

# AvatarCap: Animatable Avatar Conditioned Monocular Human Volumetric Capture

Zhe Li, Zerong Zheng, Hongwen Zhang, Chaonan Ji, and Yebin Liu

Department of Automation, Tsinghua University, China

**Abstract.** To address the ill-posed problem caused by partial observations in monocular human volumetric capture, we present AvatarCap, a novel framework that introduces animatable avatars into the capture pipeline for high-fidelity reconstruction in both visible and invisible regions. Our method firstly creates an animatable avatar for the subject from a small number ( $\sim 20$ ) of 3D scans as a prior. Then given a monocular RGB video of this subject, our method integrates information from both the image observation and the avatar prior, and accordingly reconstructs high-fidelity 3D textured models with dynamic details regardless of the visibility. To learn an effective avatar for volumetric capture from only few samples, we propose GeoTexAvatar, which leverages both geometry and texture supervisions to constrain the pose-dependent dynamics in a decomposed implicit manner. An avatar-conditioned volumetric capture method that involves a canonical normal fusion and a reconstruction network is further proposed to integrate both image observations and avatar dynamics for high-fidelity reconstruction in both observed and invisible regions. Overall, our method enables monocular human volumetric capture with detailed and pose-dependent dynamics, and the experiments show that our method outperforms state of the art.

## 1 Introduction

Human volumetric capture has been a popular research topic in computer vision for decades due to its potential value in Metaverse, holographic communication, video games, etc. Multi-view systems [4, 12, 37, 5, 38, 71, 47, 10, 51, 29, 73, 24, 62, 77, 55] can reconstruct high-resolution 3D human models using multiple RGB(D) sensors, but the sophisticated setup restricts their deployment in practice. To overcome this limitation, researchers have developed various technologies for monocular human reconstruction based on template tracking [84, 16, 19, 20], volumetric fusion [49, 74, 60] or single-image reconstruction [81, 52, 53, 25, 33, 22].

Despite the rapid development in monocular volumetric capture, most of the existing methods mainly focus on reconstructing visible surfaces according to direct observations and fail to recover the dynamic details in invisible regions. POSEFusion [35] addressed this limitation via integrating keyframes of similar poses from the whole RGBD sequence for invisible region reconstruction. However, it requires the subject to perform similar motions for multiple times facing




Fig. 1: **Overview of AvatarCap.** We present AvatarCap that leverages an animatable avatar learned from only a small number ( $\sim 20$ ) of scans for monocular human volumetric capture to realize high-fidelity reconstruction regardless of the visibility.

different directions. What’s worse, the fused invisible details are copied unaltered from other depth frames, thus suffering from poor pose generalization.

How to recover temporally coherent and pose-dependent details on invisible surfaces is an urgent and essential problem in monocular human volumetric capture. Recently, many works on pose-driven human avatars have arisen in the community. They create animatable avatars from various inputs, including scans [8, 54, 41, 43, 7], multi-view RGB videos [50, 36] and monocular depth measurements [6, 65]. In this paper, our key insight is that the pose-driven dynamics of person-specific avatars are exactly what is missing in monocular human volumetric capture. With this in mind, we propose **AvatarCap**, the first pipeline that combines person-specific animatable avatars with monocular human volumetric capture. Intuitively, the avatar encodes a data-driven prior about the pose-dependent dynamic details, which can compensate for the lack of complete observation in monocular inputs, enabling high-quality reconstruction of 3D models with dynamic details regardless of visibility.

Although introducing person-specific avatars into volumetric capture adds overhead in pipeline preparation, we believe that a data-driven prior of pose-dependent dynamics is indispensable for the future dynamic monocular human capture. In this paper, to make a trade-off between the ease of data acquisition and reconstruction quality, we choose to use only a small number ( $\sim 20$ ) of textured scans as the database. Note that it is challenging to learn a generalized avatar from only few scans, and state-of-the-art methods typically require hundreds of scans for creating one avatar [54, 43]. If only twenty scans are available, their results tend to be overfit and lack geometric details because they condition all the surface details (including pose-dependent and pose-agnostic ones) on the pose input. To address this challenge, we propose **GeoTexAvatar**, a decomposed representation that guarantees detail representation power and generalization capability. To be more specific, our representation distills pose-agnostic details as much as possible into a common implicit template [78], and models the remaining pose-driven dynamics with a pose-conditioned warping field. Such a disentanglement promotes better generalization since a large portion of geometric details are factored out as the common template and consequently the pose-dependent warping field is much easier to learn. On the other hand,

previous methods rely on solely geometric cues to learn the conditional warping fields [78], but we find that it is not enough because many types of cloth dynamics (e.g., cloth sliding) cannot be supervised by only geometry due to the ambiguity when establishing geometric correspondences. Therefore, we introduce an extra texture template represented by NeRF [46] to jointly constrain the pose-dependent warping field using both geometry and texture supervisions, which makes it possible to learn an accurate pose-conditioned warping field. As a result, the proposed GeoTexAvatar can not only preserve more details but also produce more reasonable pose-dependent dynamics for animation.

However, it is still not trivial to leverage the animatable avatar in the monocular capture pipeline. The main reason is the huge domain gap between the avatar prior and the monocular color input without any explicit 3D information. Fortunately, a 2D normal map with plentiful details can be extracted from the monocular color image [53], and we can use it to bridge the 3D avatar and the 2D RGB input. However, directly optimizing the avatar geometry using extremely dense non-rigid deformation [61] to fit the 2D normal map is difficult, if not infeasible, because it is ill-posed to force the surface to be consistent with the normal map without explicit 3D correspondences. To overcome this challenge, we propose *Avatar-conditioned Volumetric Capture* that splits the integration between the avatar and the normal maps into two steps, i.e., canonical normal fusion and model reconstruction. Specifically, the canonical normal fusion integrates the avatar normal and the image-observed one on the unified 2D canonical image plane. In this procedure, we formulate the fusion as an optimization on both the rotation grids and the normal maps to correct low-frequency normal orientation errors caused by inaccurate SMPL [39] fitting while maintaining high-frequency details. After that, a reconstruction network pretrained on a large-scale 3D human dataset [73] is used as a strong prior for producing a high-fidelity 3D human with full-body details from the fused normal maps.

This paper proposes the following contributions: 1) AvatarCap, a new framework that introduces animatable avatars into the monocular human volumetric capture pipeline to achieve detailed and dynamic capture regardless of the visibility (Sec. 3). 2) GeoTexAvatar, a new decomposed avatar representation that contains a pose-agnostic Geo-Text implicit template and a pose-dependent warping field to jointly constrain the pose-dependent dynamics using both geometry and texture supervisions for more detailed and well-generalized animation (Sec. 4). 3) Avatar-conditioned volumetric capture that contains a canonical normal fusion method and a reconstruction network to overcome the domain gap between the avatar prior and the monocular input for full-body high-fidelity reconstruction (Sec. 5). Code is available at <https://github.com/lizhe00/AvatarCap>.

