

# StereoPIFu: Depth Aware Clothed Human Digitization via Stereo Vision

Yang Hong<sup>1</sup> Juyong Zhang<sup>1\*</sup> Boyi Jiang<sup>1</sup> Yudong Guo<sup>1</sup> Ligang Liu<sup>1</sup> Hujun Bao<sup>2</sup>

<sup>1</sup>University of Science and Technology of China <sup>2</sup>Zhejiang University

{hymath@mail., juyong@, jby1993@mail., gyd2011@mail., lgliu}@ustc.edu.cn bao@cad.zju.edu.cn

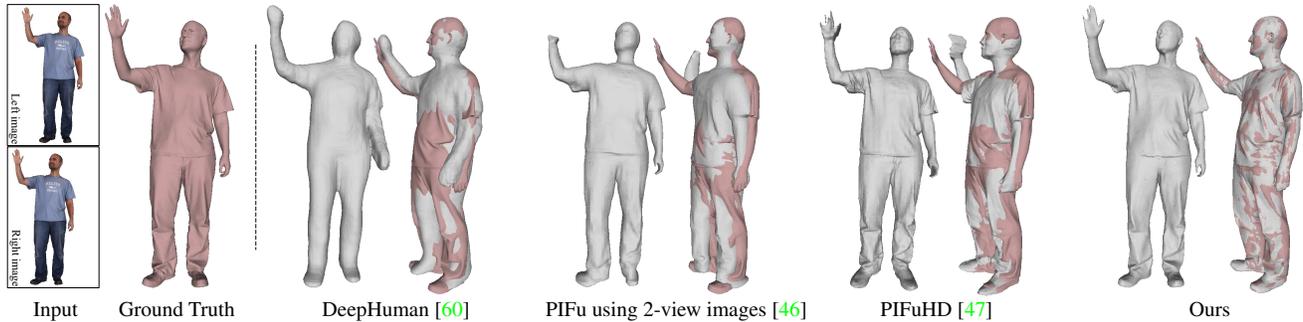


Figure 1. Our proposed StereoPIFu could recover high-fidelity and depth-aware reconstruction of clothed human body. Compared with state-of-the-art methods, our reconstruction is more accurate and maintains correct relative positions of different parts of the human body.

## Abstract

In this paper, we propose StereoPIFu, which integrates the geometric constraints of stereo vision with implicit function representation of PIFu, to recover the 3D shape of the clothed human from a pair of low-cost rectified images. First, we introduce the effective voxel-aligned features from a stereo vision-based network to enable depth-aware reconstruction. Moreover, the novel relative z-offset is employed to associate predicted high-fidelity human depth and occupancy inference, which helps restore fine-level surface details. Second, a network structure that fully utilizes the geometry information from the stereo images is designed to improve the human body reconstruction quality. Consequently, our StereoPIFu can naturally infer the human body’s spatial location in camera space and maintain the correct relative position of different parts of the human body, which enables our method to capture human performance. Compared with previous works, our StereoPIFu significantly improves the robustness, completeness, and accuracy of the clothed human reconstruction, which is demonstrated by extensive experimental results.

## 1. Introduction

Human digitization is the key to many applications like AR/VR, virtual try-on, holographic communication,

\*Corresponding author

film/game production *etc.* While high-fidelity and geometric detail preserved 3D human digitalization can be achieved with high-end acquisition equipment and well-designed capture environment [17, 22], it is not suitable for general consumers. Recently, with the popularization of consumer-level acquisition devices, human digitization with simple inputs becomes a hot research topic in the field of 3D computer vision and computer graphics.

Many methods have been proposed for clothed human reconstruction from simple inputs (e.g., single image). Among them, some methods reconstruct the 3D human body with the help of parametric models [13, 29, 33, 34, 36, 45, 27]. However, parametric models are mainly used for reconstructing naked body and can not deal with topology changes. To solve these problems, implicit function based representations have been recently introduced [25, 43, 38, 46, 47, 16, 10, 21, 23]. Representatively, Saito *et al.* [46] proposed Pixel-Aligned Implicit Function (PIFu), which performs implicit function prediction based on the  $z$ -value of a 3D query point and its projected 2D image feature. PIFu is memory-efficient and can generate a plausible surface with a single image. Later, PIFuHD [47] further improves the results of PIFu in the aspect of fine-level geometric details recovery with the aid of predicted normal maps and higher resolution. However, like other single image-based approaches, PIFu can not predict the precise spatial location and suffers from depth ambiguity, which results in inconsistent results noticeable from different perspectives. Besides, PIFu related methods are

prone to generate wrong structures like broken limbs.

A natural way to resolve the depth ambiguity is to take multi-view images as input. However, existing approaches for 3D human digitization rarely fully utilize the rich geometric relationships endowed by stereo vision. For example, Gilbert *et al.* [20] only exploits visual hull extracted from multi-view silhouettes, and PIFu-like methods [25, 46] aggregate the projected pixel features from multi-view images with a pooling layer for each query point. As shown in Fig. 1, this simple operation fails to perceive the depth information. On the other hand, some deep learning-based depth estimation methods explore geometric correlations between multi-view images and obtain satisfying results [30, 15, 31, 14, 58, 53]. The key to superior performance is the constructed cost volume feature, which encodes the correlation between a pixel and its candidate pixels in another view. Generally, these methods predict disparity, namely the horizontal displacement between a pair of corresponding pixels from rectified stereo pair images. The pixel’s predicted disparity can be obtained by calculating the weighted average of these candidate disparities, and the predicted depth can be recovered via triangulation [31, 15, 53].

