Since they non-rigidly fit the template only to the monocular input view, their results do not faithfully depict the deformation in other view points. Further, their pose estimation severely suffers from the monocular ambiguities, whereas our pose results are more robust and accurate (see supplemental video). Comparing to the other three methods [15], [16], [47] that are trained for general subjects, our approach has the following advantages: First, our method recovers the non-rigid deformations of humans in general clothes whereas the parametric model-based approaches [47], [49] only recover naked body shape. Second, our method directly provides surface correspondences over time which is important for AR/VR applications (see supplemental video). In contrast, the results of implicit representation-based methods, PIFu [15] and DeepHuman [16], lack temporal surface correspondences and do not preserve the skeletal structure of the human body, *i.e.*, they often exhibit missing arms and disconnected geometry. Furthermore, DeepHuman [16] only recovers a coarse shape in combination with a normal image of the input view, while our method can recover medium-level detailed geometry that looks plausible from all views. Last but not least, all these existing methods have problems when overlaying their reconstructions on the reference view, even though some of the methods show a very good overlay on the input

view. In contrast, our approach reconstructs accurate 3D geometry, and therefore, our results can precisely overlay on the reference views (also see Fig. 5, 8, 9, and 10).

Skeletal Pose Accuracy. We quantitatively compare our pose results (output of *PoseNet*) to existing pose estimation methods on *S1* to *S4*. To account for different types of apparel, we choose *S1* and *S2* wearing trousers and a T-shirt or a pullover and *S3* and *S4* wearing a long and short dress, respectively. We rescale the bone length for all methods to the ground truth and evaluate the following metrics on the 14 commonly used joints [44] for every 10th frame: 1) We evaluate the root joint position error or global localization error (*GLE*) to measure how good the skeleton is placed in global 3D space. Note that *GLE* can only be evaluated for LiveCap [27] and ours, since other methods only produce up-to-scale depth. 2) To evaluate the accuracy of the pose estimation, we report the 3D percentage of correct keypoints (3DPCK) with a threshold of 150mm of the root aligned poses and the area under the 3DPCK curve (AUC). 3) To factor out the errors in the global rotation, we also report the mean per joint position error (MPJPE) after Procrustes alignment. We compare our approach against the state-of-the-art pose estimation approaches including VNect [44], HMR [47], HMMR [49], and LiveCap [27]. We also compare to a multi-view baseline approach (*MVBL*), where we use our differentiable skeleton model in an optimization framework to solve for the pose per frame using the proposed multi-view losses. We can see from Tab. 3 that our approach outperforms the related monocular methods in all metrics by a large margin and is even close to *MVBL* although our method only takes a single image as input. We further compare to VNect [44] fine-tuned on our training images for *S1*. To this end, we compute the 3D joint position using The Captury [110] to provide ground truth supervision for VNect. On the evaluation sequence for *S1*, the fine-tuned VNect achieved 95.66% 3DPCK, 52.13% AUC and 47.16mm MPJPE. This shows our weakly supervised approach yields comparable or better results than supervised methods in the person-specific setting. However, our approach does not require 3D ground truth annotation that is difficult to obtain, even for only sparse keypoints, let alone the dense surfaces. Further note that even for *S3* we achieve accurate results even though she wears a long dress such that legs are mostly occluded. On *S2*, we found that our results are more accurate than *MVBL* since the classical frame-to-frame optimization can get stuck in local minima, leading to wrong poses.

Surface Reconstruction Accuracy. To evaluate the accuracy of the regressed non-rigid deformations, we compute the intersection over union (IoU) between the ground truth foreground masks and the 2D projection of the estimated shape on *S1* and *S4* for every 100th frame. We evaluate the IoU on *all views*, on *all views except the input view*, and on the *input view* which we refer to as *AMVioU*, *RVioU* and *SVioU*, respectively. To factor out the errors in global localization, we apply the ground truth translation to the reconstructed geometries. For DeepHuman [16] and PIFu [15], we cannot report the *AMVioU* and *RVioU*, since we cannot overlay their results on reference views as discussed before. Further, PIFu [15] by design achieves perfect overlay on the input view, since they regress the depth for each foreground



Fig. 4. Qualitative results. Each row shows results for a different person with varying types of apparel. We visualize input frames and our reconstruction overlaid to the corresponding frame. Note that our results precisely overlay to the input. Further, we show our reconstructions from a virtual 3D viewpoint. Note that they also look plausible in 3D.




Fig. 5. Results on our evaluation sequences where input views (IV) and reference views (RV) are available. Note that our reconstruction also precisely overlays on RV even though they are not used for tracking.

pixel. However, their reconstruction does not reflect the true 3D geometry (see Fig. 6). Therefore, it is meaningless to report their $SVIoU$. Similarly, DeepHuman [16] achieves high $SVIoU$, due to their volumetric representation. But their results are often wrong, when looking from side views. In contrast, our method consistently outperforms all other approaches in terms of $AMVIoU$ and $RVIoU$, which shows the high accuracy of our method in recovering the 3D geometry.

Further, we are again close to the multi-view baseline.