## 2 Related Work

**Template Tracking.** Given a monocular RGB(D) video, many works utilize a template to fit each frame using skeletal motion [44] or non-rigid deformation [61]. Specifically, [31, 84, 16] solved the non-rigid warp field to track the

input depth stream, while [70, 15, 82, 23] tracked the skeletal motion of the template to fit the monocular input. LiveCap [19] and DeepCap [20] jointly solved or inferred both skeletal and non-rigid motions from a monocular RGB video. MonoClothCap [68] built a statistical deformation model based on SMPL to capture visible cloth dynamics. However, these methods only focus on fitting the template to explain the image observation while neglecting the dynamics in invisible regions.

**Volumetric Fusion.** Meanwhile, to realize real-time reconstruction from a single depth sensor, Newcombe *et al.* [49] pioneered to propose DynamicFusion that tracks and completes a canonical model in an incremental manner. It inspired a lot of following works [26, 56, 17, 72, 57, 30, 74, 79, 60] to incorporate different body priors or other cues to improve the performance. However, similar to methods based on template tracking, these works do not take into account the dynamic deformations in invisible regions. SimulCap [75] introduced cloth simulation into the volumetric fusion pipeline but its reconstruction quality is limited by a simple cloth simulator. POSEFusion [35] proposed to integrate multiple keyframes of similar poses to recover the dynamic details for the whole body, but this scheme leads to poor pose generalization, i.e., only those poses that are seen in different frames can be faithfully reconstructed.

**Single-image Reconstruction.** Recently, researchers have paid more and more attention to recovering 3D humans from single RGB(D) images by volume regression [63, 27, 81], visual hull [48], depth maps [11, 58], template deformation [83, 1] and implicit functions [52, 25, 33, 21, 22, 69]. For the implicit function representation, PIFuHD [53] introduced normal estimation to produce detailed geometry. PaMIR [80] and IPNet [3] combined a parametric body model (e.g., SMPL [39]) into the implicit function to handle challenging poses. However, without direct observation, these methods only recover over-smoothed invisible geometry without details. NormalGAN [64] inferred the back-view RGBD image from the input RGBD using a GAN [13] and then seamed them together. Unfortunately, the inferred details may be inconsistent with the pose or cloth type due to the limited variation in training data.

**Animatable Human Avatar** To create animatable human avatars, previous methods usually reconstruct a template and then model the pose-dependent dynamics of the character by physical simulation [14, 59] or deep learning [2, 67, 18]. Recent works proposed to directly learn an animatable avatar from the database, including scans [42, 54, 41, 43, 7], multi-view RGB videos [36, 50] and depth frames [6, 65, 9]. These works usually require a large amount of data to train a person-specific avatar; when only a small number of scans are available, they suffer from overfitting and struggle with pose generalization. Wang *et al.* [65] learned a meta prior to overcome this issue, but it remains difficult to apply their method for texture modeling.

### 3 Overview

As shown in Fig. 1, the whole framework of AvatarCap contains two main steps:

1. **Avatar Creation.** Before performing monocular volumetric capture, we collect a small number ( $\sim 20$ ) of textured scans for the subject as the database to construct his/her animatable avatar, which will be used to facilitate dynamic detail capture. To create an avatar with realistic details and generalization capability, we propose GeoTexAvatar, a representation that decomposes the dynamic level set function [54] into an implicit template (including occupancy [45] and radiance [46] fields) and a pose-dependent warping field, as shown in Fig. 2. We train the GeoTexAvatar network by supervising both the geometry and the texture using the textured scans.
2. **Avatar-conditioned Volumetric Capture.** With the avatar prior, we perform volumetric capture given the monocular RGB video input, as illustrated in Fig. 3. To address the domain gap between the avatar and the RGB input, we propose to use the surface normals as the intermediate to bridge each other. Specifically, we firstly estimate the visible normals from each RGB image, which are then mapped to the canonical space using the estimated SMPL pose [28, 76]. Then we generate the canonical avatar with pose-dependent dynamics given the pose and render the canonical normal maps from both the front and back views. The next step is to integrate the rendered normal maps with their image-based counterparts. To do so, we propose canonical normal fusion, which aims to correct low-frequency local normal orientations while maintaining high-frequency details from image observations. Finally, a pretrained reconstruction network is used to produce a high-fidelity human model conditioned on the integrated normal maps.

## 4 Avatar Creation

In this section, our goal is to learn an animatable avatar for volumetric capture. Following the practice of SCANimate [54], we fit SMPL to the raw 3D scans and transform them to a canonical pose via inverse skinning. We aim to construct an animatable avatar, represented as a pose-conditioned implicit function, from these canonicalized scans. Since only a small number ( $\sim 20$ ) of textured scans are available, we propose a decomposed implicit function to guarantee representation power and generalization capability (Sec. 4.1), which allows us to better leverage the geometry and texture information of training data (Sec. 4.2).

### 4.1 GeoTexAvatar Representation

Our representation is built upon the pose-conditioned implicit function in SCANimate [54], which is defined as  $f(\mathbf{x}_c, \boldsymbol{\theta}) = s$ , where  $s \in [0, 1]$  is an occupancy value,  $\mathbf{x}_c$  is a 3D point in the canonical space and  $\boldsymbol{\theta}$  is the SMPL pose parameters. The pose-dependent surface is represented by the level set of this implicit function:  $f(\mathbf{x}_c, \boldsymbol{\theta}) = 0.5$ . However, such an entangled representation conditions all the surface dynamics, including the pose-dependent deformations and the pose-agnostic details, on the pose input. Consequently, the animation results tend to lack pose-agnostic details when given an unseen pose.

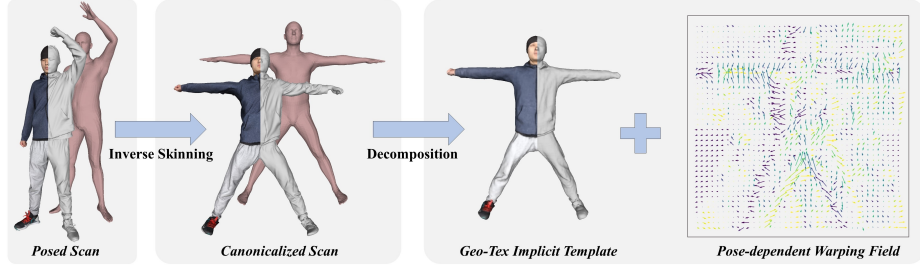


Fig. 2: **Illustration of the GeoTexAvatar representation.** We decompose the canonical scans into a pose-agnostic Geo-TeX implicit template and a pose-dependent warping field to enable joint supervisions by both geometry and texture for more detailed and well-generalized animation.