To overcome the limitations of classic PIFu-like methods mentioned in the above, we propose StereoPIFu (Stereo Vision-Based Pixel-Aligned Implicit Function) for depth-aware clothed humans digitization. Compared with PIFu, StereoPIFu has two novel designs to improve its representation ability. First, we introduce additional novel voxel-aligned features as input to the implicit function. Different from the voxel-aligned features of previous methods [23, 16], our voxel-aligned features are extracted from volume data widely used in stereo vision-based depth estimation networks. Specifically, we construct two volume data from the feature maps of the image pair. The volumes define a spatial grid and for a 3D query point, we can do trilinear interpolation on its neighboring voxels in the volumes to generate voxel-aligned features. The features contain rich geometric information, indicating the correlation between the query point and the underlying surface. With voxel-aligned features as input, StereoPIFu can predict humans’ spatial location in camera space without the need to normalize human shape into a canonical space like PIFu. Second, we introduce human shape priors to guide the representation. Specifically, a high-quality human depth map is obtained based on the above-mentioned volume data. Moreover, the relative  $z$ -offset between the query point and its projected pixel’s predicted depth is added as another input of the implicit function. Compared with using the absolute  $z$ -value, relative  $z$ -offset can generate more realistic details. More importantly, the predicted depth map provides complete surface constraints for the human body and effectively eliminates broken limbs in reconstructed results. With the

above novel designs, StereoPIFu can recover high-fidelity geometric details and accurate geometric shape of human body. In summary, the paper includes the following contributions:

- We propose StereoPIFu, a novel implicit representation integrating stereo vision to PIFu representation, which makes full use of binocular images and enables high-quality depth-aware reconstruction of the clothed human body.
- We utilize the novel well-designed voxel-aligned features and predicted depth map to help occupancy inference. The geometric correlation contained in the voxel-aligned features and depth priors significantly improve the robustness, completeness, and accuracy of clothed human reconstruction.

## 2. Related Work

**Reconstruction from Single Image.** Different strategies have been proposed to recover 3D human body shape from a single image. A common way is to fit the input image by a parametric model (e.g., SMPL [36]) via 2D landmarks and silhouette [13, 29, 45, 33, 27]. However, these methods can only recover the naked body’s shape due to the limited representation ability of the parametric model. Several methods [8, 7, 6, 9] add a displacement on each vertex to represent fine-scale details. Bhatnagar *et al.* [11] and Jiang *et al.* [28] additionally regress the clothes via parametric models. These methods have improved the quality but are still unable to recover high-fidelity geometry shape.

Several methods [51, 60] take 3D CNN to regress volumetric representation of human bodies. However, they can not recover geometric details due to the large memory consumption of 3D CNN. Instead of regressing a volume with a fixed resolution, Saito *et al.* [46] propose Pixel-Aligned Implicit Function (PIFu) representation, which is memory-efficient and can predict the occupancy of any 3D point. However, it may generate incorrect body structures and struggle to recover fine-scale geometric details. To this end, PIFuHD [47] improves PIFu with additionally extracted high-resolution features and predicted normal maps. Huang *et al.* [26] and Zheng *et al.* [59] take a feature related to the parametric human model as the input of the implicit function to improve the stability of prediction. He *et al.* [23] expands the pixel-aligned feature with geometry-aligned shape features, which serves as a shape prior for the reconstruction. Gabeur *et al.* [19] convert this problem as depth prediction of the front- and back-side, and the complete shape is obtained by merging the recovered point clouds together. However, these methods suffer from depth ambiguity due to the inherent nature of single image input.

**Reconstruction from Multi-view Images.** Previous studies extract shape cues from the silhouette, stereo, and

shading, to recover geometric shape from multi-view inputs [49, 17, 52, 8, 22]. The high-quality 3D shape can be reconstructed with hundreds of cameras [17, 22, 18, 56], but this hardware configuration is inaccessible to general consumers due to its special equipment and complexity. Gilbert *et al.* [20] predicts the high-fidelity volume of the human shape utilizing a coarse visual hull generated by sparse view silhouettes with a 3D convolutional auto-encoder network. Alldieck *et al.* [8] take a video of humans slowly rotating as input and exploit the silhouette of frames to optimize the SMPL+D representation. Alldieck *et al.* [6] uses a network to substitute the optimization, which leads to faster inference and simpler input (1-8 frames). However, these methods are hard to recover fine-level surface details due to their limited geometric representation abilities. Several works [25, 46] aggregate multi-view pixel-aligned features with a pooling layer to help reconstruction. Although improved results can be obtained, the simple aggregation with the pooling strategy does not make full use of the information endowed by multi-view images.

**Depth Estimation from Stereo Images.** Depth estimation from stereo images has been studied extensively for decades. Although traditional methods have made significant progress [12, 24], they still suffer from the edge-fattening [48, 39] and poor performance in challenging situations like textureless regions. Recently, deep learning-based methods have been proposed to alleviate these problems [37, 30, 15, 31, 14, 58, 53]. DispNet [37] constructs a large-scale synthetic dataset, Scene Flow, and builds the first end-to-end trainable framework for disparity prediction. GC-Net [30] concatenates the left and right images’ feature and applies 3D convolutions to aggregate the resulting 4D feature cost volume. PSMNet [15] further improves the accuracy by using a stacked hourglass block [41] with more 3D convolutional layers. Several methods [58, 53] design novel modules to reduce the memory consumption and can obtain competitive performance with other state-of-the-art methods. However, these methods can only recover the shape of the human body’s visible part as they do not utilize any human body shape prior.

### 3. Method

In this paper, we aim to perform depth-aware reconstruction for the clothed human. To this end, we propose StereoPIFu, which takes a pair of stereo images as input. Similar to PIFu [46], we learn an implicit function that indicates whether a 3D point is inside the human body or not. Compared with PIFu, we additionally add the self-designed voxel-aligned features as input, which effectively improves our method’s expression ability and enables depth-aware reconstruction of the clothed human body. Moreover, the predicted high-fidelity depth map further helps recover the fine-level geometric details of the visible part. An overview

of our method is shown in Fig. 2, and we will give the algorithm details for each part in the following.

#### 3.1. Stereo Vision-based PIFu

StereoPIFu takes a pair of rectified color images,  $I_l$  and  $I_r$ , as input, and can be formulated as:

$$f(F_l(\pi_l(P)), \Phi(P), \Psi(P), \psi_t(Z_E(P))) \mapsto s \in [0, 1], \quad (1)$$

where  $f$  is the implicit function represented by a multi-layer perceptron (MLP). It infers a continuous scalar  $s \in [0, 1]$  to indicate the occupancy of any 3D query point  $P$  in the left camera space.  $\pi$  represents the camera projection. Its subscripts indicate its left and right image counterparts.  $F_l$  is the feature map extracted from the left image and  $F_l(\pi_l(P))$  denotes pixel-aligned feature in PIFu [46].  $\Phi(P)$  and  $\Psi(P)$  are voxel-aligned features of  $P$ , and they will be described in detail later. We denote  $E$  as the predicted depth map corresponding to  $I_l$ , and the relative  $z$ -offset  $Z_E(P)$  of  $P$  relative to  $E$  is defined as:

$$Z_E(P) = P_z - E(\pi_l(P)), \quad (2)$$

where  $P_z$  is the  $z$ -coordinate of  $P$ , and  $E(\pi_l(P))$  is the depth value of  $\pi_l(P)$  in  $E$ . Compared with PIFu [46], we additionally input more variables to predict the occupancy value  $s$  of  $P$ . We will introduce them one by one and explain the reasons for our design.