Ablation Study. To evaluate the importance of the number of cameras, the number of training images, and our *DefNet*, we performed an ablation study on *S4* in Tab. 3. 1) In the first group of Tab. 3, we train our networks with supervision using 1 to 14 views. We can see that adding more views consistently improves the quality of the estimated poses and deformations. The most significant improvement is from one to two cameras. This is not surprising, since the single camera settings is inherently ambiguous. In Fig. 8, the importance of the number of cameras is also shown qualitatively. 2) In the second group of Tab. 3 and in Fig. 9, we reduce the training data to 1/2 and 1/4. We can see that the more frames with different poses and deformations are seen during training, the better the reconstruction quality is. This is expected since a larger number of frames may better sample the possible space of poses and deformations. 3) In the third group of Tab. 3, we evaluate the $AMVIoU$ on the template mesh animated with the results of *PoseNet*, which we refer to as *PoseNet-only*. One can see that on average, the $AMVIoU$ is improved by around 4%. Since most non-rigid deformations rather happen locally,




Fig. 6. Qualitative comparison to other methods [15], [16], [27], [47] on our green screen evaluation sequences. Note that our results overlay more accurately to the input view and also look more plausible from a reference view that was not used for tracking. Ground truth global translation is used to match the reference view for the results of [27], [47]. Since PIFu [15] and DeepHuman [16] output meshes with varying topology in a canonical volume without an attached root, it is not possible to apply the ground truth translation and therefore we show the reference view without overlay.

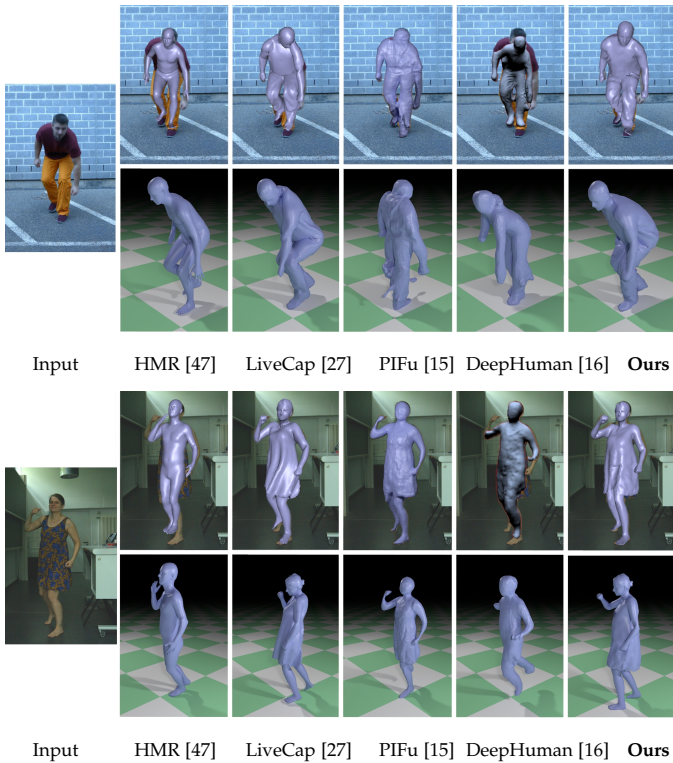


Fig. 7. Comparisons to related work [15], [16], [27], [47] on our in-the-wild sequences showing $S1$ and $S4$. Our approach can recover the deformations of clothing in contrast to [47] and gives more stable and accurate results in 3D compared to [27]. Moreover, note that in contrast to previous work [15], [16], our method regresses space-time coherent geometry, which follows the structure of the human body.

the difference is visually even more significant as shown in Fig. 11. Especially, the skirt is correctly deformed according to the input image whereas the *PoseNet-only* result cannot fit the input due to the limitation of skinning. In Fig. 10, we show the *PoseNet-only* result and our final result on one of our evaluation sequences where a reference view is available. The deformed template also looks plausible from

Method	GLE↓	3DPCK↑	AUC↑	MPJPE↓
VNect [44]	-	66.06	28.02	77.19
HMR [47]	-	82.39	43.61	72.61
HMMR [49]	-	87.48	45.33	72.40
LiveCap [27]	317.01	71.13	37.90	92.84
Ours	91.08	98.43	58.71	49.11
MVBL	76.03	99.17	57.79	45.44

Method	GLE↓	3DPCK↑	AUC↑	MPJPE↓
VNect [44]	-	80.50	39.98	66.96
HMR [47]	-	80.02	39.24	71.87
HMMR [49]	-	82.08	41.00	74.69
LiveCap [27]	142.39	79.17	42.59	69.18
Ours	75.79	94.72	54.61	52.71
MVBL	64.12	89.91	45.58	57.52

Method	GLE↓	3DPCK↑	AUC↑	MPJPE↓
VNect [44]	-	78.03	41.95	88.14
HMR [47]	-	83.37	42.37	79.02
HMMR [49]	-	79.93	36.27	91.62
LiveCap [27]	281.27	66.30	31.44	98.76
Ours	89.54	95.09	54.00	58.77
MVBL	67.82	96.37	54.99	56.08

Method	GLE↓	3DPCK↑	AUC↑	MPJPE↓
VNect [44]	-	82.06	42.73	72.62
HMR [47]	-	86.88	43.91	73.63
HMMR [49]	-	82.80	41.18	77.41
LiveCap [27]	248.67	75.11	37.35	83.48
Ours	96.56	96.74	59.25	45.40
MVBL	75.82	96.20	57.27	45.12

TABLE 1
Skeletal pose accuracy. Note that we are consistently better than other monocular approaches. Moreover, we are even close to the multi-view baseline.




