

ChallenCap: Monocular 3D Capture of Challenging Human Performances using Multi-Modal References

Yannan He¹ Anqi Pang¹ Xin Chen¹ Han Liang¹ Minye Wu¹ Yuexin Ma^{1,2} Lan Xu^{1,2}

¹ShanghaiTech University

²Shanghai Engineering Research Center of Intelligent Vision and Imaging

Abstract

Capturing challenging human motions is critical for numerous applications, but it suffers from complex motion patterns and severe self-occlusion under the monocular setting. In this paper, we propose ChallenCap — a template-based approach to capture challenging 3D human motions using a single RGB camera in a novel learning-and-optimization framework, with the aid of multi-modal references. We propose a hybrid motion inference stage with a generation network, which utilizes a temporal encoder-decoder to extract the motion details from the pair-wise sparse-view reference, as well as a motion discriminator to utilize the unpaired marker-based references to extract specific challenging motion characteristics in a data-driven manner. We further adopt a robust motion optimization stage to increase the tracking accuracy, by jointly utilizing the learned motion details from the supervised multi-modal references as well as the reliable motion hints from the input image reference. Extensive experiments on our new challenging motion dataset demonstrate the effectiveness and robustness of our approach to capture challenging human motions.

1. Introduction

The past ten years have witnessed a rapid development of markerless human motion capture [14, 24, 60, 68], which benefits various applications such as immersive VR/AR experience, sports analysis and interactive entertainment.

Multi-view solutions [60, 40, 29, 12, 30, 63] achieve high-fidelity results but rely on expensive studio setup which are difficult to be deployed for daily usage. Recent learning-based techniques enables robust human attribute prediction from monocular RGB video [31, 35, 2, 80, 55]. The state-of-the-art monocular human motion capture approaches [22, 75, 74] leverage learnable pose detections [9, 44] and template fitting to achieve space-time coherent results. However, these approaches fail to capture the specific challenging motions such as yoga or rolling on the floor,

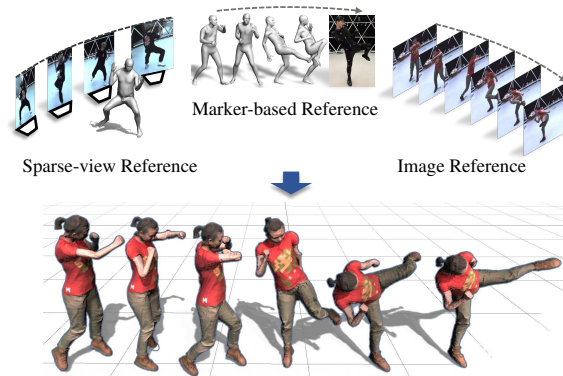


Figure 1. Our ChallenCap approach achieves robust 3D capture of challenging human motions from a single RGB video, with the aid of multi-modal references.

which suffer from extreme poses, complex motion patterns and severe self-occlusion under the monocular setting.

Capturing such challenging human motions is essential for many applications such as training and evaluation for gymnastics, sports and dancing. Currently, optical marker-based solutions like Vicon [66] are widely adopted to capture such challenging professional motions. However, directly utilizing such marker-based reference into markerless capture is inapplicable since the actor needs to re-perform the challenging motion which is temporally unsynchronized to the marker-based capture. Some data-driven human pose estimation approaches [32, 35] utilize the unpaired reference in an adversarial manner, but they only extract general motion prior from existing motion capture datasets [28, 43], which fails to recover the characteristics of specific challenging motion. The recent work [23] inspires to utilize the markerless multi-view reference in a data-driven manner to provide more robust 3D prior for monocular capture. However, this method is weakly supervised on the input images instead of the motion itself, leading to dedicated performer training. Moreover, researchers pay less attention to combine various references from both marker-based systems and sparse multi-view systems for monocular challenging motion capture.

In this paper, we tackle the above challenges and present *ChallenCap* – a template-based monocular 3D capture approach for challenging human motions from a single RGB video, which outperforms existing state-of-the-art approaches significantly (See Fig. 1 for an overview). Our novel pipeline proves the effectiveness of embracing multi-modal references from both temporally unsynchronized marker-based system and light-weight markerless multi-view system in a data-driven manner, which enables robust human motion capture under challenging scenarios with extreme poses and complex motion patterns, whilst still maintaining a monocular setup.

More specifically, we introduce a novel learning-and-optimization framework, which consists of a hybrid motion inference stage and a robust motion optimization stage. Our hybrid motion inference utilizes both the marker-based reference which encodes the accurate spatial motion characteristics but sacrifices the temporal consistency, as well as the sparse multi-view image reference which provides pairwise 3D motion priors but fails to capture extreme poses. To this end, we first obtain the initial noisy skeletal motion map from the input monocular video. Then, a novel generation network, HybridNet, is proposed to boost the initial motion map, which utilizes a temporal encoder-decoder to extract local and global motion details from the sparse-view reference, as well as a motion discriminator to utilize the unpaired marker-based reference. Besides the data-driven 3D motion characteristics from the previous stage, the input RGB video also encodes reliable motion hints for those non-extreme poses, especially for the non-occluded regions. Thus, a robust motion optimization is further proposed to refine the skeletal motions and improve the tracking accuracy and overlay performance, which jointly utilizes the learned 3D prior from the supervised multi-modal references as well as the reliable 2D and silhouette information from the input image reference. To summarize, our main contributions include:

- We propose a monocular 3D capture approach for challenging human motions, which utilizes multi-modal reference in a novel learning-and-optimization framework, achieving significant superiority to state-of-the-arts.
- We propose a novel hybrid motion inference module to learn the challenging motion characteristics from the supervised references modalities, as well as a robust motion optimization module for accurate tracking.
- We introduce and make available a new challenging human motion dataset with both unsynchronized marker-based and light-weight multi-image references, covering 60 kinds of challenging motions and 20 performers with 120k corresponding images.

2. Related Work

As an alternative to the widely used marker-based solutions [66, 70, 67], markerless motion capture [8, 15, 65] technologies alleviate the need for body-worn markers and have been widely investigated. In the following, we focus on the field of marker-less 3D human motion capture.