In order to not only model the pose-dependent deformations but also preserve the pose-agnostic details among different training samples, we propose a decomposed representation based on [78]:

$$T_{\text{Geo}}(W(\mathbf{x}_c, \boldsymbol{\theta})) = s, \quad (1)$$

where  $W(\mathbf{x}_c, \boldsymbol{\theta}) = \mathbf{x}_c + \Delta W(\mathbf{x}, \boldsymbol{\theta})$  represents the pose-dependent warping field that takes the pose parameters and a point as input and returns its template position, and  $T_{\text{Geo}}(\cdot)$  is the pose-agnostic occupancy template.

Note that previous avatars learned from scans [54, 41, 43] ignore the texture information even though their databases contain texture. However, we find the texture is essential to constrain the pose-dependent cloth deformations, because only geometrically closest constraints cannot establish correct correspondences, especially for common tangential cloth motions (e.g., cloth sliding). Therefore, we further introduce an extra texture template using the neural radiance field [46] (NeRF) in the same decomposed manner, i.e.,

$$T_{\text{Tex}}(W(\mathbf{x}_c, \boldsymbol{\theta})) = (\sigma, \mathbf{c}), \quad (2)$$

where  $T_{\text{Tex}}(\cdot)$  is a template radiance field that maps a template point to its density  $\sigma$  and color  $\mathbf{c}$ . Note that we utilize the template NeRF to represent the scan texture without view-dependent variation, so we discard the view direction input. Thanks to the decomposition (Eq. 1 & Eq. 2), our avatar representation, dubbed *GeoTexAvatar*, is able to jointly constrain the pose-dependent warping field  $W(\cdot)$  with the Geo-TeX implicit template field ( $T_{\text{Geo}}(\cdot)$  &  $T_{\text{Tex}}(\cdot)$ ) under the joint supervision of geometry and texture of training scans. Fig. 2 is an illustration of our representation.

Compared with state-of-the-art scan-based avatar methods [54, 43], our representation shows two main advantages as demonstrated in Fig. 7. 1) The decomposed representation can preserve more pose-agnostic details for animation. 2) The joint supervision of geometry and texture enables more reasonable pose-dependent deformations. What’s more, the decomposed representation allows us to finetune the texture template for high-quality rendering, which is also an advantage over other entangled methods as shown in Fig. 8.

## 4.2 GeoTexAvatar Training

The training loss for our GeoTexAvatar network contains a geometry loss, a texture loss and a regularization loss for the warping field, i.e.,  $\mathcal{L} = \lambda_{\text{geo}}\mathcal{L}_{\text{geo}} + \lambda_{\text{tex}}\mathcal{L}_{\text{tex}} + \lambda_{\text{reg}}\mathcal{L}_{\text{reg}}$ , where  $\lambda_{\text{geo}}$ ,  $\lambda_{\text{tex}}$  and  $\lambda_{\text{reg}}$  are the loss weights.

**Geometry Loss.**  $\mathcal{L}_{\text{geo}}$  penalizes the difference between the inferred occupancy and the ground truth:

$$\mathcal{L}_{\text{geo}} = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{x}_p \in \mathcal{P}} \text{BCE}(s(\mathbf{x}_p), s^*(\mathbf{x}_p)), \quad (3)$$

where  $\mathcal{P}$  is the sampled point set,  $s(\mathbf{x}_p)$  and  $s^*(\mathbf{x}_p)$  are inferred and ground-truth occupancy, respectively, and  $\text{BCE}(\cdot)$  measures the binary cross entropy.

**Texture Loss.** To jointly train the NeRF template, we render the textured scans to different views for the supervision.  $\mathcal{L}_{\text{tex}}$  measures the error between the color rendered by the network and the real one:

$$\mathcal{L}_{\text{tex}} = \frac{1}{|\mathcal{R}|} \sum_{\mathbf{r} \in \mathcal{R}} \left\| \hat{C}(\mathbf{r}) - C^*(\mathbf{r}) \right\|^2, \quad (4)$$

where  $\mathcal{R}$  is the set of ray samples in the image view frustum,  $\hat{C}(\cdot)$  is the volume rendering function as in [46], and  $C^*(\mathbf{r})$  is the ground-truth color.

**Regularization Loss.**  $\mathcal{L}_{\text{reg}}$  constrains the warped points by  $W(\cdot)$  to be close with the input because the canonical pose-dependent dynamics are usually small:

$$\mathcal{L}_{\text{reg}} = \frac{1}{|\mathcal{P} \cup \mathcal{P}_{\mathcal{R}}|} \sum_{\mathbf{x}_c \in \mathcal{P} \cup \mathcal{P}_{\mathcal{R}}} \left\| \Delta W(\mathbf{x}_c, \boldsymbol{\theta}) \right\|^2, \quad (5)$$

where  $\mathcal{P}_{\mathcal{R}}$  is the sampled points along each ray in  $\mathcal{R}$  during volume rendering.

## 5 Avatar-conditioned Volumetric Capture

Next, we move forward to the avatar-conditioned volumetric capture. The main difficulty lies in the enormous domain gap between the avatar representation and the input image, i.e., the image provides no 3D measurement to associate with the avatar geometry. As illustrated in Fig. 3, to overcome this challenge, we propose to employ the normal maps as the intermediate representation to bridge the gap between the image inputs and the avatar prior. Specifically, we conduct the integration between the two modals on a unified canonical image plane, and then split the integration as canonical normal fusion and model reconstruction.

**Initialization.** Given a RGB image, our approach firstly prepares both the canonical avatar and image normal maps as illustrated in Fig. 3. Specifically, 1) Avatar Animation: The GeoTexAvatar network outputs the animated canonical avatar using the SMPL pose, then renders the front & back canonical normal maps denoted as  $\mathbf{F}_{\text{avatar}}$  &  $\mathbf{B}_{\text{avatar}}$ , respectively. 2) Normal Map Canonicalization: In a parallel branch, the input RGB image is fed into a 2D convolutional network [66] to infer the normal map  $\mathbf{N}$  that represents the visible details. Then it is mapped to the canonical space, with the results denoted as  $\mathbf{F}_{\text{image}}$  and  $\mathbf{B}_{\text{image}}$ . Implementation details of the two steps can be found in the Supp. Mat.