**Pixel-Aligned Feature  $F_l(\pi_l(P))$ .** The pixel-aligned feature was first introduced in PIFu [46] and widely applied in the later PIFu-like methods [47, 26, 23, 35, 59]. First, an image encoder (*e.g.*, hourglass in [46]) is used to extract feature map  $F_l$  from  $I_l$ . Generally, the image encoder is specially designed to have a large receptive field to support overall perception and consistent depth inference. For a given 3D query point  $P$ , we compute its pixel-aligned feature  $F_l(\pi_l(P))$  by bilinearly interpolating  $F_l$  at  $\pi_l(P)$ . Instead of using a global feature [43], the pixel-aligned feature encodes the local detail information contained in the image and results in a more high-fidelity reconstruction.

**Voxel-Aligned Features  $\Phi(P)$ ,  $\Psi(P)$ .** Although the pixel-aligned feature leads to more surface details, it cannot perceive the absolute depth information and struggles to guarantee correct relative positions of various body parts. In addition to the pixel-aligned feature, PIFu [46] further uses the query points’  $z$ -coordinates to distinguish their occupancy values along the ray. Therefore, it needs to normalize all training data to a fixed bounding box, making PIFu unable to restore the human body’s spatial location. To this end, we specially design the novel voxel-aligned features, which contain rich spatial and geometric information. Based on these features, both the spatial location and the relative position of various body parts can be recovered.

Our observation is that the volume data, widely used in stereo vision-based depth estimation networks, naturally

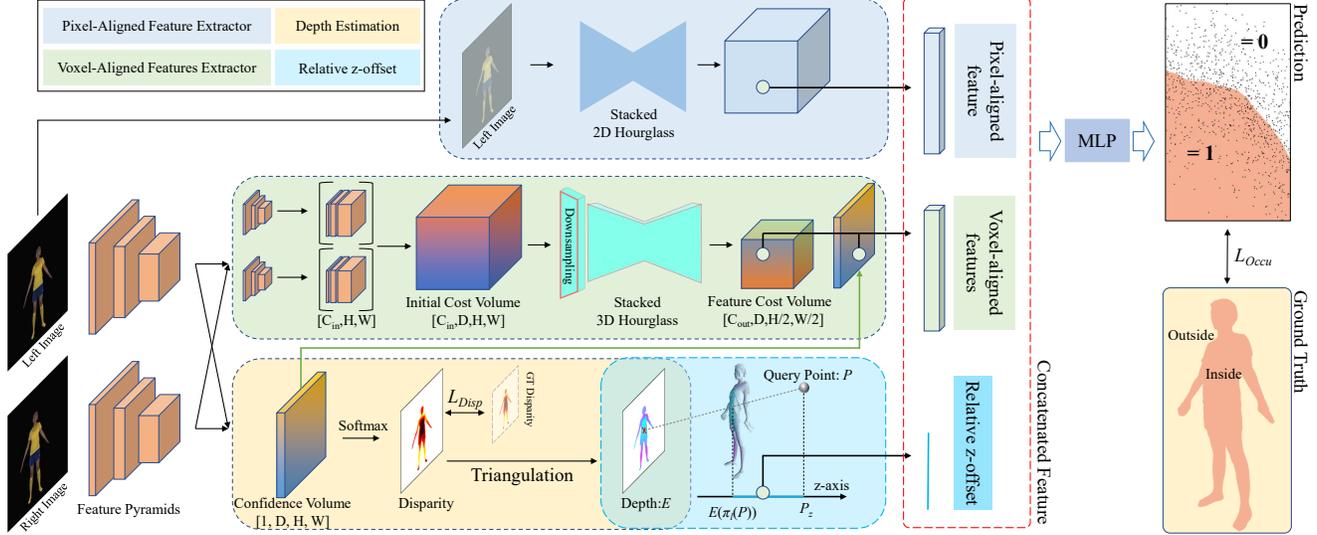


Figure 2. Overview of our StereoPIFu pipeline. Given a stereo pair, for a query point  $P$ , its pixel-aligned feature, voxel-aligned features, and relative z-offset are constructed. These features encode the information about whether  $P$  is inside the underlying surface or not and are used for inferring the occupancy of  $P$  by the MLP.

contains spatial information. For a given 3D query point  $P$ , we define its voxel-aligned features as the trilinear interpolation of these volume data depending on its spatial position. Specifically, we use AANet+, recently introduced in [53], to predict the depth map  $E$  corresponding to  $I_l$ . Its structure is shown in the lower-left part of Fig. 2. First, a shared feature extractor is used to extract the downsampled feature pyramids. Then, given the maximum disparity range  $D$ , an intermediate variable, i.e., the confidence volume  $\Psi$ , is constructed. It is a 4D tensor with size  $[1, D, H, W]$  and defines a scalar field in 3D space (please refer to [53] for details). For any element of  $\Psi$ , its index coordinate  $(d, i, j)$  corresponds to a 3D point in the space, and its value describes the probability of whether the point lies on the visible part of the underlying surface. Finally, the predicted disparity  $d^{Pred}(i, j)$  for a pixel  $(i, j)$  in the left image  $I_l$  is a combination of all the disparity values weighted by their corresponding confidence value in  $\Psi$ , and it is expressed as:

$$d^{Pred}(i, j) = \sum_{d=0}^{D-1} d \times \Psi(d, i, j). \quad (3)$$

The depth value can be computed by  $\frac{bk}{d^{Pred}(i, j)}$  [50], where  $b$  is the baseline of stereo images and  $k$  is the focal length.