Parametric Model-based Capture. Many general human parametric models [3, 41, 49, 47] learned from thousands of high-quality 3D scans have been proposed in the last decades, which factorize human deformation into pose and shape components. Deep learning is widely used to obtain skeletal pose and human shape prior through model fitting [25, 37, 7, 36] or directly regressing the model parameters from the input [32, 33, 35, 78]. Besides, various approaches [81, 64, 52, 2, 80] propose to predict detailed human geometry by utilizing parametric human model as a basic estimation. Beyond human shape and pose, recent approaches further include facial and hand models [69, 30, 49, 11] for expressive reconstruction or leverage garment and clothes modeling on top of parametric human model [51, 5, 48, 42]. But these methods are still limited to the parametric model and cannot provide space-time coherent results for loose clothes. Instead, our method is based on person-specific templates and focuses on capturing space-time coherent challenging human motions using multi-modal references.

Free-form Volumetric Capture. Free-form volumetric capture approaches with real-time performance have been proposed by combining the volumetric fusion [13] and the nonrigid tracking [62, 38, 82, 20] using depth sensors. The high-end solutions [17, 16, 73, 19] rely on multi-view studios which are difficult to be deployed. The most handy monocular approaches for general non-rigid scenes [46, 26, 21, 57, 58, 59, 71] can only capture small, controlled, and slow motions. Researchers further utilize parametric model [76, 77, 73, 61] or extra body-worn sensors [79] into the fusion pipeline to increase the tracking robustness. However, these fusion approaches rely on depth cameras which are not as cheap and ubiquitous as color cameras. Recently, the learning-based techniques enable free-form human reconstruction from monocular RGB input with various representations, such as volume [80], silhouette [45] or implicit representation [54, 55, 39, 10, 4]. However, such data-driven approaches do not recover temporal-coherent reconstruction, especially under the challenging motion setting. In contrast, our template-based approach can explicitly obtain the per-vertex correspondences over time even for challenging motions.

Template-based Capture. A good compromising settlement between from-form capture and parametric Modal-based Capture is to utilize a specific human template mesh as prior. Early solutions [18, 60, 40, 53, 50, 56, 72] require

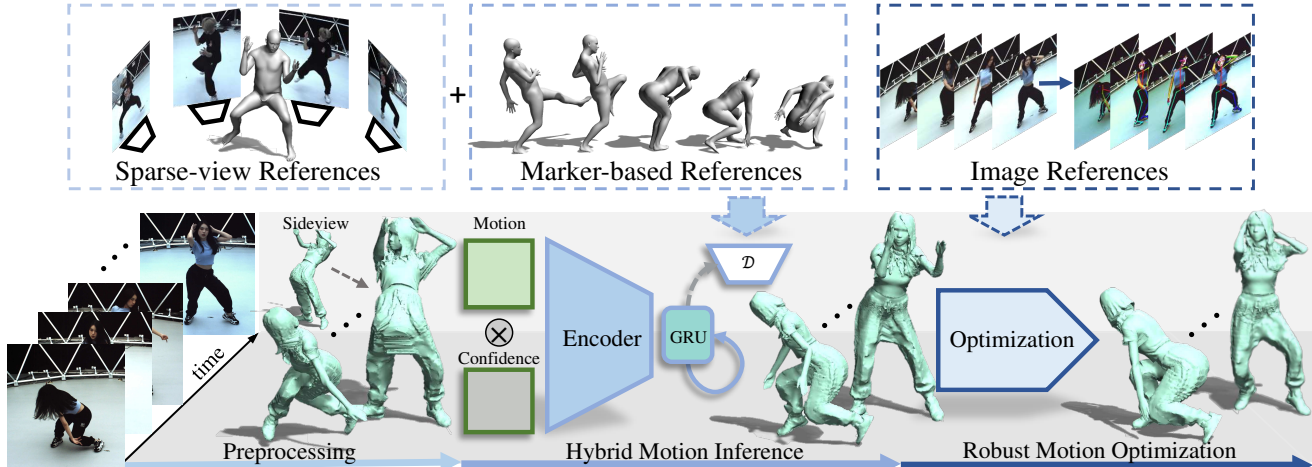


Figure 2. The pipeline of ChallenCap with multi-modal references. Assuming the video input from monocular camera, our approach consists of a hybrid motion inference stage (Sec.4.1) and a robust motion optimization stage (Sec.4.2) to capture 3D challenging motions. \mathcal{D} represents the discriminator.

multi-view capture to produce high quality skeletal and surface motions but synchronizing and calibrating multi-camera systems is still cumbersome. Recent work only relies on a single-view setup [75, 22, 74] achieve space-time coherent capture and even achieves real-time performance [22]. However, these approaches fail to capture the challenging motions such as yoga or rolling on the floor, which suffer from extreme poses and severe self-occlusion under the monocular setting. The recent work [23] utilizes weekly supervision on multi-view images directly so as to improve the 3D tracking accuracy during test time. However, their training strategy leads to dedicated per-performer training. Similarly, our approach also employs a person-specific template mesh. Differently, we adopt a specific learning-and-optimization framework for challenging human motion capture. Our learning module is supervised on the motion itself instead of the input images of specific performers for improving the generation performance to various performers and challenging motions.

3. Overview

Our goal is to capture challenging 3D human motions from a single RGB video, which suffers from extreme poses, complex motion patterns and severe self-occlusion. Fig. 2 provides an overview of ChallenCap, which relies on a template mesh of the actor and makes full usage of multi-modal references in a learning-and-optimization framework. Our method consists of a hybrid motion inference module to learn the challenging motion characteristics from the supervised references modalities, and a robust motion optimization module to further extract the reliable motion hints in the input images for more accurate tracking.

Template and Motion Representation. We use a 3D

body scanner to generate the template mesh of the actor and rig it by fitting the Skinned Multi-Person Linear Model (SMPL)[41] to the template mesh and transferring the SMPL skinning weights to our scanned mesh. The kinematic skeleton is parameterized as $\mathbf{S} = [\boldsymbol{\theta}, \mathbf{R}, \mathbf{t}]$, including the joint angles $\boldsymbol{\theta} \in \mathbb{R}^{30}$ of the N_J joints, the global rotation $\mathbf{R} \in \mathbb{R}^3$ and translation $\mathbf{t} \in \mathbb{R}^3$ of the root. Furthermore, let \mathbf{Q} denotes the quaternions representation of the skeleton. Thus, we can formulate $\mathbf{S} = \mathcal{M}(\mathbf{Q}, \mathbf{t})$ where \mathcal{M} denotes the motion transformation between various representations.

Hybrid Motion Inference. Our novel motion inference scheme extracts the challenging motion characteristics from the supervised marker-based and sparse multi-view references in a data-driven manner. We first obtain the initial noisy skeletal motion map from the monocular video. Then, a novel generation network, HybridNet, is adopted to boost the initial motion map, which consists of a temporal encoder-decoder to extract local and global motion details from the sparse-view references, as well as a motion discriminator to utilize the motion characteristics from the unpaired marker-based references. To train our HybridNet, a new dataset with rich references modalities and various challenging motions is introduced (Sec. 4.1).