Fig. 3: **Avatar-conditioned volumetric capture pipeline.** Given a RGB image from the monocular video, we firstly infer the SMPL pose and normal map. Then the pose-driven GeoTexAvatar generates and renders canonical avatar normal maps, while the image-observed normal map is warped into the canonical space. The canonical normal fusion integrates both avatar and observed normals together and feeds the fused normal maps into the reconstruction network to output a high-fidelity 3D human model. Finally, a high-resolution texture is generated using the GeoTexAvatar network.

### 5.1 Canonical Normal Fusion

Given the prepared avatar & image-observed normal maps, we integrate them on the 2D canonical image plane. However, directly replacing the avatar normals with the corresponding visible image-based ones is not feasible, because the canonicalized normal orientations may be incorrect due to the inaccurate SMPL estimation (e.g., rotation of the forearm) as shown in Fig. 4(a), leading to severe artifacts in the reconstruction (Fig. 4(d)). Therefore, we propose a new canonical normal fusion method to not only preserve high-frequency image-observed normals but also correct low-frequency local batch orientations.

Without loss of generality, we take the front avatar normal map and the image-based map ( $\mathbf{F}_{\text{avatar}}$  and  $\mathbf{F}_{\text{image}}$ ) as the example to introduce the formulation. As illustrated in Fig. 4(a),  $\mathbf{F}_{\text{image}}$  contains plentiful observed details estimated from the input color, but the orientations of normals are possibly incorrect due to SMPL estimation error. On the other hand, even though the visible region of  $\mathbf{F}_{\text{avatar}}$  does not completely follow the image observation, the low-frequency normal orientations are accurate in the canonical space as shown in Fig. 4(b). To this end, we propose to optimize the avatar normal map  $\mathbf{F}_{\text{avatar}}$  to integrate high-frequency details from the image-observed one  $\mathbf{F}_{\text{image}}$  while maintaining its initial correct low-frequency orientations as shown in Fig. 4(d). To do so, we introduce 2D rotation grids to factor out the low-frequency orientation differences between  $\mathbf{F}_{\text{avatar}}$  and  $\mathbf{F}_{\text{image}}$ , so that the remaining high-frequency details on  $\mathbf{F}_{\text{image}}$  can be rotated back to  $\mathbf{F}_{\text{avatar}}$  with correct orientations. As illustrated in Fig. 4(b), each grid is assigned a rotation matrix  $\mathbf{R}_i \in SO(3)$ , and the rotation of a 2D point  $\mathbf{p} = (x, y)$  on the map is defined as  $\mathbf{R}(\mathbf{p}) = \sum_i w_i(\mathbf{p})\mathbf{R}_i$ , a linear combination of  $\{\mathbf{R}_i\}$  using bilinear interpolation, where  $w_i(\mathbf{p})$  is the interpolation weight. With such a parameterization, we optimize the rotation grids  $\{\mathbf{R}_i\}$






Fig. 4: **Illustration of canonical normal fusion.** Directly replacing visible regions using image-observed normals causes severe reconstructed artifacts (d), while the proposed canonical normal fusion corrects the low-frequency local batch orientations and preserves high-frequency details for robust and high-fidelity reconstruction (c).

and avatar normal map  $\mathbf{F}_{\text{avatar}}$  by minimizing

$$E(\mathbf{R}_i, \mathbf{F}_{\text{avatar}}) = \lambda_{\text{fitting}} E_{\text{fitting}}(\mathbf{R}_i, \mathbf{F}_{\text{avatar}}) + \lambda_{\text{smooth}} E_{\text{smooth}}(\mathbf{R}_i), \quad (6)$$

where  $E_{\text{fitting}}$  and  $E_{\text{smooth}}$  are energies of misalignment between rotated avatar normals and observed ones and smooth regularization of grids, respectively.

**Fitting Term.** The fitting term measures the residuals between the avatar normal rotated by its transformation matrix and the target image-observed one:

$$E_{\text{fitting}}(\mathbf{R}_i, \mathbf{F}_{\text{avatar}}) = \sum_{\mathbf{p} \in \mathcal{D}} \|\mathbf{R}(\mathbf{p}) \mathbf{F}_{\text{avatar}}(\mathbf{p}) - \mathbf{F}_{\text{image}}(\mathbf{p})\|^2, \quad (7)$$

where  $\mathcal{D}$  is the valid intersection region of  $\mathbf{F}_{\text{avatar}}$  and  $\mathbf{F}_{\text{image}}$ .

**Smooth Term.** The smooth term regularizes the rotation grids to be low-frequency by constrain the rotation similarity between adjacent grids:

$$E_{\text{smooth}}(\mathbf{R}_i) = \sum_i \sum_{j \in \mathcal{N}(i)} \|\text{Rod}(\mathbf{R}_i) - \text{Rod}(\mathbf{R}_j)\|^2, \quad (8)$$

where  $\mathcal{N}(i)$  is the neighbors of the  $i$ -th grid, and  $\text{Rod} : SO(3) \rightarrow so(3)$  maps the rotation matrix to the axis-angle vector.

**Delayed Optimization of  $\mathbf{F}_{\text{avatar}}$ .** We firstly initialize  $\{\mathbf{R}_i\}$  as identity matrices. Note that both the avatar normal map  $\mathbf{F}_{\text{avatar}}$  and rotation grids  $\{\mathbf{R}_i\}$  are optimizable variables, so that the solutions are not unique. If we jointly optimizes both variables,  $\mathbf{F}_{\text{avatar}}$  tends to be equal with  $\mathbf{F}_{\text{image}}$  which is not desired. To this end, we firstly solve the low-frequency rotation grids  $\{\mathbf{R}_i\}$ , then optimize  $\mathbf{F}_{\text{avatar}}$  to integrate high-frequency details from  $\mathbf{F}_{\text{image}}$  with  $\{\mathbf{R}_i\}$  fixed. As a result, we obtain the optimized  $\mathbf{F}_{\text{avatar}}$  as the fused normal map  $\mathbf{F}_{\text{fused}}$  with high-frequency details and correct low-frequency orientations as shown in Fig. 4(c).