As the confidence volume mainly encodes the visible part of the underlying surfaces, we introduce another feature cost volume to further encode the 3D space containing the entire underlying surface. As shown in the voxel-aligned features extractor of Fig. 2, we first upsample the multi-scale features of two feature pyramids to the same size. Then, we concatenate them together to form two 3D tensors with size  $[C_{in}, H, W]$ , denoted as  $\mathcal{F}_l$  and  $\mathcal{F}_r$ , re-

spectively. Next, we follow Khamis *et al.* [31] to construct the initial feature cost volume  $\Phi'$  with size  $[C_{in}, D, H, W]$  from  $\mathcal{F}_l$  and  $\mathcal{F}_r$ , by computing the differences between a feature in  $\mathcal{F}_l$  and its corresponding features with some disparities in  $\mathcal{F}_r$ , i.e.:

$$\Phi'(:, d, i, j) = \mathcal{F}_l(:, i, j) - \mathcal{F}_r(:, i, j - d). \quad (4)$$

Based on  $\Phi'$ , we use a variation of the stacked 3D hourglass network in [15] to aggregate context across the spatial and disparity domain to generate the feature cost volume  $\Phi$ . Specifically, we add a downsampling layer at the beginning of the network for reducing memory consumption. All 3D deconvolutions are replaced by the combination of upsampling and convolution for eliminating the checkerboard artifacts [42]. The repeated bottom-up and top-down architecture of the hourglass network extends the 3D receptive field, making features in  $\Phi$  more consistent and robust. Like the above-mentioned confidence volume  $\Psi$ ,  $\Phi$  naturally defines a vector field in 3D space, the index coordinates of each element of  $\Phi$  corresponds to a 3D point, and its feature vector encodes the information whether the 3D point lies on the underlying surface, including the occluded parts.

For a given query point  $P$ , its voxel-aligned features consist of cost feature  $\Phi(P)$  and confidence value  $\Psi(P)$ . They are obtained by trilinearly interpolating  $\Phi$  and  $\Psi$  at the coordinate  $(\pi_l(P)_x - \pi_r(P)_x, \pi_l(P)_y, \pi_l(P)_z)$ , respectively. **Relative z-offset  $Z_E(P)$ .** As stated in PIFuHD [47], the pixel-aligned feature with low resolution is hard to encode fine-level surface details. The above-mentioned voxel-aligned features also suffer a similar problem. Fortunately, AANet+ can provide a high-resolution depth map. Thus the relative z-offset naturally encodes the fine-level geometric

information. Moreover, thanks to the human shape prior from our constructed large-scale human body dataset (see Sec. 4.1), the AANet+ retrained by the dataset can generate a high-quality, complete, and consistent human depth map. Furthermore, the relative z-offset can exploit the human priors from predicted depth to guide occupancy inference. Specifically, the relative z-offset endows the query points close to the underlying surface with a similar relative z-offset value, making the network easier to train and enabling our StereoPIFu to generate fine-level surface details and avoids broken limbs. Therefore, the relative z-offset actually acts as a bridge between the predicted depth and occupancy prediction.

For some cases such as a hand in front of the torso, there will be some discontinuous regions in the predicted depth map  $E$ . In these cases,  $Z_E(P)$  of the occluded query points in the back will change discontinuously, resulting in unnatural geometric copy in the occluded region, as shown in Fig. 4. To tackle this issue, we design a transformation function  $\psi_t(x)$  to normalize  $Z_E(P)$  to the interval  $(-1.0, 1.0)$ :

$$\psi_t(x) = \frac{2.0}{1.0 + \exp(-t \cdot x)} - 1.0, \quad (5)$$

where  $t$  is a hyperparameter. For query points far away from the predicted depth,  $\psi_t(Z_E(P))$  will be close to 1.0 or  $-1.0$ , which effectively eliminates the geometric copy of the occluded part of human body and helps to produce more natural reconstruction results, see Sec. 4.3.

### 3.2. Loss Function

With our constructed dataset (see Sec. 4.1), our StereoPIFu is trained with ground truth as supervision. In the following, we give details of the loss terms.

**Loss on Depth Estimation.** We adopt the same multi-scale loss function as [53] to retrain AANet+, i.e., depth estimation module of our StereoPIFu, especially for human body type images,

$$L_{Disp} = \sum_{\omega} \lambda_{\omega} \cdot \mathcal{L}(d_{\omega}^{GT}(p), d_{\omega}^{Pred}(p)) \quad (6)$$

where  $d_{\omega}^{Pred}(p)$  is  $\omega$ -scale predicted disparity in pixel  $p$ , and  $d_{\omega}^{GT}(p)$  is the corresponding ground truth value.  $\mathcal{L}$  is the smooth L1 loss and  $\lambda_{\omega}$  is the weight to balance different terms.

**Loss on Implicit Function.** We adopt the same sampling strategy with [46, 47] to generate sampling points for training our implicit function. Differently, we do not normalize our training data to a canonical space and they are accompanied by random spatial locations (see Sec. 4.1). The Binary Cross Entropy (BCE) loss is used for occupancy prediction.

$$L_{Occu} = \sum_{P \in \mathcal{S}} f^{GT}(P) \cdot \log(f^{Pred}(P)) + (1 - f^{GT}(P)) \cdot \log(1 - f^{Pred}(P)) \quad (7)$$

where  $\mathcal{S}$  denotes the set of sampling points,  $f^{Pred}(P)$  is the predicted occupancy value for sampling point  $P$  and  $f^{GT}(P)$  is the corresponding ground truth occupancy value. The pixel-aligned feature extractor, the voxel-aligned features extractor, and the MLP of Fig. 2 are trained using  $L_{Occu}$ .

## 4. Experiments

### 4.1. Datasets

We collected 171 and 103 rigged human models from AXYZ [1] and RenderPeople [4], respectively. They all have high-fidelity geometry and realistic texture. Then non-water-tight meshes are converted into water-tight meshes with blender [2] for efficiently generating sampling points and their ground truth occupancy values. To construct a large-scale dataset for training, these human models are animated with Mixamo [3] to generate mesh sequences with various actions and postures. Finally, we generate 68419 meshes from 246 human models as the training data, and 7578 meshes from 28 human models as the testing data.

To efficiently synthesize binocular image pairs from the generated meshes, we assume the mesh surface to be predominantly Lambertian and fix the binocular camera parameters. For a given mesh in our dataset, we randomly rotate, translate, and scale the mesh within a certain range and randomly disturb the light direction and intensity to further enhance our data. The synthetic images are rendered on-the-fly with CUDA acceleration and directly used for training. This strategy effectively improves our network’s generalization ability.

### 4.2. Implementation Details

We train StereoPIFu in two steps. First we train the AANet+ for depth prediction based on Eq. (6). Then we fix the parameters of the AANet+ and train the other parts of StereoPIFu for occupancy prediction based on Eq. (7).