Robust Motion Optimization. Besides the data-driven 3D motion characteristics from the previous stage, the input RGB video also encodes reliable motion hints for those non-extreme poses, especially for the non-occluded regions. Thus, a robust motion optimization is introduced to refine the skeletal motions so as to increase the tracking accuracy and overlay performance, which jointly utilizes the learned 3D prior from the supervised multi-modal references as well as the reliable 2D and silhouette information from the input image references (Sec. 4.2).

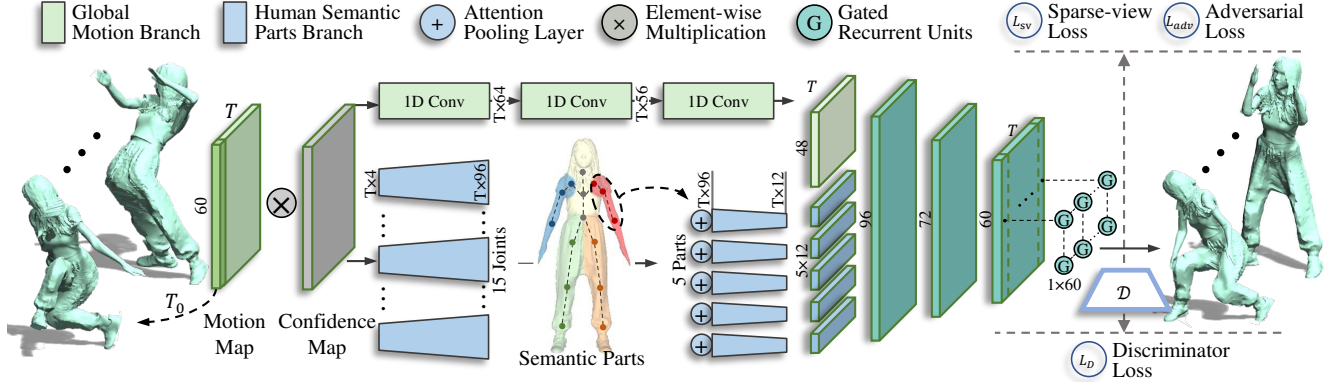


Figure 3. Illustration of our hybrid motion network, HybridNet, which encodes the global and local temporal motion information with the losses on both the generator and the discriminator. Note that the attention pooling operation is performed by applying element-wise addition (blue branches) on the features of the adjacent body joints. The features after attention pooling are concatenated together with the global feature as input to the fully connected layers.

4. Approach

4.1. Hybrid Motion Inference

Preprocessing. Given an input monocular image sequence $I_t, t \in [1, T]$ of length T and a well-scanned template model, we first adopt the off-the-shelf template-based motion capture approach [75] to obtain the initial skeletal motion \mathbf{S}_t and transform it into quaternions format, denoted as \mathbf{Q}_t^{init} . More specifically, we only adopt the 2D term from [75] using OpenPose [9] to obtain 2D joint detections. Please refer to [75] for more optimization details. Note that such initial skeletal motions suffer from severe motion ambiguity since no 3D prior is utilized, as illustrated in the pre-processing stage in Fig. 2. After the initial optimization, we concatenate the \mathbf{Q}_t^{init} and the detection confidence from OpenPose [9] for all the T frames into a motion map $\mathcal{Q} \in \mathbb{R}^{T \times 4N_J}$ as well as a confidence map $\mathcal{C} \in \mathbb{R}^{T \times N_J}$.

HybridNet Training. Based on the initial noisy motion map \mathcal{Q} and confidence map \mathcal{C} , we propose a novel generation network, HybridNet, to boost the initial capture results for challenging human motions.

As illustrated in Fig. 3, our HybridNet learns the challenging motion characteristics by the supervision from multi-modal references. To avoid tedious 3D pose annotation, we utilize the supervision from the optimized motions using sparse multi-view image reference. Even though such sparse-view reference still cannot recover all the extreme poses, it provides rich pair-wise overall 3D motion prior. To further extract the fine challenging motion details, we utilize adversarial supervision from the marker-based reference since it only provides accurate but temporally unsynchronized motion characteristics due to re-performing. To this end, we utilize the well-known Generative Adversarial Network (GAN) structure in our HybridNet with the generative network \mathcal{G} and the discriminator network \mathcal{D} .

Our generative module consists of a global-local motion encoder with a hierarchical attention pooling block and a GRU-based decoder to extract motion details from the sparse-view references, which takes the concatenated \mathcal{Q} and \mathcal{C} as input. In our encoder, we design two branches to encode the global and local skeletal motion features independently. Note that we mainly use 1D convolution layers during the encoding process to extract corresponding temporal information. We apply three layers of 1D convolution for the global branch while splitting the input motion map into N_J local quaternions for the local branch inspired by [6]. Differently, we utilize a hierarchy attention pooling layer to connect the feature of adjacent joints and compute the latent codes from five local body regions, including the four limbs and the torso. We concatenate the global and local feature of two branches as the final latent code, and decode them to the original quaternions domain with three linear layers in our GRU-based decoder (see Fig. 3 for detail). Here, the loss of our generator \mathcal{G} is formulated as:

$$\mathcal{L}_{\mathcal{G}} = \mathcal{L}_{sv} + \mathcal{L}_{adv}, \quad (1)$$

where \mathcal{L}_{sv} is the sparse-view loss and \mathcal{L}_{adv} is the adversarial loss. Our sparse-view loss is formulated as:

$$\mathcal{L}_{sv} = \sum_{t=1}^T \left\| \hat{\mathbf{Q}}_t - \mathbf{Q}_t^{sv} \right\|_2^2 + \lambda_{quat} \sum_{t=1}^T \sum_{i=1}^{N_J} (\|\hat{\mathbf{Q}}_t^{(i)}\| - 1)^2. \quad (2)$$

Here, the first term is the L_2 loss between the regressed output motion $\hat{\mathbf{Q}}_t$ and the 3D motion prior \mathbf{Q}_t^{sv} from sparse-views reference. Note that we obtain \mathbf{Q}_t^{sv} from the reference sparse multi-view images by directly extending the same optimization process to the multi-view setting. The second regular term forces the network output quaternions to represent a rotation. Besides, N_J denotes the number of joints which is 15 in our case while λ_{quat} is set to be 1×10^{-5} .