## 5.2 Model Reconstruction

**Geometric Reconstruction.** To reconstruct 3D geometry from the fused canonical normal maps  $\mathbf{F}_{\text{fused}}$  &  $\mathbf{B}_{\text{fused}}$ , we pretrain a reconstruction network on a large-scale 3D human dataset [73]. With such a strong data prior, we can efficiently and robustly recover the 3D geometry with high-fidelity full body details



Fig. 5: Example volumetric captured results of our method. From top to bottom are the monocular RGB input, geometric and textured results, respectively.

from the complete normal maps. The reconstruction network is formulated as an image-conditioned implicit function  $g(h(\pi(\mathbf{x}); \mathbf{F}_{\text{fused}}, \mathbf{B}_{\text{fused}}), \mathbf{x}_z)$ , where  $\mathbf{x}$  is a 3D point in the canonical space,  $h(\cdot)$  is a function to sample convoluted image features,  $\pi(\cdot)$  is the orthographic projection,  $\mathbf{x}_z$  is the z-axis value, and  $g(\cdot)$  is an implicit function that maps the image feature and  $\mathbf{x}_z$  to an occupancy value. We perform Marching Cubes [40] on this implicit function to reconstruct the canonical model, then deform it to the posed space by forward skinning.

**Texture Generation.** Based on the GeoTexAvatar representation, we can generate the texture of the reconstructed geometry by mapping the radiance field to it. Specifically, given a vertex  $\mathbf{v}$  of the canonical model and its normal  $\mathbf{n}_v$ , based on Eq. 2, we can calculate its color using volume rendering in NeRF [46] with the camera ray  $\mathbf{r}(t) = \mathbf{v} - t\mathbf{n}_v$  and near and far bounds  $-\delta$  and  $\delta$  ( $\delta > 0$ ).

## 6 Results

The volumetric captured results of our method are demonstrated in Fig. 5. For the experiments, we collect textured scans of 10 subjects and their monocular videos, and partial scans are utilized as the evaluation dataset. More details about data preprocessing and implementation can be found in the Supp. Mat.

### 6.1 Comparison

**Volumetric Capture.** As shown in Fig. 6, we compare AvatarCap, our whole volumetric capture framework, against state-of-the-art fusion and single-RGB(D)-image reconstruction methods, including POSEFusion [35], PIFuHD [53] and NormalGAN [64]. We conduct this comparison on the sequences captured by one Kinect Azure to also compare against RGBD-based methods [64, 35], and




Fig. 6: **Qualitative comparison against monocular volumetric capture methods.** We show reconstructed results of our method (AvatarCap), POSEFusion [35], PIFuHD [53] and NormalGAN [64]. And our method outperforms others on the capture of pose-dependent dynamics in the invisible regions (red circles).

Metric/Method	AvatarCap (ours)	PIFuHD [53]	NormalGAN [64]
Chamfer Distance	<b>1.097</b>	3.400	2.852
Scan-to-Mesh Distance	<b>1.096</b>	3.092	2.855

Table 1: **Quantitative comparison of AvatarCap with PIFuHD [53] and NormalGAN [64].** We report the averaged Chamfer and Scan-to-Model distance errors ( $\times 10^{-2}$  m) of different methods on the whole testing dataset.

all the learning-based methods are finetuned on the person-specific scans used in our avatar creation for fairness. Fig. 6 shows that our method can achieve high-fidelity reconstruction with detailed observations and reasonable pose-dependent invisible dynamics. Though POSEFusion [35] can integrate invisible surfaces from other frames, it entirely relies on each time captured sequence without pose generalization. PIFuHD [53] only considers to recover visible details from the normal map inferred by the color input without pose-conditioned person-specific dynamics, so the invisible regions are usually oversmoothed. Though NormalGAN [64] can infer a plausible back RGBD map from the RGBD input, the inferred invisible appearance may be inconsistent with the person-specific dynamics. We also conduct quantitative comparison on the testing dataset with ground-truth scans, and report the averaged errors in Tab. 1. Note that POSEFusion is a sequence-based method, but the testing scans are under discrete poses, so we only compare with the other methods. Overall, our method achieves state-of-the-art capture on both quality and accuracy.

**Animatable Avatar.** As shown in Fig. 7, we compare our avatar module, GeoTexAvatar, against state-of-the-art avatar works based on person-specific scans, SCANimate [54], SCALE [41] and POP [43]. Note that POP is a multi-subject-outfit representation, in this comparison we train it from scratch using the same few ( $\sim 20$ ) scans as other methods. Our method outperforms these methods on the recovery of dynamic details as well as tangential cloth motion benefiting from the proposed decomposed representation and joint supervisions of both geometry and texture, respectively. We further quantitatively evaluate




Fig. 7: **Qualitative comparison against animatable avatar methods.** We show animated results of our method (also with high-quality texture), SCANimate [54], SCALE [41] and POP [43]. And our method shows the superiority on the modeling of wrinkles (solid circles) and pose-dependent cloth tangential motions (dotted circles).

Case/Method	GeoTexAvatar (Ours)	SCANimate [54]	SCALE [41]	POP [43]
HOODY_1	<b>6.29</b>	7.38	8.19	6.83
SHIRT_1	<b>2.80</b>	5.72	4.72	3.08

Table 2: **Quantitative comparison of GeoTexAvatar with SCANimate [54], SCALE [41] and POP [43].** We report the averaged Chamfer distance errors ( $\times 10^{-3}$  m) between the animated results of different methods and the ground-truth scans.

the animation accuracy of GeoTexAvatar and other works on the testing dataset in Tab. 2, and our method achieves more accurate animated results.

## 6.2 Evaluation

**Decomposed Representation of GeoTexAvatar.** We evaluate the proposed decomposed representation compared with the entangled one [54] in Fig. 8. Firstly, similar to SCANimate [54], we choose the local pose vector as the pose encoding of the warping field in our representation. Compared with the entangled representation (Fig. 8 (a)), the decomposed one (Fig. 8 (b)) produces more detailed animation results, e.g., the zippers, facial and leg details, thanks to the decomposition of pose-dependent dynamics and pose-agnostic details. Besides, the decomposition allows us to finetune the texture template on a single scan to recovery high-quality texture, while the texture is totally blurred in the entangled learning. Furthermore, we empirically find that a SMPL positional map defined in the canonical space shows more powerful expression for pose-dependent dynamics than the local pose vector as shown in Fig. 8 (b) and (c).

**Texture Supervision in GeoTexAvatar.** We evaluate the effectiveness of texture supervision to the pose-dependent warping field by visualizing the correspondences during animation in Fig. 9. We firstly train the avatar network with and without texture template individually. To visualize the correspondences among animated results by different poses, we firstly generate the geometric tem-




Fig. 8: **Evaluation of the decomposed representation.** (a), (b) and (c) are the animated geometric and textured results of the entangled representation [54] and decomposed representations with pose-vector and positional-map encoding, respectively.