We implemented our StereoPIFu in Pytorch [44]. Adam optimizer [32] ( $\beta_1 = 0.9, \beta_2 = 0.999$ ) is used and the weight decay is set to 0.0001. The learning rate starts at 0.0001 and is decayed by the factor of 0.1 after every 10 epochs. We train our network on 3 NVIDIA V100 GPUs with batch size of 15. The parameter  $t$  of transformation function  $\psi_t(\cdot)$  is set to 50. The loss weights in Eq. (6) are set to  $\lambda_1 = \lambda_2 = \lambda_3 = 1.0, \lambda_4 = \frac{2}{3}, \lambda_5 = \frac{1}{3}$ . The resolution of our stereo image pair is set to  $576 \times 576$  and their maximum disparity range is set to 72. Due to the downsampling operator, the sizes of confidence volume  $\Psi$  and cost volume  $\Phi$  are  $[1, 24, 192, 192], [64, 24, 96, 96]$ , respectively.

### 4.3. Evaluations

We evaluate our StereoPIFu on three datasets, including AXYZ pose dataset [1], the synthetic meshes from our test-

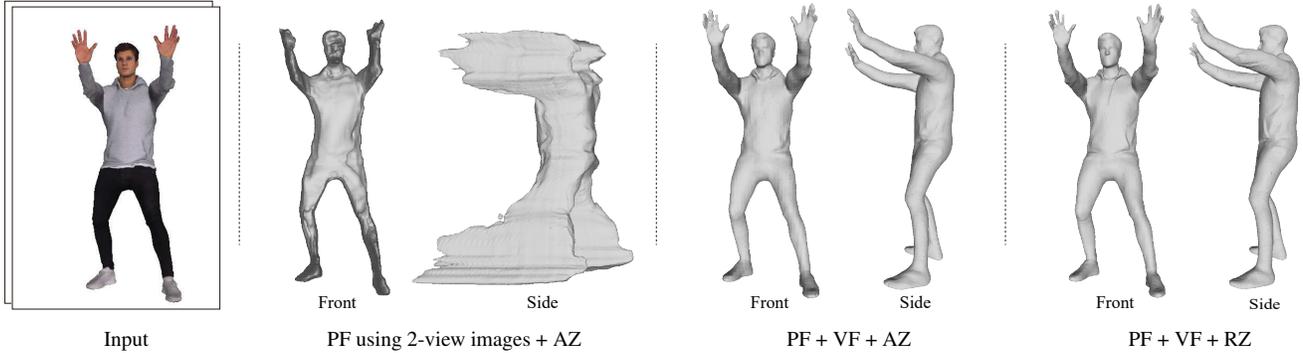


Figure 3. Ablation study on our StereoPIFu. ‘PF’: Pixel-aligned Feature; ‘PF using 2-view images’: The average of PFs from 2-view images; ‘VF’: Voxel-aligned Features; ‘AZ’: Absolute z-value; ‘RZ’: Relative z-offset. ‘PF using 2-view images + AZ’ fails to infer the plausible geometry due to the complexity and diverse spatial locations of our training data. ‘PF + VF + AZ’ is a variant of our StereoPIFu. Compared with PIFu[46], only our self-designed voxel-aligned features is added as the input of the implicit function. ‘PF + VF + RZ’ is another variant of our StereoPIFu. As relative z-offset replaces absolute z-value, finer surface details can be recovered.

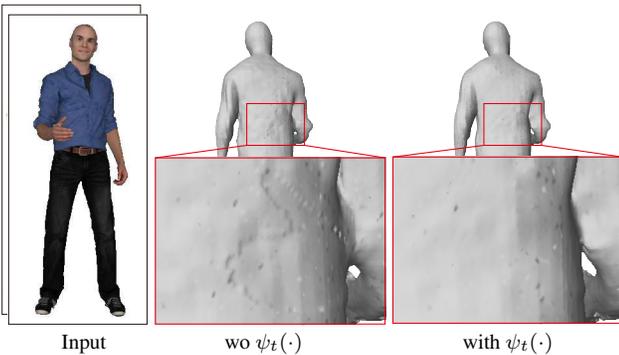


Figure 4. Ablation study on transformation function  $\psi_t(\cdot)$ . The transformation function  $\psi_t(\cdot)$  can effectively eliminate the back geometry copy artifact caused by the discontinuity of  $Z_E(P)$ .

ing data, and BUFF dataset [57]. AXYZ pose dataset contains 100 high-quality meshes with obvious cloth wrinkles under general posture. For our testing dataset, we randomly select 100 meshes that are accompanied by some challenging postures and complex occlusions. BUFF dataset contains 100 human meshes scanned by high-end equipment. The material and lighting of these meshes look realistic. These datasets have ground-truth measurements and are not used in our training stage. The corresponding color images are obtained by using our renderer. In the following, we first show some ablation studies to analyze the role of each module in our method, and then do comparisons with other state-of-the-art methods. In the final, we show the reconstruction results on real captured data to demonstrate the generalization of our method.

**Ablation Study.** We conduct ablation studies to demonstrate the importance of inputs in Eq. (1) that we design for occupancy inference and high-fidelity reconstruction. First, we retrain PIFu using 2-view images [46] with our constructed dataset, where PIFu’s network takes the same input

Method	AXYZ	BUFF	Synthetic
AANet+	0.7268	0.8969	0.7699
Retrained AANet+	0.1635	0.1785	0.1600
PF+VF+ $\psi_t$ (RZ)	0.1754	0.1907	0.1659

Table 1. Comparisons between AANet+ and StereoPIFu. The numerical values represent the average distance (cm) between ground truth geometry and the visible area generated by AANet+ predicting depth or our reconstructed result.