Our motion discriminator \mathcal{D} further utilizes the motion characteristics from the unpaired marker-based references, which maps the motion map $\hat{\mathcal{Q}}$ corrected by the generator to a value $\in [0, 1]$ to represent the probability $\hat{\mathcal{Q}}$ is a plausible human challenging motion. Specifically, we follow the video motion capture approach VIBE [35] to design two losses, including the adversarial loss \mathcal{L}_{adv} to backpropagate to the generator \mathcal{G} and the discriminator loss $\mathcal{L}_{\mathcal{D}}$ for the discriminator \mathcal{D} :

$$\mathcal{L}_{adv} = \mathbb{E}_{\hat{\mathcal{Q}} \sim p_{\mathcal{G}}} \left[\left(\mathcal{D}(\hat{\mathcal{Q}}) - 1 \right)^2 \right], \quad (3)$$

$$\mathcal{L}_{\mathcal{D}} = \mathbb{E}_{\mathcal{Q}^{mb} \sim p_V} \left[\left(\mathcal{D}(\mathcal{Q}^{mb}) - 1 \right)^2 \right] + \mathbb{E}_{\hat{\mathcal{Q}} \sim p_{\mathcal{G}}} \left[\left(\mathcal{D}(\hat{\mathcal{Q}}) \right)^2 \right]. \quad (4)$$

Here, the adversarial loss \mathcal{L}_{adv} is the expectation that $\hat{\mathcal{Q}}$ belongs to a plausible human challenging motion, while $p_{\mathcal{G}}$ and p_V represents the corrected motion sequence the corresponding captured marker-based challenging motion sequence, respectively. Note that \mathcal{Q}^{mb} denotes the accurate but temporally unsynchronized motion map captured by the marker-based system. Compared to VIBE [35] which extracts general motion prior, our scheme can recover the characteristics of more specific challenging motion.

Training Details. We train our HybridNet for 500 epochs with Adam optimizer [34], and set the dropout ratio as 0.1 for the GRU layers. We apply Exponential Linear Unit (ELU) activation and batch normalization layer after every our 1D convolutional layer with kernel size 7, except the final output layer before the decoder. During training, four NVidia 2080Ti GPUs are utilized. The batch size is set to be 32, while the learning rate is set to be 1×10^{-3} for the generator and 1×10^{-2} for the discriminator, the decay rate is 0.1 (final 100 epochs). To train our HybridNet, a new dataset with rich references modalities and various challenging motions and performers is further introduced and more details about our dataset are provided in Sec. 5

Our hybrid motion inference utilizes multi-modal references to extract fine motion details for challenging human motions in a data-driven manner. At test time, our method can robustly boost the tracking accuracy of the initial noisy skeletal motions via a novel generation network. Since our learning scheme is not directly supervised on the input images of specific performers, it's not restricted by per-performer training. Instead, our approach focus on extracting the characteristics of challenging motions directly.

4.2. Robust Motion Optimization

Besides the data-driven 3D motion characteristics from the previous stage, the input RGB video also encodes reliable motion hints for those non-extreme poses, especially for the non-occluded regions. We thus introduce this robust motion optimization to refine the skeletal motions so as

to increase the tracking accuracy and overlay performance, which jointly utilizes the learned 3D prior from the supervised multi-modal references as well as the reliable 2D and silhouette information from the input image references. The optimization to refine the skeletal pose is formulated as:

$$\mathbf{E}_{total}(\mathbf{S}_t) = \mathbf{E}_{3D} + \lambda_{2D} \mathbf{E}_{2D} + \lambda_T \mathbf{E}_T + \lambda_S \mathbf{E}_S. \quad (5)$$

Here, \mathbf{E}_{3D} makes the final motion sequence close to the output of network on occluded and invisible joints while \mathbf{E}_{2D} adds a re-projection constraint on high-confidence 2D keypoints detected. \mathbf{E}_T enforces the final motion to be temporally smooth, while the \mathbf{E}_S enforces alignment of the projected 3D model boundary with the detected silhouette. Specifically, the 3D term \mathbf{E}_{3D} is as following:

$$\mathbf{E}_{3D} = \sum_{t=1}^T \|\mathbf{S}_t - \mathcal{M}(\hat{\mathbf{Q}}_t, \mathbf{t}_t)\|_2^2, \quad (6)$$

where $\hat{\mathbf{Q}}_t$ is the regressed quaternions motion from our previous stage; \mathcal{M} is the mapping from quaternions to skeletal poses; \mathbf{t}_t is the global translation of \mathbf{S}_t . Note that the joint angles θ_t of \mathbf{S}_t locate in the pre-defined range $[\theta_{min}, \theta_{max}]$ of physically plausible joint angles to prevent unnatural poses. We then propose the projected 2D term as:

$$\mathbf{E}_{2D} = \frac{1}{T} \sum_{t=1}^T \frac{1}{|C_t|} \sum_{i \in C_t} \|\Pi(J_i(\mathbf{S}_t)) - \mathbf{p}_t^{(i)}\|_2^2, \quad (7)$$

where $C_t = \{i \mid c_t^{(i)} \geq \text{thred}\}$ is the set of indexes of high-confidence keypoints on the image I_t ; $c_t^{(i)}$ is the confidence value of the i^{th} keypoint $\mathbf{p}_t^{(i)}$; thred is 0.8 in our implementation. The projection function Π maps 3D joint positions to 2D coordinates while J_i computes the 3D position of the i^{th} joint. Then, the temporal term \mathbf{E}_T is formulated as:

$$\mathbf{E}_T = \sum_{t=1}^{T-1} \|\mathcal{M}(\hat{\mathbf{Q}}_t, \mathbf{t}_t) - \mathcal{M}(\hat{\mathbf{Q}}_{t+1}, \mathbf{t}_{t+1})\|_2^2, \quad (8)$$

where $\hat{\mathbf{Q}}_t$ and $\hat{\mathbf{Q}}_{t+1}$ are two adjacent regressed quaternions motion from our hybrid motion inference module. We utilize temporal smoothing to enable globally consistent capture in 3D space. Moreover, we follow [75] to formulate the silhouette term \mathbf{E}_S and please refer to [75] for more detail.