Fig. 9: **Evaluation of the effectiveness of texture supervision in GeoTexAvatar.** We visualize the correspondences among different frames by the vertex color which indicates whether the vertex belongs to the upper or lower body.

plate using [40] on  $T_{\text{Geo}}(\cdot)$ , then manually segment the template mesh as upper and lower body parts. Given a new pose, the avatar network outputs a canonical avatar model, then each vertex on this model can be warped to the template using the pose-dependent warping field. Finally, we determine whether the warped vertex belongs to the upper or lower part by its closest point on the template. Fig. 9 demonstrates that the texture supervision can implicitly constrain the warping field by jointly learning an extra texture template  $T_{\text{Tex}}(\cdot)$ , thus enabling more reasonable pose generalization for animation. However, training only with geometry supervision results in overfitted animation due to the ambiguity when establishing correspondences under only the geometric closest constraint.

**Training Scan Number.** We quantitatively evaluate the effect of the training scan number on the animation accuracy of GeoTexAvatar. We choose one subset (“SHORT\_SLEEVE\_1”) of our dataset that contains 100 scans, and randomly choose 80 items as the training dataset, and the rest for evaluation. Fig. 10 shows the visualized and numerical animation errors using different numbers of scans. More training samples does not always lead to more accurate results because the mapping from poses to cloth details may be one-to-many in the training dataset.

**Canonical Normal Fusion.** We evaluate the proposed canonical normal fusion compared with the directly covering method both qualitatively and quantitatively. Fig. 11 (c) and (d) show the reconstructed results using directly covering and canonical normal fusion, respectively, as well as their per-vertex point-to-surface errors to the ground-truth scans. Due to the inaccurate SMPL estimation and the camera view difference with the orthographic hypothesis in normal map inference, the canonicalized normal tends to be fallacious. The directly covering method maintains the wrong image-observed normal, thus leading to inaccurate reconstruction and ghosting artifacts as shown in Fig. 11(c). By contrary,




Fig. 10: **Evaluation of the effect of the training scan number on the animation accuracy of the GeoTexAvatar.** From (a) to (d) are visualized vertex-to-surface error of animated results trained by 20, 40, 60 and 80 scans, respectively, (e) is the ground-truth scan, and (f) is the chart of the averaged Chamfer and Scan-to-Mesh distance errors on the whole testing dataset.




Fig. 11: **Evaluation of canonical normal fusion.** We visualize the per-vertex point-to-surface errors between the reconstructed models and the ground truth, (f) is the percentage of correct vertices under different thresholds

the canonical normal fusion not only corrects the low-frequency orientations of canonicalized image normal, but also maintains the high-frequency details from the image observation, thus enabling the following accurate and high-fidelity reconstruction as shown in Fig. 11 (d) and (f).

## 7 Discussion

**Conclusion.** We present AvatarCap, a novel monocular human volumetric capture framework, that leverages an animatable avatar learned from only few scans to capture body dynamics regardless of the visibility. Based on the proposed GeoTexAvatar and avatar-conditioned volumetric capture, our method effectively integrates the information from image observations and the avatar prior. Overall, our method outperforms other state-of-the-art capture approaches, and we believe that the avatar-conditioned volumetric capture will make progress towards dynamic and realistic 3D human with the advance of animatable avatars.

**Limitation.** The main limitation of our method is the 3D scan collection, a possible solution is to capture scans using 3D self-portrait methods [32, 34] with an RGBD camera. Moreover, our method may fail for loose clothes, e.g., long skirts, because the SMPL skeletons cannot correctly deform such garments.

**Acknowledgement.** This paper is supported by National Key R&D Program of China (2021ZD0113501) and the NSFC project No.62125107.

## References

1. Alldieck, T., Pons-Moll, G., Theobalt, C., Magnor, M.: Tex2shape: Detailed full human body geometry from a single image. In: ICCV. pp. 2293–2303 (2019)
2. Bagautdinov, T., Wu, C., Simon, T., Prada, F., Shiratori, T., Wei, S.E., Xu, W., Sheikh, Y., Saragih, J.: Driving-signal aware full-body avatars. TOG **40**(4), 1–17 (2021)
3. Bhatnagar, B.L., Sminchisescu, C., Theobalt, C., Pons-Moll, G.: Combining implicit function learning and parametric models for 3d human reconstruction. In: ECCV. pp. 311–329. Springer (2020)
4. Bradley, D., Popa, T., Sheffer, A., Heidrich, W., Boubekeur, T.: Markerless garment capture. TOG **27**(3), 1–9 (2008)
5. Brox, T., Rosenhahn, B., Gall, J., Cremers, D.: Combined region and motion-based 3d tracking of rigid and articulated objects. IEEE T-PAMI **32**(3), 402–415 (2009)
6. Burov, A., Nießner, M., Thies, J.: Dynamic surface function networks for clothed human bodies. In: ICCV. pp. 10754–10764 (2021)
7. Chen, X., Zheng, Y., Black, M.J., Hilliges, O., Geiger, A.: Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In: ICCV. pp. 11594–11604 (2021)
8. Deng, B., Lewis, J.P., Jeruzalski, T., Pons-Moll, G., Hinton, G., Norouzi, M., Tagliasacchi, A.: Nasa neural articulated shape approximation. In: ECCV. pp. 612–628. Springer (2020)
9. Dong, Z., Guo, C., Song, J., Chen, X., Geiger, A., Hilliges, O.: Pina: Learning a personalized implicit neural avatar from a single rgb-d video sequence. In: CVPR (2022)
10. Dou, M., Khamis, S., Degtyarev, Y., Davidson, P., Fanello, S.R., Kowdle, A., Escolano, S.O., Rhemann, C., Kim, D., Taylor, J., et al.: Fusion4d: Real-time performance capture of challenging scenes. TOG **35**(4), 1–13 (2016)
11. Gabeur, V., Franco, J.S., Martin, X., Schmid, C., Rogez, G.: Moulding humans: Non-parametric 3d human shape estimation from single images. In: ICCV. pp. 2232–2241 (2019)
12. Gall, J., Stoll, C., De Aguiar, E., Theobalt, C., Rosenhahn, B., Seidel, H.P.: Motion capture using joint skeleton tracking and surface estimation. In: CVPR. pp. 1746–1753. IEEE (2009)
13. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. NeurIPS **27** (2014)
14. Guan, P., Reiss, L., Hirshberg, D.A., Weiss, A., Black, M.J.: Drape: Dressing any person. TOG **31**(4), 1–10 (2012)
15. Guo, C., Chen, X., Song, J., Hilliges, O.: Human performance capture from monocular video in the wild. In: 3DV. pp. 889–898. IEEE (2021)
16. Guo, K., Xu, F., Wang, Y., Liu, Y., Dai, Q.: Robust non-rigid motion tracking and surface reconstruction using l0 regularization. In: ICCV. pp. 3083–3091 (2015)
17. Guo, K., Xu, F., Yu, T., Liu, X., Dai, Q., Liu, Y.: Real-time geometry, albedo and motion reconstruction using a single rgb-d camera. TOG **36**(3), 32:1–32:13 (2017)
18. Habermann, M., Liu, L., Xu, W., Zollhoefer, M., Pons-Moll, G., Theobalt, C.: Real-time deep dynamic characters. TOG **40**(4), 1–16 (2021)
19. Habermann, M., Xu, W., Zollhoefer, M., Pons-Moll, G., Theobalt, C.: LiveCap: Real-time human performance capture from monocular video. TOG **38**(2), 1–17 (2019)