Method	AXYZ Pose		BUFF		Synthetic	
	P2S	Chamfer	P2S	Chamfer	P2S	Chamfer
DeepHuman	2.656	2.670	3.875	3.454	2.761	3.502
PIFu With 1-View	1.760	1.980	2.010	2.033	2.729	3.680
PIFu With 2-Views	1.739	1.971	1.975	2.013	2.749	3.706
PIFuHD	1.551	1.666	1.816	1.735	2.544	3.219
PF+VF+AZ	0.612	0.668	0.639	0.667	0.726	0.641
PF+VF+RZ	0.556	0.611	0.591	0.614	0.469	0.485
PF+VF+ $\psi_t$ (RZ)	<b>0.547</b>	<b>0.603</b>	<b>0.568</b>	<b>0.612</b>	<b>0.417</b>	<b>0.436</b>

Table 2. Quantitative evaluation on several datasets. Point-to-Surface distance and Chamfer distance (cm) between reconstructed and ground truth meshes are computed. Please refer to Fig. 3 for ‘PF,VF,AZ,RZ’.

images as our StereoPIFu. As shown in Fig. 3, due to the complexity and diverse spatial locations of our training data, PIFu [46] using the pixel-aligned feature from 2 view images and absolute  $z$ -coordinate fails to reconstruct reasonable human geometry. In contrast, a variant version of our StereoPIFu successfully learns human priors from the same dataset by taking the pixel-aligned feature, voxel-aligned features, and the absolute  $z$ -coordinate of  $P$  as input. This experiment shows that our voxel-aligned features indeed encode the query point’s depth-scale information and further enhance the representation ability of previous works.

Fig. 3 also shows that geometric details can be better recovered by replacing the absolute  $z$  value with the relative  $z$ -offset. Besides, as shown in Tab. 1, the errors of our StereoPIFu and retrained AANet+ are similar (slight error

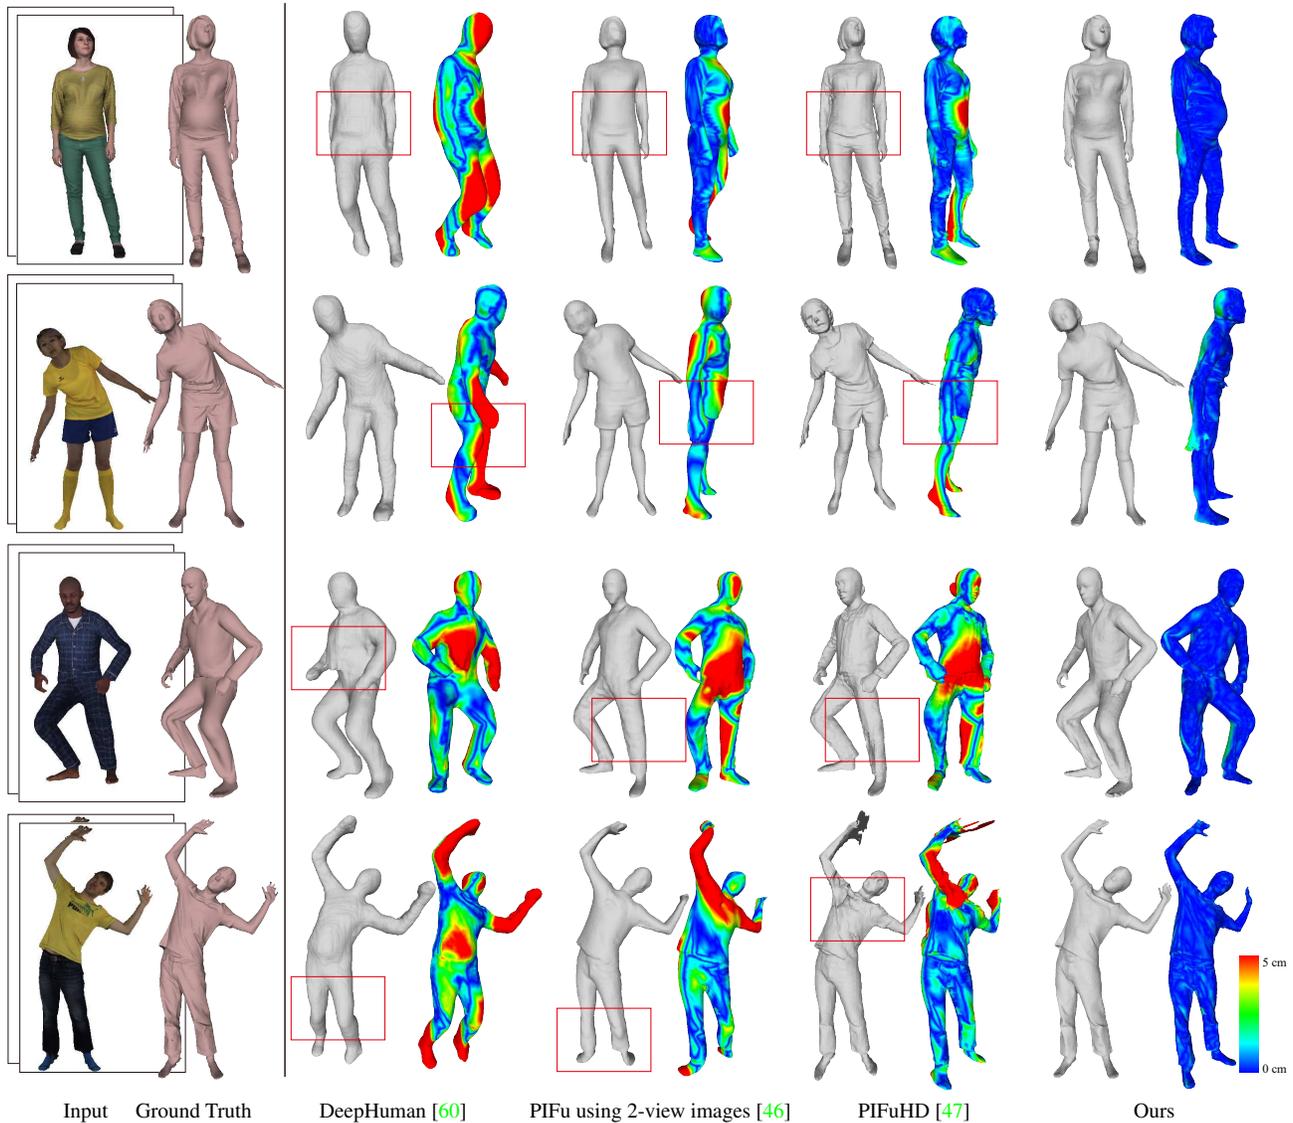


Figure 5. Qualitative comparison with other methods. Our StereoPIFu can depth-aware reconstruct the high-fidelity clothed human body and is more robust to extreme postures. The incorrect reconstruction parts of other methods are marked with red rectangles.

increase of StereoPIFu may come from the process of reconstruction from occupancy fields), which verifies that our self-designed relative  $z$ -offset indeed effectively utilizes the human priors from the predicted depth map to guide occupancy inference. In addition, the significant accuracy improvement of the retrained AANet+ demonstrates the effectiveness of our constructed dataset. Moreover, Fig. 4 shows that our transformation function  $\psi_t(\cdot)$  can help infer occluded back geometry and eliminate unreasonable artifacts.