The constrained optimization problem to minimize the Eqn. 5 is solved using the Levenberg-Marquardt (LM) algorithm of *ceres* [1]. In all experiments, we use the following empirically determined parameters: $\lambda_{2D} = 1.0$, $\lambda_T = 20.0$ and $\lambda_S = 0.3$. Note that the initial \mathbf{t}_t for the mapping from quaternions to skeletal poses is obtained through the pre-processing stage in Sec. 4.1. To enable more robust optimization, we first optimize the global translation \mathbf{t}_t for all

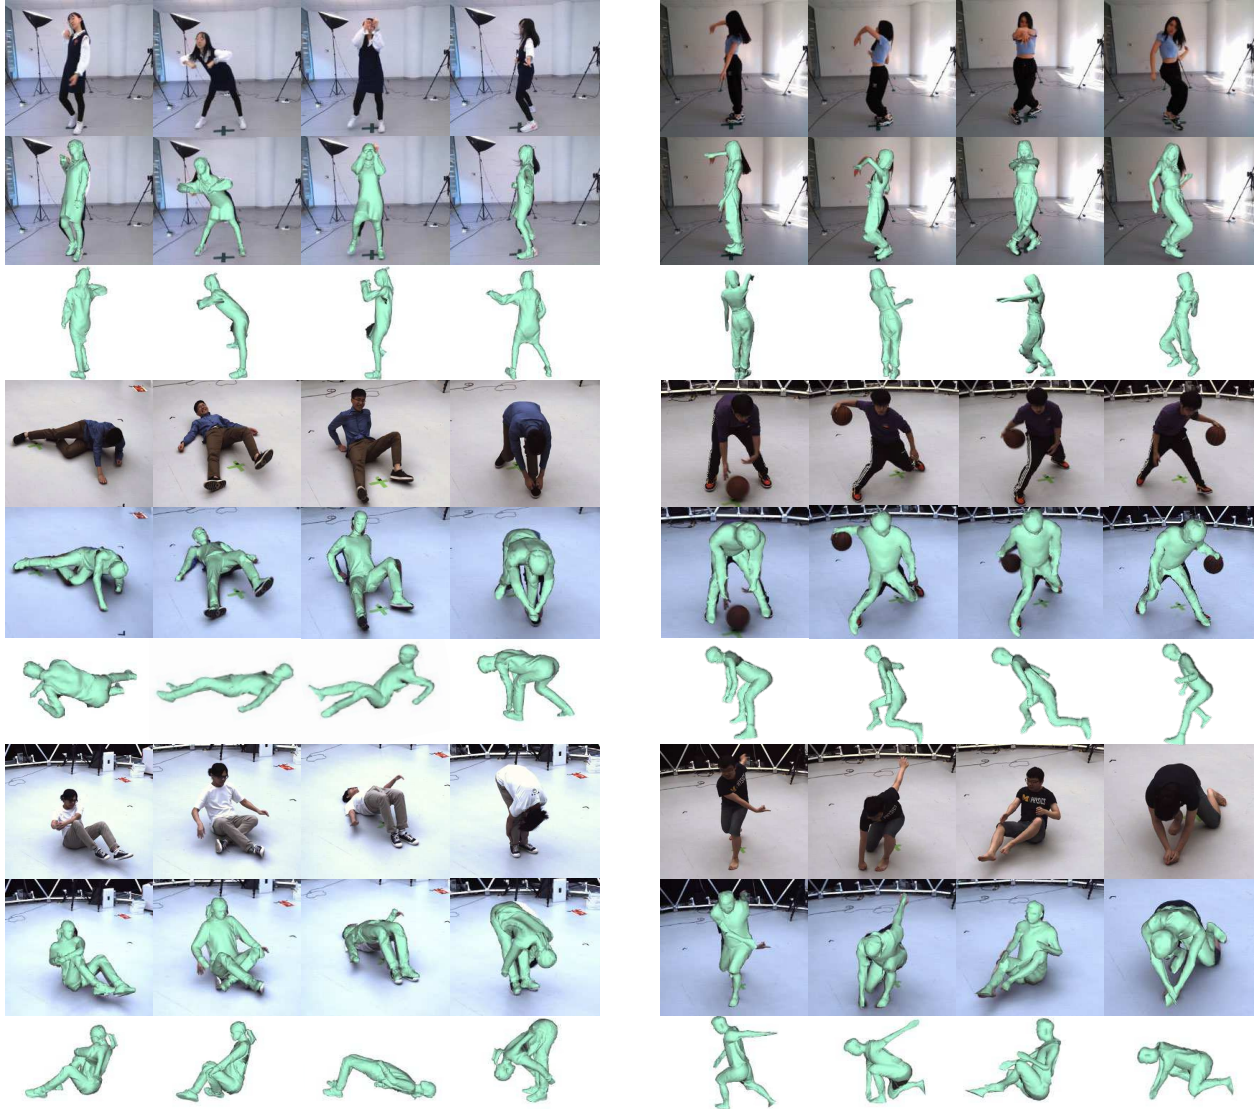


Figure 4. 3D capturing results on challenging human motions. For each body motions, the top row shows the input images, the middle row shows the captured body results on camera view, and the bottom row shows the rendering result of the captured body from side view.

the frames and then optimize the S_t . Such a flip-flop optimization strategy improves the overlay performance and tracking accuracy, which jointly utilizes the 3D challenging motion characteristics from the supervised multi-modal references as well as the reliable 2D and silhouette information from the input image references.

5. Experimental Results

In this section, we introduce our new dataset and evaluate our ChallenCap in a variety of challenging scenarios.

ChallenCap Dataset. There are existing datasets for 3D pose estimation and human performance capture, such as the Human3.6M[27] dataset that contains 3D human poses

in daily activities, but it lacks challenging motions and template human meshes. The AMASS [43] dataset provides a variety of human motions using marker-based approach, but the corresponding RGB videos are not provided. To evaluate our method, we propose a new challenging human motion dataset containing 20 different characters with a wide range of challenging motions such as dancing, boxing, gymnastic, exercise, basketball, yoga, rolling, leap, etc(see Fig.5). We adopt a 12-view marker-based Vicon motion capture system to capture challenging human motions. Each challenging motion consists of data from two modalities: synchronized sparse-view RGB video sequences and marker-based reference motion captured in an unsynchronized manner.




Figure 5. Illustration of our capturing system and examples of our dataset. The left shows our capturing system, including four RGB cameras (blue) for sparse-view image sequences and Vicon cameras (red) for marker-based motion capture (partially annotated). The right shows the animated meshes from the rigged character-wise template models with various challenging motions.




Figure 6. Qualitative comparison. Our results overlay better with the input video frames than the results of other methods.