Fig. 8. Ablation for number of *cameras* used during training. The most significant improvement happens when adding one additional camera to the monocular setting. But also adding further cameras consistently improves the result as the yellow circles indicate.

a reference view that was not used for tracking. Importantly, *DefNet* can correctly regress deformations that are along the camera viewing direction of the input camera (see reference view in second column) and surface parts that are even occluded (see reference view in fourth column). This implies that our weak multi-view supervision during training let the network learn the entire 3D surface deformation of the human body. For more results, we refer to the supplemental video. 4) finally, in Fig. 12, we visually demonstrate the impact of our domain adaptation step. It becomes obvious

Method	AMVIoU \uparrow	RVIoU \uparrow	SVIoU \uparrow
HMR [47]	62.25	61.7	68.85
HMMR [49]	65.98	65.58	70.77
LiveCap [27]	56.02	54.21	77.75
DeepHuman [16]	-	-	91.57
Ours	87.2	87.03	89.26
MVBL	91.74	91.72	92.02

Method	AMVIoU \uparrow	RVIoU \uparrow	SVIoU \uparrow
HMR [47]	59.79	59.1	66.78
HMMR [49]	62.64	62.03	68.77
LiveCap [27]	60.52	58.82	77.75
DeepHuman [16]	-	-	91.57
Ours	83.73	83.49	89.26
MVBL	89.62	89.67	92.02

Method	AMVIoU \uparrow	RVIoU \uparrow	SVIoU \uparrow
HMR [47]	59.05	58.73	63.12
HMMR [49]	61.73	61.32	67.14
LiveCap [27]	61.55	60.47	75.6
DeepHuman [16]	-	-	79.66
Ours	85.75	85.55	88.27
MVBL	90.31	90.21	91.53

Method	AMVIoU \uparrow	RVIoU \uparrow	SVIoU \uparrow
HMR [47]	65.1	64.66	70.84
HMMR [49]	63.79	63.29	70.23
LiveCap [27]	59.96	59.02	72.16
DeepHuman [16]	-	-	84.15
Ours	82.53	82.22	86.66
MVBL	88.14	88.03	89.66

TABLE 2

Surface deformation accuracy. Note that we again outperform all other monocular methods and are close to the multi-view baseline. Further note, that for [16] an evaluation of the multi-view IoU is not possible since their output is always in local image space that cannot be brought to global space.




Fig. 9. Ablation for number of *frames* used during training. The more frames we used during training the better the result becomes as the network can better sample the possible pose and deformation space.

that the refinement drastically improves the pose as well as the non-rigid deformations so that the input can be matched at much higher accuracy. Further, we do not require any additional input for the refinement as our losses can be directly adapted to the monocular setting.

Applications. Our method enables driving 3D characters just from a monocular RGB video (see Fig. 4). As the only device that is needed at test time is a single color camera, our method can be easily used in daily life scenarios once




Fig. 10. Our result from the input view and a reference view that was not used for tracking. Note that our *DefNet* can even regress deformations along the camera viewing axis of the input camera (second column) and it can correctly deform surface parts that are occluded (fourth column).

Method	3DPCK \uparrow	AMVIoU \uparrow
1 camera view	62.11	65.11
2 camera views	93.52	78.44
3 camera views	94.70	79.75
7 camera views	95.95	81.73
6500 frames	85.19	73.41
13000 frames	92.25	78.97
PoseNet-only	96.74	78.51
Ours(14 views, 26000 frames)	96.74	82.53

TABLE 3

Ablation study. We evaluate the number of cameras and the number of frames used during training in terms of the *3DPCK* and *AMVIoU* metrics. Adding more cameras and frames consistently improves the quality of reconstruction. Further, *DefNet* improves the *AMVIoU* compared to pure pose estimation.

the multi-view video and the template are acquired and the training of the model was performed. Further, as we also account for non-rigid surface deformations, our method also enhances the realism of the virtual characters. Our approach also allows augmenting a video as shown in Fig. 13. Since we track the entire 3D geometry, the augmented texture is also aware of occlusions in contrast to pure image-based augmentation techniques.

6 CONCLUSION

Limitations. Conceptually, both representations, pose and the non-rigid deformations, are decoupled. Nevertheless, since the predicted poses during training are not perfect, our *DefNet* also deforms the graph to account for wrong poses to a certain degree. In our supplemental video, we also tested our method on subjects that were not used for training but which wear the same clothing as the training subject. Although, our method still performs reasonable, the overall accuracy drops as the subjects appearance was never observed during training. Further, our method can fail for extreme poses, e.g. a hand stand, that were not observed during training.

We have presented a learning-based approach for monocular dense human performance capture using only weak multi-view supervision. In contrast to existing methods, our approach directly regresses poses and surface deformations from neural networks, produces temporal sur-




Fig. 11. *PoseNet + DefNet vs. PoseNet-only*. *DefNet* can deform the template to accurately match the input, especially for loose clothing. In addition, *DefNet* also corrects slight errors in the pose and typical skinning artifacts.




Fig. 12. Impact of the in-the-wild domain adaption step. Note that after the network refinement, both, the pose as well as the deformations better match the input.

face correspondences, preserves the skeletal structure of the human body, and can handle loose clothes. Our qualitative and quantitative results in different scenarios show that our method produces more accurate 3D reconstruction of pose and non-rigid deformation than existing methods. In the future, we plan to incorporate hands and the face to our mesh representation to enable joint tracking of body, facial expressions and hand gestures. We are also interested in physically more correct multi-layered representations to model the garments even more realistically.