Tab. 2 shows our quantitative evaluation using the above mentioned three types of dataset. We compute Point-to-Surface distance and Chamfer distance from the reconstructed mesh to the ground truth mesh. The results from Tab. 2 demonstrate that our novel voxel-aligned feature sig-

nificantly improves the reconstruction accuracy.

**Comparisons.** Fig. 1, 5 and Tab. 2 show our qualitative and quantitative comparison with state-of-the-art 3D human reconstruction methods. DeepHuman [60] first predicts a parametric human model, i.e., SMPL [36], then the entire volume’s occupancy value is regressed based on the predicted SMPL mesh. Its results are over-smooth due to memory limitation. Besides, incorrect body structure prediction in the first stage will result in large error in the final. PIFu’s results can recover rough shape but still suffer from lacking fine-level details. PIFuHD [47] additionally uses the feature extracted from a high resolution image to guide the fine-level reconstruction. Also, it further enhances geometric details using predicted front- and back- side normal

maps. However, its results are still depth-ambiguous and cannot ensure the correctness of different parts of the human body’s relative positions due to single image input. Tab. 2 shows the quantitative comparison between our method and these methods. Our method greatly outperforms the state-of-the-art method of PIFuHD in terms of geometric errors. On average, the Point-to-Surface distance and Chamfer distance decrease from 1.97cm to 0.51cm (74.1% reduction) and from 2.21cm to 0.55cm (75.1% reduction) respectively. It shows that the introduction of our well-designed voxel-aligned features can bring a substantial gain. It is worth noting that, compared with single image based methods, our method does not need to perform scaling operation before computing the reconstruction accuracy. For other methods, we need to align their results to ground-truth models based on rigid transformation with scaling.

Different from single image based methods, our method is depth-aware. As shown in the first row of Fig. 5, the belly of the pregnant women can be accurately recovered by our method. StereoPIFu is also aware of the bending motion and can accurately reconstruct the plausible result as shown in the second row of Fig. 5. Similarly, our method also successfully maintains the relative position of the right hand and the human body as shown in Fig. 1. As pointed out via the red rectangles, the competing methods can not handle these challenging cases. Besides, our method is more robust than other methods. For some extreme postures as shown in the bottom two rows of Fig. 5, our method can still obtain correct results and our reconstructions are significantly better than other methods.

**Results on Real Data.** We also evaluate our method on actual captured data. We use a binocular camera [5] to capture binocular images where people do various actions. As shown in Fig. 6, although the camera parameters and lighting conditions may be inconsistent with our training data, our StereoPIFu can still accurately reconstruct the geometry shape of human bodies. We also show the reconstruction results of PIFuHD [47] using a single view image. Although their results look quite good from the input view, inaccurate human body structure can be easily observed from another view. As shown in the 1st row of Fig. 6, the left leg of their result is longer than the right one. Besides, we can find that PIFuHD generates incorrect shape especially for the leg as shown in the last two rows, a significantly enlarged right leg can be observed from the side view of the third example. In contrast, with the help of binocular images and the well-designed neural network, our StereoPIFu can accurately recover the human body’s geometry shape and relative positions of different body parts. Therefore, our method can be directly extended to capture human performance, and a comparison video is supplied as supplementary material. As the video shows, our results are more stable and robust than others.

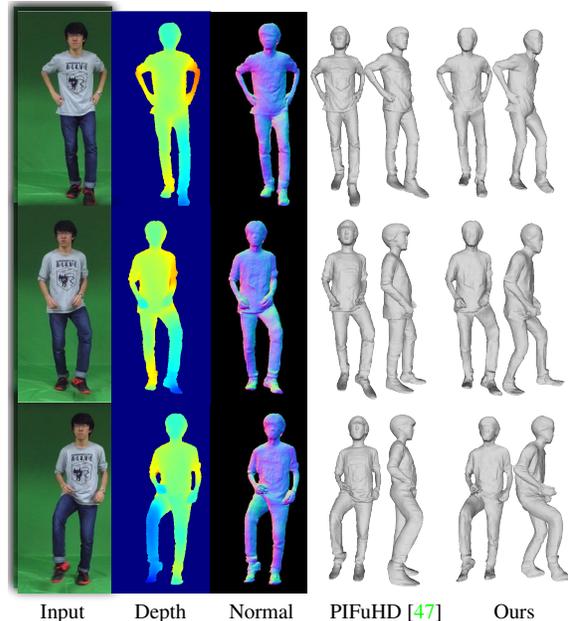


Figure 6. Comparison results of our method and PIFuHD [47] on real data. From left to right: input images, depth maps recovered by our network, normal maps computed from the recovered depth maps, results of PIFuHD, results of ours. Please refer to the supplementary video for sequence comparison.

## 5. Discussion and Conclusion

Our work still has several limitations. First, the careful calibration of the binocular camera is essential for our method. When the camera parameters are not accurate, our method may be affected or even fail to reconstruct body shapes. Second, for the invisible area, our StereoPIFu can only predict a plausible result while can not guarantee its accuracy. In the future, we plan to utilize several MVS systems [54, 55, 40] to alleviate the problem.

In this paper, we proposed StereoPIFu, a novel clothed human digitization method that integrates stereo vision to implicit function representation. First, we introduced the novel voxel-aligned features, which enables our StereoPIFu to depth-aware clothed human body reconstruction. Second, the transformed relative z-offset of the query point is used to recover the geometric details and eliminate the back region’s geometry-copy artifacts, and it further improves the reconstruction accuracy of our method. Extensive experiments demonstrate that the proposed method outperforms existing state-of-the-art methods and achieves more robust and accurate 3D human digitization.

**Acknowledgements.** This work was supported by the Youth Innovation Promotion Association CAS (No. 2018495) and the Fundamental Research Funds for the Central Universities.