Table 1. Quantitative comparison of several methods in terms of tracking accuracy and template mesh overlay.

| Method | MPJPE (<i>mm</i>) \downarrow | PCK0.5(%) \uparrow | PCK0.3 \uparrow | mIoU(%) \uparrow |
|------------------|----------------------------------|----------------------|-------------------|--------------------|
| HMR [32] | 154.3 | 77.2 | 68.9 | 57.0 |
| VIBE [35] | 116.7 | 83.7 | 71.8 | 73.7 |
| MonoPerfCap [75] | 134.7 | 77.4 | 65.6 | 65.5 |
| Ours | 52.6 | 96.6 | 87.4 | 83.6 |

5.1. Comparison

Our method enables more accurate motion capture for challenging human motions. For further comparison, we do experiments to demonstrate its effectiveness. We compare the proposed ChallenCap method with several monocular 3D human motion capture methods. Specifically, we apply MonoPerfCap [75] which is based on optimization. We also apply HMR [32] and VIBE [35] where the latter also relies on an adversarial learning framework. For fair comparisons, we fine-tune HMR and VIBE with part of manually annotated data from our dataset. As shown in Fig. 6, our method outperforms other methods in motion capture quality. Ben-

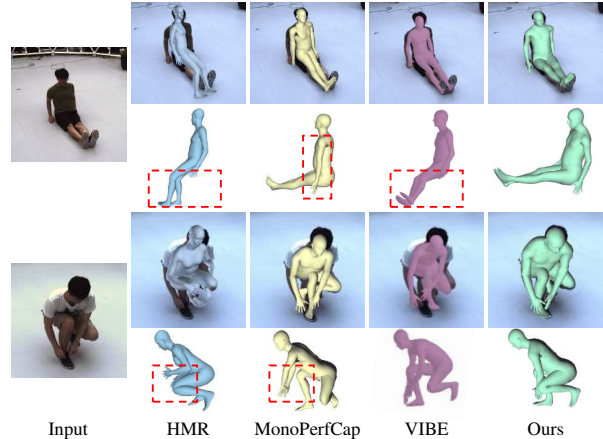


Figure 7. Qualitative comparison with side views. Our method maintains projective consistency in the side views while other methods have misalignment errors.




Figure 8. Qualitative comparison between ChallenCap (green) and MonoPerfCap (yellow) with reference-view verification. As marked in the figure, MonoPerfCap misses limbs in the reference view.

efiting from multi-modal references, our method performs better on the overall scale and also gets better overlays of the captured body.

As illustrated in Fig. 7, we perform a qualitative comparison with other methods on the main camera view and a corresponding side view. The figure shows that the side view results of other methods wrongly estimated the global position of arms or legs. This is mainly because our HybridNet promotes capture results for challenging human motions in the world frame.

Fig. 8 shows the reference view verification results. We capture challenging motions in the main view and verify the result in a reference view with a rendered mesh. The figure shows that even MonoPerfCap gets almost right 3D capture results in the main camera view, the results in the reference view show the misalignments on the limbs of the human body.

Tab. 1 shows the quantitative comparisons between our method and state-of-the-art methods using different evaluation metrics. We report the mean per joint position error (MPJPE), the Percentage of Correct Keypoints (PCK), and

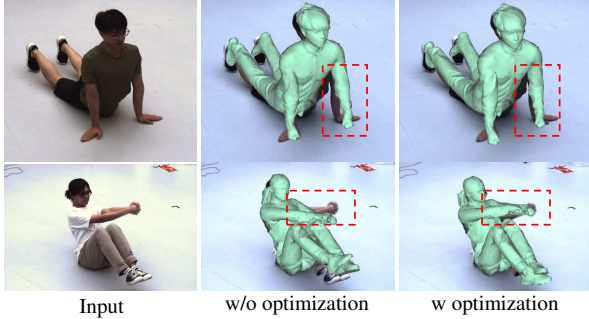


Figure 9. Evaluation for the optimization stage. The figure shows that our robust optimization stage improves the overlay performance.

Table 2. Quantitative evaluations on different optimization configurations.

| Method | MPJPE↓ | PCK0.5↑ | PCK0.3↑ |
|---------------------|-------------|-------------|-------------|
| MonoPerfCap | 134.7 | 77.4 | 65.6 |
| Ours | 106.5 | 92.8 | 82.2 |
| Ours + optimization | 52.6 | 96.6 | 87.4 |

the Mean Intersection-Over-Union (mIoU) results. Benefiting from our multi-modal references, our method gets better performance than optimization-based methods and other data-driven methods.

5.2. Evaluation

We apply two ablation experiments. One verifies that it is more effective to apply the robust optimization stage, the other validates that the diligently-designed network structure and loss functions gain for the challenge human motions by comparing with other network structures.

Evaluation on optimization. Tab.2 shows the performance of models with or without the optimization module. The baseline 3D capture method is MonoPerfCap [75]. The table demonstrates that whether the robust optimization stage is applied, our method outperforms MonoPerfCap. As illustrated in Fig.9, the obvious misalignment on the limb is corrected when the robust optimization stage is applied.

Evaluation on network structure. We experiment with several different network structures and loss design configurations, the results are demonstrated in Tab.3. The table shows that even only using sparse-view loss or adversarial loss, our method performs better than the simple encoder-decoder network structure without the attention pooling design. It also outperforms VIBE. When using both sparse-view loss and adversarial loss, our method gets 4% to 5% increase for PCK-0.5 and 12 to 32 decrease for MPJPE, compared with using only sparse-view or adversarial loss. The experiment illustrates multi-modal references contribute a lot to the improvement of results as illustrated in Fig. 10.




Figure 10. Evaluation of our network structure. The figure shows the effectiveness of both of our losses. Note that all experiments are applied with the robust optimization stage. Results of the full pipeline overlay more accurately with the input video frames.

Table 3. Quantitative evaluation of different network structure configurations. Our full pipeline achieves the lowest error.

| Method | MPJPE↓ | PCK0.5↑ | PCK0.3↑ |
|---|-------------|-------------|-------------|
| Encoder-Decoder | 109.1 | 85.7 | 81.5 |
| VIBE | 94.2 | 90.2 | 82.0 |
| Ours + \mathcal{L}_{sv} | 64.3 | 92.5 | 84.1 |
| Ours + \mathcal{L}_{adv} | 84.6 | 91.2 | 84.1 |
| Ours + $\mathcal{L}_{sv} + \mathcal{L}_{adv}$ | 52.6 | 96.6 | 87.4 |

6. Discussion

Limitation. As the first trial to explore challenging human motion capture with multi-modal references, the proposed ChallenCap still owns limitations as follows. First, our method relies on a pre-scanned template and cannot handle topological changes like clothes removal. Our method is also restricted to human reconstruction, without modeling human-object interactions. It’s interesting to model the human-object scenarios in a physically plausible way to capture more challenging and complicated motions. Besides, our current pipeline turns to utilize the references in a two-stage manner. It’s a promising direction to formulate the challenging motion capture problem in an end-to-end learning-based framework.