20. Habermann, M., Xu, W., Zollhofer, M., Pons-Moll, G., Theobalt, C.: Deepcap: Monocular human performance capture using weak supervision. In: CVPR. pp. 5052–5063 (2020)
21. He, T., Collomosse, J., Jin, H., Soatto, S.: Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. NeurIPS **33**, 9276–9287 (2020)
22. He, T., Xu, Y., Saito, S., Soatto, S., Tung, T.: Arch++: Animation-ready clothed human reconstruction revisited. In: ICCV. pp. 11046–11056 (2021)
23. He, Y., Pang, A., Chen, X., Liang, H., Wu, M., Ma, Y., Xu, L.: Challengcap: Monocular 3d capture of challenging human performances using multi-modal references. In: CVPR. pp. 11400–11411 (2021)
24. Hong, Y., Zhang, J., Jiang, B., Guo, Y., Liu, L., Bao, H.: Stereopifu: Depth aware clothed human digitization via stereo vision. In: CVPR. pp. 535–545 (2021)
25. Huang, Z., Xu, Y., Lassner, C., Li, H., Tung, T.: Arch: Animatable reconstruction of clothed humans. In: CVPR. pp. 3093–3102 (2020)
26. Innmann, M., Zollhöfer, M., Nießner, M., Theobalt, C., Stamminger, M.: Volumedeform: Real-time volumetric non-rigid reconstruction. In: ECCV. vol. 9912, pp. 362–379. SPRINGER (2016)
27. Jackson, A.S., Manafas, C., Tzimiropoulos, G.: 3d human body reconstruction from a single image via volumetric regression. In: ECCV Workshops. pp. 0–0 (2018)
28. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In: ICCV. pp. 2252–2261 (2019)
29. Leroy, V., Franco, J.S., Boyer, E.: Multi-view dynamic shape refinement using local temporal integration. In: ICCV. pp. 3094–3103 (2017)
30. Li, C., Zhang, Z., Guo, X.: Articulatedfusion: Real-time reconstruction of motion, geometry and segmentation using a single depth camera. In: ECCV. pp. 324–40. SPRINGER, Munich (2018)
31. Li, H., Adams, B., Guibas, L.J., Pauly, M.: Robust single-view geometry and motion reconstruction. TOG **28**(5), 1–10 (2009)
32. Li, H., Vouga, E., Gudym, A., Luo, L., Barron, J.T., Gusev, G.: 3d self-portraits. TOG **32**(6), 1–9 (2013)
33. Li, R., Xiu, Y., Saito, S., Huang, Z., Olszewski, K., Li, H.: Monocular real-time volumetric performance capture. In: ECCV (2020)
34. Li, Z., Yu, T., Pan, C., Zheng, Z., Liu, Y.: Robust 3d self-portraits in seconds. In: CVPR. pp. 1344–1353 (2020)
35. Li, Z., Yu, T., Zheng, Z., Guo, K., Liu, Y.: Posefusion: Pose-guided selective fusion for single-view human volumetric capture. In: CVPR. pp. 14162–14172 (2021)
36. Liu, L., Habermann, M., Rudnev, V., Sarkar, K., Gu, J., Theobalt, C.: Neural actor: Neural free-view synthesis of human actors with pose control. TOG **40**(6), 1–16 (2021)
37. Liu, Y., Dai, Q., Xu, W.: A point-cloud-based multiview stereo algorithm for free-viewpoint video. TVCG **16**(3), 407–418 (2009)
38. Liu, Y., Stoll, C., Gall, J., Seidel, H.P., Theobalt, C.: Markerless motion capture of interacting characters using multi-view image segmentation. In: CVPR. pp. 1249–1256. IEEE (2011)
39. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. TOG **34**(6), 1–16 (2015)
40. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. TOG **21**(4), 163–169 (1987)



41. Ma, Q., Saito, S., Yang, J., Tang, S., Black, M.J.: Scale: Modeling clothed humans with a surface codec of articulated local elements. In: CVPR. pp. 16082–16093 (2021)
42. Ma, Q., Yang, J., Ranjan, A., Pujades, S., Pons-Moll, G., Tang, S., Black, M.J.: Learning to dress 3d people in generative clothing. In: CVPR. pp. 6469–6478 (2020)
43. Ma, Q., Yang, J., Tang, S., Black, M.J.: The power of points for modeling humans in clothing. In: ICCV. pp. 10974–10984 (2021)
44. Magnenat-Thalmann, N., Laperrire, R., Thalmann, D.: Joint-dependent local deformations for hand animation and object grasping. In: In Proceedings on Graphics Interface. Citeseer (1988)
45. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: CVPR. pp. 4460–4470 (2019)
46. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV. pp. 405–421. Springer (2020)
47. Mustafa, A., Kim, H., Guillemot, J.Y., Hilton, A.: General dynamic scene reconstruction from multiple view video. In: ICCV. pp. 900–908 (2015)
48. Natsume, R., Saito, S., Huang, Z., Chen, W., Ma, C., Li, H., Morishima, S.: Siclope: Silhouette-based clothed people. In: CVPR. pp. 4480–4490 (2019)
49. Newcombe, R.A., Fox, D., Seitz, S.M.: Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In: CVPR. pp. 343–352 (2015)
50. Peng, S., Dong, J., Wang, Q., Zhang, S., Shuai, Q., Zhou, X., Bao, H.: Animatable neural radiance fields for modeling dynamic human bodies. In: ICCV. pp. 14314–14323 (2021)
51. Pons-Moll, G., Pujades, S., Hu, S., Black, M.J.: Clothcap: Seamless 4d clothing capture and retargeting. TOG **36**(4), 1–15 (2017)
52. Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: ICCV. pp. 2304–2314 (2019)
53. Saito, S., Simon, T., Saragih, J., Joo, H.: Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In: CVPR (June 2020)
54. Saito, S., Yang, J., Ma, Q., Black, M.J.: Scanimate: Weakly supervised learning of skinned clothed avatar networks. In: CVPR. pp. 2886–2897 (2021)
55. Shao, R., Zhang, H., Zhang, H., Chen, M., Cao, Y., Yu, T., Liu, Y.: Doublefield: Bridging the neural surface and radiance fields for high-fidelity human reconstruction and rendering. In: CVPR (2022)
56. Slavcheva, M., Baust, M., Cremers, D., Ilic, S.: Killingfusion: Non-rigid 3d reconstruction without correspondences. In: CVPR. pp. 1386–1395 (2017)
57. Slavcheva, M., Baust, M., Ilic, S.: Sobolevfusion: 3d reconstruction of scenes undergoing free non-rigid motion. In: CVPR. pp. 2646–2655. IEEE, Salt Lake City (June 2018)
58. Smith, D., Loper, M., Hu, X., Mavroidis, P., Romero, J.: Facsimile: Fast and accurate scans from an image in less than a second. In: ICCV. pp. 5330–5339 (2019)
59. Stoll, C., Gall, J., De Aguiar, E., Thrun, S., Theobalt, C.: Video-based reconstruction of animatable human characters. TOG **29**(6), 1–10 (2010)
60. Su, Z., Xu, L., Zheng, Z., Yu, T., Liu, Y., Fang, L.: Robustfusion: human volumetric capture with data-driven visual cues using a rgbd camera. In: ECCV (2020)
61. Sumner, R.W., Schmid, J., Pauly, M.: Embedded deformation for shape manipulation. TOG **26**(3), 80–es (2007)