ACKNOWLEDGMENT

This work was funded by the ERC Consolidator Grant 4DRepLy (770784) and the Deutsche Forschungsgemeinschaft (Project Nr. 409792180, Emmy Noether Programme, project: Real Virtual Humans).

REFERENCES

- [1] M. Habermann, W. Xu, M. Zollhoefer, G. Pons-Moll, and C. Theobalt, "Deepcap: Monocular human performance capture using weak supervision," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2020.
- [2] M. Bray, P. Kohli, and P. H. Torr, "Posecut: Simultaneous segmentation and 3d pose estimation of humans using dynamic graphcuts," in *European conference on computer vision*. Springer, 2006, pp. 642–655.




Fig. 13. Video augmentation. Our method can be used to augment a video with textures like the logo on the T-shirt. Since we track the underlying 3D geometry our method also accounts for occlusions of the augmented texture.

- [3] T. Brox, B. Rosenhahn, D. Cremers, and H.-P. Seidel, "High accuracy optical flow serves 3-d pose tracking: exploiting contour and flow based constraints," in *European Conference on Computer Vision*. Springer, 2006, pp. 98–111.
- [4] E. De Aguiar, C. Stoll, C. Theobalt, N. Ahmed, H.-P. Seidel, and S. Thrun, "Performance capture from sparse multi-view video," in *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3. ACM, 2008, p. 98.
- [5] D. Vlasic, I. Baran, W. Matusik, and J. Popović, "Articulated mesh animation from multi-view silhouettes," in *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3. ACM, 2008, p. 97.
- [6] J. Gall, C. Stoll, E. De Aguiar, C. Theobalt, B. Rosenhahn, and H.-P. Seidel, "Motion capture using joint skeleton tracking and surface estimation," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 2009, pp. 1746–1753.
- [7] D. Vlasic, P. Peers, I. Baran, P. Debevec, J. Popović, S. Rusinkiewicz, and W. Matusik, "Dynamic shape capture using multi-view photometric stereo," *ACM Transactions on Graphics (TOG)*, vol. 28, no. 5, p. 174, 2009.
- [8] T. Brox, B. Rosenhahn, J. Gall, and D. Cremers, "Combined region and motion-based 3d tracking of rigid and articulated objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 3, pp. 402–415, 2010.
- [9] C. Cagniard, E. Boyer, and S. Ilic, "Free-form mesh tracking: a patch-based approach," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2010, pp. 1339–1346.
- [10] Y. Liu, C. Stoll, J. Gall, H.-P. Seidel, and C. Theobalt, "Markerless motion capture of interacting characters using multi-view image segmentation," in *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. IEEE, 2011, pp. 1249–1256.
- [11] C. Wu, C. Stoll, L. Valgaerts, and C. Theobalt, "On-set Performance Capture of Multiple Actors With A Stereo Camera," in *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2013)*, vol. 32, no. 6, November 2013, pp. 161:1–161:11. [Online]. Available: <http://doi.acm.org/10.1145/2508363.2508418>
- [12] A. Mustafa, H. Kim, J.-Y. Guillemaut, and A. Hilton, "General dynamic scene reconstruction from multiple view video," in *ICCV*, 2015.
- [13] G. Pons-Moll, J. Romero, N. Mahmood, and M. J. Black, "Dyna: a model of dynamic human shape in motion," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, p. 120, 2015.
- [14] G. Pons-Moll, S. Pujades, S. Hu, and M. Black, "ClothCap: Seamless 4D clothing capture and retargeting," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 36, no. 4, 2017. [Online]. Available: <http://dx.doi.org/10.1145/3072959.3073711>
- [15] S. Saito, Z. Huang, R. Natsume, S. Morishima, A. Kanazawa, and H. Li, "Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization," *CoRR*, vol. abs/1905.05172, 2019. [Online]. Available: <http://arxiv.org/abs/1905.05172>
- [16] Z. Zheng, T. Yu, Y. Wei, Q. Dai, and Y. Liu, "Deephuman: 3d human reconstruction from a single image," *CoRR*, vol. abs/1903.06473, 2019. [Online]. Available: <http://arxiv.org/abs/1903.06473>
- [17] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2019, pp. 1175–1186.
- [18] T. Alldieck, M. Magnor, W. Xu, C. Theobalt, and G. Pons-Moll, "Video based reconstruction of 3d people models," in *IEEE*