Conclusion. We present a robust template-based approach to capture challenging 3D human motions using only a single RGB camera, which is in a novel two-stage learning-and-optimization framework to make full usage of multi-modal references. Our hybrid motion inference learns the challenging motion details from various supervised references modalities, while our robust motion optimization further improves the tracking accuracy by extracting the reliable motion hints in the input image reference. Our experimental results demonstrate the effectiveness and robustness of ChallenCap in capturing challenging human motions in various scenarios. We believe that it is a significant step to enable robust 3D capture of challenging human motions, with many potential applications in VR/AR, and performance evaluation for gymnastics, sports, and dancing.

References

- [1] Sameer Agarwal, Keir Mierle, and Others. Ceres solver. <http://ceres-solver.org>. 5
- [2] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 1, 2
- [3] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: Shape completion and animation of people. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, page 408–416, New York, NY, USA, 2005. Association for Computing Machinery. 2
- [4] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Combining implicit function learning and parametric models for 3d human reconstruction. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 311–329, Cham, 2020. Springer International Publishing. 2
- [5] 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
- [6] Uttaran Bhattacharya, Christian Roncal, Trisha Mittal, Rohan Chandra, Aniket Bera, and Dinesh Manocha. Take an emotion walk: Perceiving emotions from gaits using hierarchical attention pooling and affective mapping. In *Proceedings of the European Conference on Computer Vision (ECCV)*, August 2020. 4
- [7] 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 Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 561–578, Cham, 2016. Springer International Publishing. 2
- [8] Chris Bregler and Jitendra Malik. Tracking people with twists and exponential maps. In *Computer Vision and Pattern Recognition (CVPR)*, 1998. 2
- [9] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 4
- [10] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [11] Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J. Black. Monocular expressive body regression through body-driven attention. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 20–40, Cham, 2020. Springer International Publishing. 2
- [12] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (TOG)*, 34(4):69, 2015. 1
- [13] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, pages 303–312, New York, NY, USA, 1996. ACM. 2
- [14] Andrew J. Davison, Jonathan Deutscher, and Ian D. Reid. Markerless motion capture of complex full-body movement for character animation. In *Eurographics Workshop on Computer Animation and Simulation*, 2001. 1
- [15] Edilson De Aguiar, Carsten Stoll, Christian Theobalt, Naveed Ahmed, Hans-Peter Seidel, and Sebastian Thrun. Performance capture from sparse multi-view video. pages 1–10, 2008. 2
- [16] Mingsong Dou, Philip Davidson, Sean Ryan Fanello, Sameh Khamis, Adarsh Kowdle, Christoph Rhemann, Vladimir Tankovich, and Shahram Izadi. Motion2fusion: Real-time volumetric performance capture. *ACM Trans. Graph.*, 36(6):246:1–246:16, Nov. 2017. 2
- [17] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, Pushmeet Kohli, Vladimir Tankovich, and Shahram Izadi. Fusion4D: Real-time Performance Capture of Challenging Scenes. In *ACM SIGGRAPH Conference on Computer Graphics and Interactive Techniques*, 2016. 2
- [18] Juergen Gall, Bodo Rosenhahn, Thomas Brox, and Hans-Peter Seidel. Optimization and filtering for human motion capture. *International Journal of Computer Vision (IJCV)*, 87(1–2):75–92, 2010. 2
- [19] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, and et al. The re-lightables: Volumetric performance capture of humans with realistic relighting. *ACM Trans. Graph.*, 38(6), Nov. 2019. 2
- [20] Kaiwen Guo, Feng Xu, Yangang Wang, Yebin Liu, and Qionghai Dai. Robust Non-Rigid Motion Tracking and Surface Reconstruction Using L0 Regularization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3083–3091, 2015. 2
- [21] Kaiwen Guo, Feng Xu, Tao Yu, Xiaoyang Liu, Qionghai Dai, and Yebin Liu. Real-time geometry, albedo and motion reconstruction using a single rgbd camera. *ACM Transactions on Graphics (TOG)*, 2017. 2
- [22] Marc Habermann, Weipeng Xu, Michael Zollhöfer, Gerard Pons-Moll, and Christian Theobalt. Livecap: Real-time human performance capture from monocular video. *ACM Transactions on Graphics (TOG)*, 38(2):14:1–14:17, 2019. 1, 3
- [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 (CVPR)*, June 2020. 1, 3
- [24] Nils Hasler, Bodo Rosenhahn, Thorsten Thormahlen, Michael Wand, Juergen Gall, and Hans-Peter Seidel. Markerless motion capture with unsynchronized moving cameras.