62. Suo, X., Jiang, Y., Lin, P., Zhang, Y., Wu, M., Guo, K., Xu, L.: Neuralhumanfvv: Real-time neural volumetric human performance rendering using rgb cameras. In: CVPR. pp. 6226–6237 (2021)
63. Varol, G., Ceylan, D., Russell, B., Yang, J., Yumer, E., Laptev, I., Schmid, C.: Bodynet: Volumetric inference of 3d human body shapes. In: ECCV. pp. 20–36 (2018)
64. Wang, L., Zhao, X., Yu, T., Wang, S., Liu, Y.: Normalgan: Learning detailed 3d human from a single rgb-d image. In: ECCV. pp. 430–446. Springer (2020)
65. Wang, S., Mihajlovic, M., Ma, Q., Geiger, A., Tang, S.: Metaavatar: Learning animatable clothed human models from few depth images. *NeurIPS* **34** (2021)
66. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: CVPR. pp. 8798–8807 (2018)
67. Xiang, D., Prada, F., Bagautdinov, T., Xu, W., Dong, Y., Wen, H., Hodgins, J., Wu, C.: Modeling clothing as a separate layer for an animatable human avatar. *TOG* **40**(6), 1–15 (2021)
68. Xiang, D., Prada, F., Wu, C., Hodgins, J.: Monoclothcap: Towards temporally coherent clothing capture from monocular rgb video. In: 3DV. pp. 322–332. IEEE (2020)
69. Xiu, Y., Yang, J., Tzionas, D., Black, M.J.: Icon: Implicit clothed humans obtained from normals. In: CVPR (2022)
70. Xu, W., Chatterjee, A., Zollhöfer, M., Rhodin, H., Mehta, D., Seidel, H.P., Theobalt, C.: Monoperfcap: Human performance capture from monocular video. *TOG* **37**(2), 1–15 (2018)
71. Ye, G., Liu, Y., Hasler, N., Ji, X., Dai, Q., Theobalt, C.: Performance capture of interacting characters with handheld kinects. In: ECCV. pp. 828–841. Springer (2012)
72. Yu, T., Guo, K., Xu, F., Dong, Y., Su, Z., Zhao, J., Li, J., Dai, Q., Liu, Y.: Bodyfusion: Real-time capture of human motion and surface geometry using a single depth camera. In: ICCV. pp. 910–919. IEEE, Venice (2017)
73. Yu, T., Zheng, Z., Guo, K., Liu, P., Dai, Q., Liu, Y.: Function4d: Real-time human volumetric capture from very sparse consumer rgb-d sensors. In: CVPR. pp. 5746–5756 (2021)
74. Yu, T., Zheng, Z., Guo, K., Zhao, J., Dai, Q., Li, H., Pons-Moll, G., Liu, Y.: Doublefusion: Real-time capture of human performances with inner body shapes from a single depth sensor. In: CVPR. pp. 7287–7296. IEEE, Salt Lake City (June 2018)
75. Yu, T., Zheng, Z., Zhong, Y., Zhao, J., Dai, Q., Pons-Moll, G., Liu, Y.: Simulcap: Single-view human performance capture with cloth simulation. In: CVPR. pp. 5499–5509. IEEE (2019)
76. Zhang, H., Tian, Y., Zhou, X., Ouyang, W., Liu, Y., Wang, L., Sun, Z.: Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In: ICCV. pp. 11446–11456 (2021)
77. Zheng, Y., Shao, R., Zhang, Y., Yu, T., Zheng, Z., Dai, Q., Liu, Y.: Deepmulticap: Performance capture of multiple characters using sparse multiview cameras. In: ICCV (2021)
78. Zheng, Z., Yu, T., Dai, Q., Liu, Y.: Deep implicit templates for 3d shape representation. In: CVPR. pp. 1429–1439 (2021)
79. Zheng, Z., Yu, T., Li, H., Guo, K., Dai, Q., Fang, L., Liu, Y.: Hybridfusion: Real-time performance capture using a single depth sensor and sparse imus. In: ECCV. pp. 389–406. SPRINGER, Munich (Sept 2018)

- 80. Zheng, Z., Yu, T., Liu, Y., Dai, Q.: Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE T-PAMI* (2021)
- 81. Zheng, Z., Yu, T., Wei, Y., Dai, Q., Liu, Y.: Deephuman: 3d human reconstruction from a single image. In: *ICCV*. pp. 7739–7749 (2019)
- 82. Zhi, T., Lassner, C., Tung, T., Stoll, C., Narasimhan, S.G., Vo, M.: Texmesh: Reconstructing detailed human texture and geometry from rgb-d video. In: *ECCV* (2020)
- 83. Zhu, H., Zuo, X., Wang, S., Cao, X., Yang, R.: Detailed human shape estimation from a single image by hierarchical mesh deformation. In: *CVPR*. pp. 4491–4500 (2019)
- 84. Zollhöfer, M., Nießner, M., Izadi, S., Rehmann, C., Zach, C., Fisher, M., Wu, C., Fitzgibbon, A., Loop, C., Theobalt, C., et al.: Real-time non-rigid reconstruction using an rgb-d camera. *TOG* **33**(4), 1–12 (2014)