- Conference on Computer Vision and Pattern Recognition*, 2018, CVPR Spotlight Paper.
- [19] 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.
- [20] C. Patel, Z. Liao, and G. Pons-Moll, "Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2020.
- [21] Q. Ma, J. Yang, A. Ranjan, S. Pujades, G. Pons-Moll, S. Tang, and M. Black, "Learning to dress 3d people in generative clothing," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2020.
- [22] 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*, 2019, pp. 2232–2241.
- [23] 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 Conference on Computer Vision and Pattern Recognition*, June 2020.
- [24] T. Alldieck, G. Pons-Moll, C. Theobalt, and M. Magnor, "Tex2shape: Detailed full human body geometry from a single image," in *IEEE International Conference on Computer Vision (ICCV)*. IEEE, oct 2019.
- [25] A. Pumarola, J. Sanchez-Riera, G. P. T. Choi, A. Sanfeliu, and F. Moreno-Noguer, "3dpeople: Modeling the geometry of dressed humans," in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [26] W. Xu, A. Chatterjee, M. Zollhöfer, H. Rhodin, D. Mehta, H.-P. Seidel, and C. Theobalt, "Monoperfcap: Human performance capture from monocular video," *ACM Trans. Graph.*, vol. 37, no. 2, pp. 27:1–27:15, May 2018. [Online]. Available: <http://doi.acm.org/10.1145/3181973>
- [27] M. Habermann, W. Xu, M. Zollhoefer, G. Pons-Moll, and C. Theobalt, "Livecap: Real-time human performance capture from monocular video," *ACM Trans. Graph.*, 2019.
- [28] R. Plänkers and P. Fua, "Tracking and modeling people in video sequences," *Computer Vision and Image Understanding*, vol. 81, no. 3, pp. 285–302, 2001.
- [29] C. Sminchisescu and B. Triggs, "Estimating articulated human motion with covariance scaled sampling," *The International Journal of Robotics Research*, vol. 22, no. 6, pp. 371–391, 2003.
- [30] L. Sigal, S. Bhatia, S. Roth, M. J. Black, and M. Isard, "Tracking loose-limbed people," in *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, vol. 1. IEEE, 2004, pp. I–421.
- [31] D. Metaxas and D. Terzopoulos, "Shape and nonrigid motion estimation through physics-based synthesis," *IEEE Trans. PAMI*, vol. 15, no. 6, pp. 580–591, 1993.
- [32] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis, "SCAPE: Shape Completion and Animation of People," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 408–416, 2005.
- [33] N. Hasler, H. Ackermann, B. Rosenhahn, T. Thormählen, and H.-P. Seidel, "Multilinear pose and body shape estimation of dressed subjects from image sets," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2010, pp. 1823–1830.
- [34] S. I. Park and J. K. Hodgins, "Data-driven modeling of skin and muscle deformation," in *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3. ACM, 2008, p. 96.
- [35] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "SMPL: A skinned multi-person linear model," *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, vol. 34, no. 6, pp. 248:1–248:16, Oct. 2015.
- [36] P. Kadlecik, A.-E. Ichim, T. Liu, J. Krivanek, and L. Kavan, "Reconstructing personalized anatomical models for physics-based body animation," *ACM Trans. Graph.*, vol. 35, no. 6, 2016.
- [37] M. Kim, G. Pons-Moll, S. Pujades, S. Bang, J. Kim, M. Black, and S.-H. Lee, "Data-driven physics for human soft tissue animation," *ACM Transactions on Graphics, (Proc. SIGGRAPH)*, vol. 36, no. 4, 2017. [Online]. Available: <http://dx.doi.org/10.1145/3072959.3073685>
- [38] A. Weiss, D. Hirshberg, and M. J. Black, "Home 3d body scans from noisy image and range data," in *Proc. ICCV*. IEEE, 2011, pp. 1951–1958.
- [39] T. Helten, M. Muller, H.-P. Seidel, and C. Theobalt, "Real-time body tracking with one depth camera and inertial sensors," in *The IEEE International Conference on Computer Vision (ICCV)*, December 2013.
- [40] Q. Zhang, B. Fu, M. Ye, and R. Yang, "Quality dynamic human body modeling using a single low-cost depth camera," in *2014 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2014, pp. 676–683.
- [41] F. Bogo, M. J. Black, M. Loper, and J. Romero, "Detailed full-body reconstructions of moving people from monocular RGB-D sequences," in *International Conference on Computer Vision (ICCV)*, Dec. 2015, pp. 2300–2308.
- [42] Y. Huang, F. Bogo, C. Lassner, A. Kanazawa, P. V. Gehler, J. Romero, I. Akhter, and M. J. Black, "Towards accurate markerless human shape and pose estimation over time," in *International Conference on 3D Vision (3DV)*, 2017.
- [43] C. Lassner, J. Romero, M. Kiefel, F. Bogo, M. J. Black, and P. V. Gehler, "Unite the people: Closing the loop between 3d and 2d human representations," in *Proc. CVPR*, 2017.
- [44] D. Mehta, S. Sridhar, O. Sotnychenko, H. Rhodin, M. Shafiee, H.-P. Seidel, W. Xu, D. Casas, and C. Theobalt, "Vnect: Real-time 3d human pose estimation with a single rgb camera," in *ACM Transactions on Graphics*, vol. 36, no. 4, July 2017. [Online]. Available: <http://gvv.mpi-inf.mpg.de/projects/VNect/>
- [45] 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 (ECCV)*, 2016.
- [46] N. Kolotouros, G. Pavlakos, and K. Daniilidis, "Convolutional mesh regression for single-image human shape reconstruction," in *CVPR*, 2019.
- [47] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik, "End-to-end recovery of human shape and pose," *CoRR*, vol. abs/1712.06584, 2017. [Online]. Available: <http://arxiv.org/abs/1712.06584>
- [48] G. Pavlakos, L. Zhu, X. Zhou, and K. Daniilidis, "Learning to estimate 3D human pose and shape from a single color image," in *CVPR*, 2018.
- [49] A. Kanazawa, J. Y. Zhang, P. Felsen, and J. Malik, "Learning 3d human dynamics from video," in *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [50] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. A. Osman, D. Tzionas, and M. J. Black, "Expressive body capture: 3d hands, face, and body from a single image," in *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019. [Online]. Available: <http://smpl-x.is.tue.mpg.de>
- [51] D. Xiang, H. Joo, and Y. Sheikh, "Monocular total capture: Posing face, body, and hands in the wild," *CoRR*, vol. abs/1812.01598, 2018. [Online]. Available: <http://arxiv.org/abs/1812.01598>
- [52] H. Joo, T. Simon, and Y. Sheikh, "Total capture: A 3d deformation model for tracking faces, hands, and bodies," *CoRR*, vol. abs/1801.01615, 2018.
- [53] J. Romero, D. Tzionas, and M. J. Black, "Embodied hands: Modeling and capturing hands and bodies together," *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, vol. 36, no. 6, pp. 245:1–245:17, Nov. 2017. [Online]. Available: <http://doi.acm.org/10.1145/3130800.3130883>
- [54] Y. Zhou, M. Habermann, I. Habibie, A. Tewari, C. Theobalt, and F. Xu, "Monocular real-time full body capture with inter-part correlations," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [55] A. O. Balan, L. Sigal, M. J. Black, J. E. Davis, and H. W. Haussecker, "Detailed human shape and pose from images," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007, pp. 1–8.
- [56] A. O. Balan and M. J. Black, "The naked truth: Estimating body shape under clothing," in *European Conference on Computer Vision*. Springer, 2008, pp. 15–29.
- [57] C. Zhang, S. Pujades, M. Black, and G. Pons-Moll, "Detailed, accurate, human shape estimation from clothed 3D scan sequences," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, spotlight.
- [58] J. Yang, J.-S. Franco, F. Héttroy-Wheeler, and S. Wuhrer, "Estimation of Human Body Shape in Motion with Wide Clothing," in *European Conference on Computer Vision 2016*, Amsterdam, Netherlands, Oct. 2016.
- [59] H. Rhodin, N. Robertini, D. Casas, C. Richardt, H.-P. Seidel, and C. Theobalt, "General automatic human shape and motion cap-