- In *Computer Vision and Pattern Recognition (CVPR)*, pages 224–231, 2009. 1
- [25] Y. Huang, F. Bogo, C. Lassner, A. Kanazawa, P. V. Gehler, J. Romero, I. Akhter, and M. J. Black. Towards accurate marker-less human shape and pose estimation over time. In *2017 International Conference on 3D Vision (3DV)*, pages 421–430, 2017. 2
- [26] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. VolumeDeform: Real-time Volumetric Non-rigid Reconstruction. October 2016. 2
- [27] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 6
- [28] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(7):1325–1339, 2014. 1
- [29] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic Studio: A Massively Multiview System for Social Motion Capture. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3334–3342, 2015. 1
- [30] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1, 2
- [31] 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
- [32] 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, 7
- [33] Angjoo Kanazawa, Jason Y. Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [34] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [35] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 2, 5, 7
- [36] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6050–6059, 2017. 2
- [38] Hao Li, Bart Adams, Leonidas J Guibas, and Mark Pauly. Robust single-view geometry and motion reconstruction. *ACM Transactions on Graphics (ToG)*, 28(5):1–10, 2009. 2
- [39] Ruilong Li, Yuliang Xiu, Shunsuke Saito, Zeng Huang, Kyle Olszewski, and Hao Li. Monocular real-time volumetric performance capture. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 49–67, Cham, 2020. Springer International Publishing. 2
- [40] Yebin Liu, Juergen Gall, Carsten Stoll, Qionghai Dai, Hans-Peter Seidel, and Christian Theobalt. Markerless motion capture of multiple characters using multiview image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(11):2720–2735, 2013. 1, 2
- [41] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: A skinned multi-person linear model. *ACM Trans. Graph.*, 34(6):248:1–248:16, Oct. 2015. 2, 3
- [42] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J. Black. Learning to dress 3d people in generative clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [43] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 1, 6
- [44] 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. *ACM Transactions on Graphics (TOG)*, 36(4), 2017. 1
- [45] Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, and Shigeo Morishima. Siclepe: Silhouette-based clothed people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [46] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 343–352, 2015. 2
- [47] Ahmed A. A. Osman, Timo Bolkart, and Michael J. Black. Star: Sparse trained articulated human body regressor. In *European Conference on Computer Vision (ECCV)*, volume LNCS 12355, pages 598–613, Aug. 2020. 2
- [48] 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 (CVPR)*, June 2020. 2
- [49] 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. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, June 2019. 2
- [50] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Harvesting multiple views for marker-less 3d human pose annotations. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [51] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J. Black. Clothcap: Seamless 4d clothing capture and retargeting. *ACM Trans. Graph.*, 36(4), July 2017. 2
- [52] Albert Pumarola, Jordi Sanchez-Riera, Gary P. T. Choi, Alberto Sanfeliu, and Francesc Moreno-Noguer. 3dpeople: Modeling the geometry of dressed humans. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [53] Nadia Robertini, Dan Casas, Helge Rhodin, Hans-Peter Seidel, and Christian Theobalt. Model-based outdoor performance capture. In *International Conference on 3D Vision (3DV)*, 2016. 2
- [54] 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 *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [55] 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 (CVPR)*, June 2020. 1, 2
- [56] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [57] M. Slavcheva, M. Baust, D. Cremers, and S. Ilic. Killing-Fusion: Non-rigid 3D Reconstruction without Correspondences. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [58] M. Slavcheva, M. Baust, and S. Ilic. SobolevFusion: 3D Reconstruction of Scenes Undergoing Free Non-rigid Motion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [59] M. Slavcheva, M. Baust, and S. Ilic. Variational Level Set Evolution for Non-rigid 3D Reconstruction from a Single Depth Camera. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2020. 2
- [60] Carsten Stoll, Nils Hasler, Juergen Gall, Hans-Peter Seidel, and Christian Theobalt. Fast articulated motion tracking using a sums of Gaussians body model. In *International Conference on Computer Vision (ICCV)*, 2011. 1, 2
- [61] Zhuo Su, Lan Xu, Zerong Zheng, Tao Yu, Yebin Liu, and Lu Fang. Robustfusion: Human volumetric capture with data-driven visual cues using a rgbd camera. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 246–264, Cham, 2020. Springer International Publishing. 2
- [62] Robert W Sumner, Johannes Schmid, and Mark Pauly. Embedded deformation for shape manipulation. *ACM Transactions on Graphics (TOG)*, 26(3):80, 2007. 2
- [63] Xin Suo, Yuheng Jiang, Pei Lin, Yingliang Zhang, Kaiwen Guo, Minye Wu, and Lan Xu. Neuralhumanfvv: Real-time neural volumetric human performance rendering using rgb cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021. 1
- [64] Sicong Tang, Feitong Tan, Kelvin Cheng, Zhaoyang Li, Siyu Zhu, and Ping Tan. A neural network for detailed human depth estimation from a single image. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [65] Christian Theobalt, Edilson de Aguiar, Carsten Stoll, Hans-Peter Seidel, and Sebastian Thrun. Performance capture from multi-view video. In *Image and Geometry Processing for 3-D Cinematography*, pages 127–149. Springer, 2010. 2
- [66] Vicon Motion Systems. <https://www.vicon.com/>, 2019. 1, 2
- [67] Daniel Vlasic, Rolf Adelsberger, Giovanni Vannucci, John Barnwell, Markus Gross, Wojciech Matusik, and Jovan Popović. Practical motion capture in everyday surroundings. *ACM transactions on graphics (TOG)*, 26(3):35–es, 2007. 2
- [68] Yangang Wang, Yebin Liu, Xin Tong, Qionghai Dai, and Ping Tan. Outdoor markerless motion capture with sparse handheld video cameras. *Transactions on Visualization and Computer Graphics (TVCG)*, 2017. 1
- [69] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [70] Xsens Technologies B.V. <https://www.xsens.com/>, 2019. 2
- [71] L. Xu, W. Cheng, K. Guo, L. Han, Y. Liu, and L. Fang. Fly-fusion: Realtime dynamic scene reconstruction using a flying depth camera. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2019. 2
- [72] Lan Xu, Yebin Liu, Wei Cheng, Kaiwen Guo, Guyue Zhou, Qionghai Dai, and Lu Fang. Flycap: Markerless motion capture using multiple autonomous flying cameras. *IEEE Transactions on Visualization and Computer Graphics*, 24(8):2284–2297, Aug 2018. 2
- [73] L. Xu, Z. Su, L. Han, T. Yu, Y. Liu, and L. FANG. Unstructuredfusion: Realtime 4d geometry and texture reconstruction using commercialrgbd cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019. 2
- [74] Lan Xu, Weipeng Xu, Vladislav Golyanik, Marc Habermann, Lu Fang, and Christian Theobalt. Eventcap: Monocular 3d capture of high-speed human motions using an event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 3
- [75] 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):27:1–27:15, 2018. [1](#), [3](#), [4](#), [5](#), [7](#), [8](#)
- [76] Tao Yu, Kaiwen Guo, Feng Xu, Yuan Dong, Zhaoqi Su, Jianhui Zhao, Jianguo Li, Qionghai Dai, and Yebin Liu. Body-fusion: 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. [2](#)
- [77] 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. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2019. [2](#)
- [78] Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Neural descent for visual 3d human pose and shape. *arXiv preprint arXiv:2008.06910*, 2020. [2](#)
- [79] Zerong Zheng, Tao Yu, Hao Li, Kaiwen Guo, Qionghai Dai, Lu Fang, and Yebin Liu. Hybridfusion: Real-time performance capture using a single depth sensor and sparse imus. In *European Conference on Computer Vision (ECCV)*, Sept 2018. [2](#)
- [80] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. [1](#), [2](#)
- [81] Hao Zhu, Xinxin Zuo, Sen Wang, Xun Cao, and Ruigang Yang. Detailed human shape estimation from a single image by hierarchical mesh deformation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [82] Michael Zollhöfer, Matthias Nießner, Shahram Izadi, Christoph Rehmann, Christopher Zach, Matthew Fisher, Chenglei Wu, Andrew Fitzgibbon, Charles Loop, Christian Theobalt, et al. Real-time Non-rigid Reconstruction using an RGB-D Camera. *ACM Transactions on Graphics (TOG)*, 33(4):156, 2014. [2](#)