## References

- [1] axyz-design. <https://secure.axyz-design.com/>, 2020. 5
- [2] blender. <https://www.blender.org/>, 2020. 5
- [3] mixamo. <https://www.mixamo.com/>, 2020. 5
- [4] renderpeople. <https://renderpeople.com/>, 2020. 5
- [5] Zed stereo camera. <https://www.stereolabs.com/zed/>, 2020. 8
- [6] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, jun 2019. 2, 3
- [7] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *2018 International Conference on 3D Vision (3DV)*, pages 98–109. IEEE, 2018. 2
- [8] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8387–8397, 2018. 2, 3
- [9] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2293–2303, 2019. 2
- [10] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Combining implicit function learning and parametric models for 3d human reconstruction. *arXiv preprint arXiv:2007.11432*, 2020. 1
- [11] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5420–5430, 2019. 2
- [12] Michael Bleyer, Christoph Rhemann, and Carsten Rother. Patchmatch stereo - stereo matching with slanted support windows. In Jesse Hoey, Stephen J. McKenna, and Emanuele Trucco, editors, *British Machine Vision Conference (BMVC)*, pages 1–11, 2011. 3
- [13] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578. Springer, 2016. 1, 2
- [14] Rohan Chabra, Julian Straub, Christopher Sweeney, Richard Newcombe, and Henry Fuchs. Stereodnet: Dilated residual stereonet. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11786–11795, 2019. 2, 3
- [15] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5410–5418, 2018. 2, 3, 4
- [16] 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*, pages 6970–6981, 2020. 1, 2
- [17] 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):1–13, 2015. 1, 3
- [18] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2009. 3
- [19] Valentin Gabeur, Jean-Sébastien Franco, Xavier Martin, Cordelia Schmid, and Gregory Rogez. Moulding humans: Non-parametric 3d human shape estimation from single images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2232–2241, 2019. 2
- [20] Andrew Gilbert, Marco Volino, John Collomosse, and Adrian Hilton. Volumetric performance capture from minimal camera viewpoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 566–581, 2018. 2, 3
- [21] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020. 1
- [22] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, et al. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics (TOG)*, 38(6):1–19, 2019. 1, 3
- [23] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. In *Annual Conference on Neural Information Processing Systems*, 2020. 1, 2, 3
- [24] Asmaa Hosni, Christoph Rhemann, Michael Bleyer, Carsten Rother, and Margrit Gelautz. Fast cost-volume filtering for visual correspondence and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(2):504–511, 2012. 3
- [25] Zeng Huang, Tianye Li, Weikai Chen, Yajie Zhao, Jun Xing, Chloe LeGendre, Linjie Luo, Chongyang Ma, and Hao Li. Deep volumetric video from very sparse multi-view performance capture. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 336–354, 2018. 1, 2, 3
- [26] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3093–3102, 2020. 2, 3
- [27] Boyi Jiang, Juyong Zhang, Jianfei Cai, and Jianmin Zheng. Disentangled human body embedding based on deep hierarchical neural network. *IEEE Transactions on Visualization and Computer Graphics*, 2020. 1, 2
- [28] Boyi Jiang, Juyong Zhang, Yang Hong, Jinhao Luo, Ligang Liu, and Hujun Bao. Bcnet: Learning body and cloth shape

- from a single image. In *European Conference on Computer Vision*. Springer, 2020. 2
- [29] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7122–7131, 2018. 1, 2
- [30] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 66–75, 2017. 2, 3
- [31] Sameh Khamis, Sean Fanello, Christoph Rhemann, Adarsh Kowdle, Julien Valentin, and Shahram Izadi. Stereonet: Guided hierarchical refinement for real-time edge-aware depth prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 573–590, 2018. 2, 3, 4
- [32] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *International Conference on Learning Representations (ICLR)*, 2015. 5
- [33] 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*, pages 5253–5263, 2020. 1, 2
- [34] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2252–2261, 2019. 1
- [35] Zhe Li, Tao Yu, Chuanyu Pan, Zerong Zheng, and Yebin Liu. Robust 3d self-portraits in seconds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1344–1353, 2020. 3
- [36] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 1, 2, 7
- [37] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016. 3
- [38] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 1
- [39] Dongbo Min, Jiangbo Lu, and Minh N Do. A revisit to cost aggregation in stereo matching: How far can we reduce its computational redundancy? In *2011 International Conference on Computer Vision*, pages 1567–1574. IEEE, 2011. 3
- [40] Zak Murez, Tarrence van As, James Bartolozzi, Ayan Sinha, Vijay Badrinarayanan, and Andrew Rabinovich. Atlas: End-to-end 3d scene reconstruction from posed images. *arXiv preprint arXiv:2003.10432*, 2020. 8
- [41] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 3
- [42] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *distill* (2016), 2016. 4
- [43] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 1, 3
- [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019. 5
- [45] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10975–10985, 2019. 1, 2
- [46] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2304–2314, 2019. 1, 2, 3, 5, 6, 7
- [47] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020. 1, 2, 3, 4, 5, 7, 8
- [48] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International journal of computer vision*, 47(1-3):7–42, 2002. 3
- [49] Peng Song, Xiaojun Wu, and Michael Yu Wang. Volumetric stereo and silhouette fusion for image-based modeling. *The Visual Computer*, 26(12):1435–1450, 2010. 3
- [50] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010. 4
- [51] Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 20–36, 2018. 2
- [52] Di Xu, Qi Duan, Jianmin Zheng, Juyong Zhang, Jianfei Cai, and Tat-Jen Cham. Shading-based surface detail recovery under general unknown illumination. *IEEE transactions on pattern analysis and machine intelligence*, 40(2):423–436, 2017. 3
- [53] Haoifei Xu and Juyong Zhang. A Janet: Adaptive aggregation network for efficient stereo matching. In *Proceedings of*

- the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1959–1968, 2020. [2](#), [3](#), [4](#), [5](#)
- [54] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018. [8](#)
- [55] Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent mvsnet for high-resolution multi-view stereo depth inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5525–5534, 2019. [8](#)
- [56] Jure Žbontar and Yann LeCun. Stereo matching by training a convolutional neural network to compare image patches. *The journal of machine learning research*, 17(1):2287–2318, 2016. [3](#)
- [57] Chao Zhang, Sergi Pujades, Michael J Black, and Gerard Pons-Moll. Detailed, accurate, human shape estimation from clothed 3d scan sequences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4191–4200, 2017. [6](#)
- [58] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 185–194, 2019. [2](#), [3](#)
- [59] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *arXiv preprint arXiv:2007.03858*, 2020. [2](#), [3](#)
- [60] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7739–7749, 2019. [1](#), [2](#), [7](#)