- ture using volumetric contour cues,” in *ECCV*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 509–526.
- [60] T. Alldieck, M. Magnor, W. Xu, C. Theobalt, and G. Pons-Moll, “Detailed human avatars from monocular video,” in *International Conference on 3D Vision*, Sep 2018, pp. 98–109.
- [61] D. Xiang, F. Prada, C. Wu, and J. Hodgins, “Monoclothcap: Towards temporally coherent clothing capture from monocular rgb video,” in *2020 International Conference on 3D Vision (3DV)*, 2020, pp. 322–332.
- [62] Y. Tao, Z. Zheng, Y. Zhong, J. Zhao, D. Quionhai, G. Pons-Moll, and Y. Liu, “Simulcap: Single-view human performance capture with cloth simulation,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, jun 2019.
- [63] Y. Li, M. Habermann, B. Thomaszewski, S. Coros, T. Beeler, and C. Theobalt, “Deep physics-aware inference of cloth deformation for monocular human performance capture,” 2020.
- [64] S. Shimada, V. Golyanik, W. Xu, and C. Theobalt, “Physcap: Physically plausible monocular 3d motion capture in real time,” *ACM Transactions on Graphics*, vol. 39, no. 6, dec 2020.
- [65] S. Shimada, V. Golyanik, W. Xu, P. Pérez, and C. Theobalt, “Neural monocular 3d human motion capture,” *ACM Transactions on Graphics*, vol. 40, no. 4, aug 2021.
- [66] M. Slavcheva, M. Baust, D. Cremers, and S. Ilic, “Killingfusion: Non-rigid 3d reconstruction without correspondences,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 3, no. 4, 2017, p. 7.
- [67] K. Guo, F. Xu, T. Yu, X. Liu, Q. Dai, and Y. Liu, “Real-time geometry, albedo, and motion reconstruction using a single rgb-d camera,” *ACM Transactions on Graphics (TOG)*, vol. 36, no. 3, p. 32, 2017.
- [68] R. A. Newcombe, D. Fox, and S. 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.
- [69] M. Innmann, M. Zollhöfer, M. Nießner, C. Theobalt, and M. Stamminger, “VolumeDeform: Real-time Volumetric Non-rigid Reconstruction,” *Proceedings of European Conference on Computer Vision (ECCV)*, October 2016.
- [70] K. Guo, J. Taylor, S. Fanello, A. Tagliasacchi, M. Dou, P. Davidson, A. Kowdle, and S. Izadi, “Twinfusion: High framerate non-rigid fusion through fast correspondence tracking,” in *3DV*, 09 2018.
- [71] A. Kowdle, C. Rhemann, S. Fanello, A. Tagliasacchi, J. Taylor, P. Davidson, M. Dou, K. Guo, C. Keskin, S. Khamis, D. Kim, D. Tang, V. Tankovich, J. Valentin, and S. Izadi, “The need 4 speed in real-time dense visual tracking,” in *SIGGRAPH Asia 2018 Technical Papers*, ser. SIGGRAPH Asia ’18. New York, NY, USA: ACM, 2018, pp. 220:1–220:14. [Online]. Available: <http://doi.acm.org/10.1145/3272127.3275062>
- [72] G. Ye, Y. Liu, N. Hasler, X. Ji, Q. Dai, and C. Theobalt, “Performance capture of interacting characters with handheld kinects,” in *ECCV*, vol. 7573 LNCS, no. PART 2, 2012, pp. 828–841.
- [73] M. Dou, S. Khamis, Y. Degtyarev, P. Davidson, S. R. Fanello, A. Kowdle, S. O. Escolano, C. Rhemann, D. Kim, J. Taylor *et al.*, “Fusion4d: Real-time performance capture of challenging scenes,” *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, p. 114, 2016.
- [74] S. Orts-Escolano, C. Rhemann, S. Fanello, W. Chang, A. Kowdle, Y. Degtyarev, D. Kim, P. L. Davidson, S. Khamis, M. Dou *et al.*, “Holoportation: Virtual 3d teleportation in real-time,” in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 2016, pp. 741–754.
- [75] M. Dou, P. Davidson, S. R. Fanello, S. Khamis, A. Kowdle, C. Rhemann, V. Tankovich, and S. Izadi, “Motion2fusion: Real-time volumetric performance capture,” *ACM Trans. Graph.*, vol. 36, no. 6, pp. 246:1–246:16, Nov. 2017.
- [76] P. Zhang, K. Siu, J. Zhang, C. K. Liu, and J. Chai, “Leveraging depth cameras and wearable pressure sensors for full-body kinematics and dynamics capture,” *ACM Transactions on Graphics (TOG)*, vol. 33, no. 6, p. 14, 2014.
- [77] T. Yu, K. Guo, F. Xu, Y. Dong, Z. Su, J. Zhao, J. Li, Q. Dai, and Y. Liu, “Bodyfusion: Real-time capture of human motion and surface geometry using a single depth camera,” in *The IEEE International Conference on Computer Vision (ICCV)*. ACM, October 2017.
- [78] 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 *The IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2018.
- [79] M. Ye and R. 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*, 2014, pp. 2345–2352.
- [80] X. Wei, P. Zhang, and J. Chai, “Accurate realtime full-body motion capture using a single depth camera,” *ACM TOG (Proc. SIGGRAPH Asia)*, vol. 31, no. 6, pp. 188:1–188:12, 2012.
- [81] Z. Zheng, T. Yu, H. Li, K. Guo, Q. Dai, L. Fang, and Y. Liu, “HybridFusion: Real-Time Performance Capture Using a Single Depth Sensor and Sparse IMUs,” in *Proceedings of the 15th European Conference on Computer Vision*. Munich, Germany: Computer Vision Foundation, Sep. 2018.
- [82] C.-H. Huang, B. Allain, J.-S. Franco, N. Navab, S. Ilic, and E. Boyer, “Volumetric 3d tracking by detection,” in *Proc. CVPR*, 2016.
- [83] B. Allain, J.-S. Franco, and E. Boyer, “An Efficient Volumetric Framework for Shape Tracking,” in *CVPR 2015 - IEEE International Conference on Computer Vision and Pattern Recognition*. Boston, United States: IEEE, Jun. 2015, pp. 268–276. [Online]. Available: <https://hal.inria.fr/hal-01141207>
- [84] V. Leroy, J.-S. Franco, and E. Boyer, “Multi-View Dynamic Shape Refinement Using Local Temporal Integration,” in *IEEE, International Conference on Computer Vision 2017*, Venice, Italy, Oct. 2017. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-01567758>
- [85] A. Collet, M. Chuang, P. Sweeney, D. Gillett, D. Evseev, D. Calabrese, H. Hoppe, A. Kirk, and S. Sullivan, “High-quality streamable free-viewpoint video,” *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, p. 69, 2015.
- [86] F. Prada, M. Kazhdan, M. Chuang, A. Collet, and H. Hoppe, “Spatiotemporal atlas parameterization for evolving meshes,” *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, p. 58, 2017.
- [87] G. Varol, D. Ceylan, B. Russell, J. Yang, E. Yumer, I. Laptev, and C. Schmid, “BodyNet: Volumetric inference of 3D human body shapes,” in *ECCV*, 2018.
- [88] 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 15th European Conference on Computer Vision*. Computer Vision Foundation, Sep. 2018.
- [89] R. Natsume, S. Saito, Z. Huang, W. Chen, C. Ma, H. Li, and S. Morishima, “Siclope: Silhouette-based clothed people,” *arXiv preprint arXiv:1901.00049*, 2018.
- [90] R. Li, Y. Xiu, S. Saito, Z. Huang, K. Olszewski, and H. Li, “Monocular real-time volumetric performance capture,” in *European Conference on Computer Vision*. Springer, 2020, pp. 49–67.
- [91] Z. Huang, Y. Xu, C. Lassner, H. Li, and T. Tung, “Arch: Animatable reconstruction of clothed humans,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020.
- [92] J. Starck and A. Hilton, “Surface capture for performance-based animation,” *IEEE Computer Graphics and Applications*, vol. 27, no. 3, pp. 21–31, 2007.
- [93] C. Wu, K. Varanasi, and C. Theobalt, “Full body performance capture under uncontrolled and varying illumination: A shading-based approach,” in *ECCV*, 2012, pp. 757–770.
- [94] J. Carranza, C. Theobalt, M. A. Magnor, and H.-P. Seidel, “Free-viewpoint video of human actors,” *ACM Trans. Graph.*, vol. 22, no. 3, Jul. 2003.
- [95] M. Zollhöfer, M. Nießner, S. Izadi, C. Rhemann, C. Zach, M. Fisher, C. Wu, A. Fitzgibbon, C. Loop, C. Theobalt, and M. Stamminger, “Real-time non-rigid reconstruction using an rgb-d camera,” *ACM Transactions on Graphics (TOG)*, vol. 33, no. 4, 2014.
- [96] “Treedys,” <https://www.treedys.com/>.
- [97] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, “Realtime multi-person 2d pose estimation using part affinity fields,” in *CVPR*, 2017.
- [98] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh, “OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields,” in *arXiv preprint arXiv:1812.08008*, 2018.
- [99] T. Simon, H. Joo, I. Matthews, and Y. Sheikh, “Hand keypoint detection in single images using multiview bootstrapping,” in *CVPR*, 2017.

- [100] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh, "Convolutional Pose Machines," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [101] G. Borgefors, "Distance transformations in digital images," *Computer Vision, Graphics, and Image Processing*, vol. 34, no. 3, pp. 344 – 371, 1986. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0734189X86800470>
- [102] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [103] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*, 2009.
- [104] R. W. Sumner, J. Schmid, and M. Pauly, "Embedded deformation for shape manipulation," *ACM Trans. Graph.*, vol. 26, no. 3, Jul. 2007.
- [105] L. Kavan, S. Collins, J. Žára, and C. O'Sullivan, "Skinning with dual quaternions," in *Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games*, ser. I3D '07, 2007.
- [106] O. Sorkine and M. Alexa, "As-rigid-as-possible surface modeling," in *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, ser. SGP '07. Eurographics Association, 2007.
- [107] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: <http://tensorflow.org/>
- [108] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014.
- [109] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool, "One-shot video object segmentation," in *Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [110] "The Captury," <http://www.thecaptury.com/>.



Marc Habermann works as a PhD student in the Graphics, Vision and Video group at the Max Planck Institute for Informatics. Within his thesis, he explores the modeling and tracking of non-rigid deformations of surfaces, e.g. capturing the performance of humans in their everyday clothing. In his previous works, he showed that this is possible at real-time frame rates and that the 3D performance can be further improved using deep learning techniques. He received the Guenter-Hotz-Medal for the best Master graduates in Computer Science at Saarland University in 2017 and his work, DeepCap, received the CVPR Best Student Paper Honorable Mention in 2020.



Weipeng Xu is a research scientist at Facebook Reality Labs in Pittsburgh. He was a postdoctoral researcher at the Graphic, Vision & Video group of Max Planck Institute for Informatics in Saarbruecken, Germany. He received B.E. and Ph.D. degrees from Beijing Institute of Technology in 2009 and 2016, respectively. He studied as a long-term visiting student at NICTA and Australian National University from 2013 to 2015. His research interests include virtual human character, human pose estimation and machine learning for vision/graphics.



Michael Zollhoefer is a Visiting Assistant Professor at Stanford University. His stay at Stanford is funded by a postdoctoral fellowship of the Max Planck Center for Visual Computing and Communication (MPC-VCC), which he received for his work in the fields of computer vision, computer graphics, and machine learning. Before, Michael was a Postdoctoral Researcher in the Graphics, Vision & Video group at the Max Planck Institute for Informatics in Saarbruecken, Germany. He received his PhD in 2014 from the University of Erlangen-Nuremberg for his work on real-time static and dynamic scene reconstruction. His research is focused on teaching computers to reconstruct and analyze our world at frame rate based on visual input. To this end, he develops key technology to invert the image formation models of computer graphics based on data-parallel optimization and state-of-the-art deep learning techniques.



Gerard Pons-Moll is the head of the Emmy Noether independent research group "Real Virtual Humans", senior researcher at the Max Planck for Informatics (MPII) in Saarbrücken, Germany, and Junior Faculty at Saarland Informatics Campus. His research lies at the intersection of computer vision, computer graphics and machine learning – with special focus on analyzing people in videos, and creating virtual human models by "looking" at real ones. His research has produced some of the most advanced statistical human body models of pose, shape, soft-tissue and clothing (which are currently used for a number of applications in industry and research), as well as algorithms to track and reconstruct 3D people models from images, video, depth, and IMUs. His work has received several awards including the prestigious Emmy Noether Grant (2018), a Google Faculty Research Award (2019), a Facebook Reality Labs Faculty Award (2018), and recently the German Pattern Recognition Award (2019), which is given annually by the German Pattern Recognition Society to one outstanding researcher in the fields of Computer Vision and Machine Learning. In 2020 he received a Snap-Research gift. His work got Best Papers Awards BMVC'13, Eurographics'17, 3DV'18 and CVPR'20 and has been published at the top venues and journals including CVPR, ICCV, Siggraph, Eurographics, 3DV, IJCV and PAMI. He served as Area Chair for ECCV'18, 3DV'19, SCA'18'19, FG'20, ECCV'20. He will serve as Area Chair for CVPR'21 and 3DV'20.



Christian Theobalt is a Professor of Computer Science and the head of the research group "Graphics, Vision, & Video" at the Max-Planck-Institute for Informatics, Saarbruecken, Germany. He is also a professor at Saarland University. His research lies on the boundary between Computer Vision and Computer Graphics. For instance, he works on 4D scene reconstruction, marker-less motion and performance capture, machine learning for graphics and vision, and new sensors for 3D acquisition. Christian received several awards, for instance the Otto Hahn Medal of the Max-Planck Society (2007), the EUROGRAPHICS Young Researcher Award (2009), the German Pattern Recognition Award (2012), an ERC Starting Grant (2013), an ERC Consolidator Grant (2017), and the Eurographics Outstanding Technical Contributions Award (2020). In 2015, he was elected one of Germany's top 40 innovators under 40 by the magazine Capital. He is a co-founder of theCaptury (www.thecaptury.com).