



significant amount of missing body parts in an image is unavoidable from a single-view perspective during multi-person interaction. Such incomplete information in 2D, in turn, affects the 3D reconstruction with missing geometry. While existing methods [5, 18, 20, 21] have utilized a 3D body model, e.g., SMPL [26], to capture a complete human, they are limited to expressing unclothed human appearance; and the reconstruction of the multi-person is often physically implausible since it involves unrealistic penetration artifacts and scene-space positional misalignment (i.e., a person’s position in world coordinates is unusually distant from others compared to their actual physical distance, as shown in Fig. 2, bottom right).

To address this challenge, we propose to use human priors for complete 3D geometry and surface contacts. For the complete geometry, we effectively utilize a 3D generative model (e.g., gDNA [6]) which is designed to generate 3D features from a latent vector, for a coarse and looking-plausible human geometry. We upgrade this 3D geometric coarse prior to 3D geometric detail prior by developing an implicit network that combines the 3D features [6] and a surface normal map. It predicts the 3D geometry augmented in its style, shape, and details (e.g., wrinkles). We finally glue the different modalities between the image and the 3D model by designing an image regression network that can predict the appearance vectors whose distribution matches the latent space of the 3D generative model [6]. For understanding the physical state of multiple people in 3D, we develop a data-driven prior for surface contact map—a probability distribution of 3D surface contacts of each body part across people. A neural network learns to predict such probability maps from the images of interacting multi-person, which provides strong guidance for accurate pose refinement in scene space. In particular, it is designed to prevent surface penetration; and to enforce accurate scene-space localization, e.g., the contact around the torso of two people in Fig. 2 provides a strong cue to decide their ordinal depth relation.

Using these two priors, we introduce a new design of the monocular multi-human capturing system. Given an image of multiple people, we first predict the detailed geometry for each person in the canonical body poses using our 3D geometric detail prior. The coarse scene-space 3D body location for each person are estimated by solving Perspective-n-Point (PnP) between the predicted 3D keypoints and a 3D body model (SMPL [26]). The multi-person poses are then refined by optimizing their global relation based on the predicted contact maps and penetration loss.

The experimental results demonstrate the strength of our method in geometric completeness with the upgraded 3D geometric detail priors; in global coherence thanks to the contact priors; and in physical plausibility with minimal penetration with the effective refinement pipeline.

Our contribution includes (1) a novel pipeline that realizes 3D reconstruction of interacting multi-person using only a single image with an effective global pose refinement; (2) the innovation to upgrade existing geometric priors that enable a complete and detailed 3D human reconstruction from the image of a partial body; and (3) an effective scene-space pose refinement framework guided by surface contact priors, which enforces the physical plausibility.

## 2. Related Work

**3D Clothed Single Person Reconstruction.** Many existing methods have explored the 3D reconstruction of human meshes from a single-view image. Some works [18, 20, 21, 33] predict SMPL parameters [26] from RGB images of a person with occlusion to reconstruct a complete human mesh. However, these works cannot express the shape beyond the naked body.

Some explicit-based approaches [1, 2, 22, 48] solved this problem by estimating the offsets at each vertex of the body mesh to represent the humans in more detail. Those explicit approaches have demonstrated sub-optimal performance when dealing with a single-view input, mainly because of the fixed 3D mesh topology [26].

To overcome the limitation of the fixed topology constrained by explicit 3D mesh models, some existing methods have utilized an implicit signed distance field (SDF) [3, 40, 45, 46] and occupancy field [7, 11–13, 15, 19, 29–31, 35, 37, 39, 47, 49]. To generate a high-resolution mesh from a coarse one in the canonical space, Neural-GIF [35] uses a backward mapping network and a displacement network. LaplacianFusion [19] employed Laplacian coordinates extracted from the mesh to generate the detailed surface. While the above methods demonstrate promising reconstruction results, those methods often suffer from missing information in 3D when the input images of humans involve occlusion.

**3D Clothed Multi-Person Reconstruction.** The multi-person 3D reconstruction is often considered as a special case of a single person with occlusion where several methods [5, 8, 10, 16, 17, 33, 34, 42–44] separately handle each individual from the scenes of human interaction. However, the focus of these methods is lying on the reconstruction of unclothed human models (e.g., SMPL [26]) that cannot represent the details of human geometry such as clothing, hair, and shoes.

While DeepMultiCap [45] utilizes the attention module and temporal fusion to capture high-quality 3D models of interacting people in clothing, they require multiview images for the best quality of 3D reconstruction in scene space.

Mustafa et al. [28] introduced an implicit approach for 3D reconstruction of multiple people with 6DOF spatial location estimation to accurately estimate the position of peo-

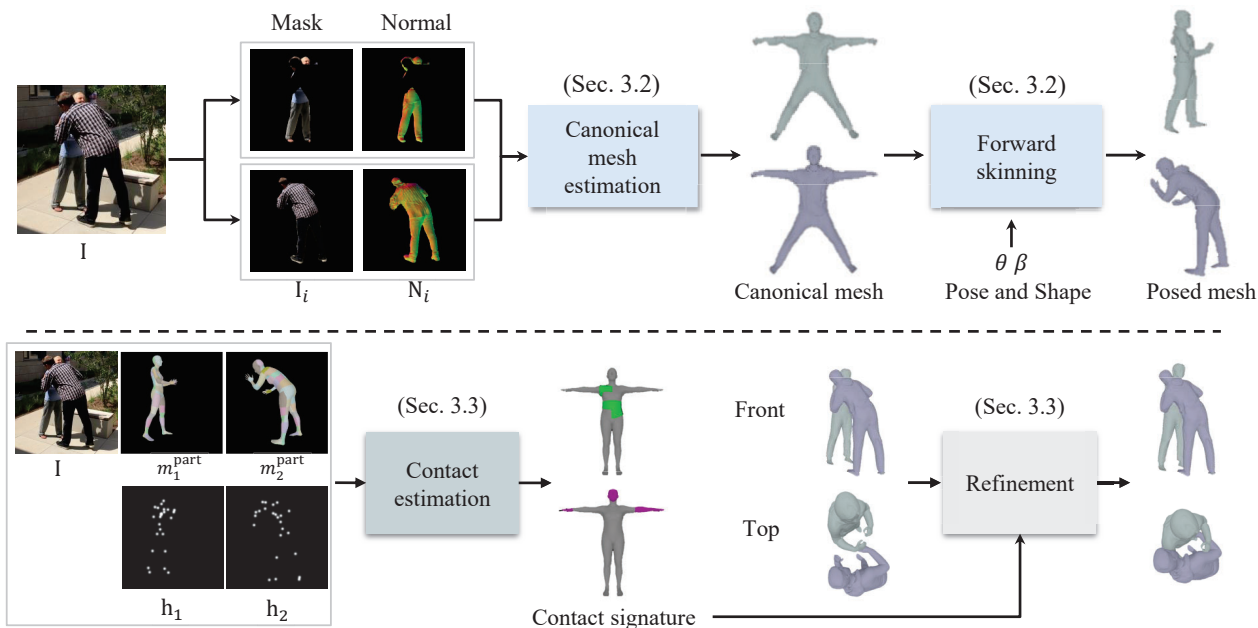


Figure 2. **System overview.** Given an image of interacting people, we aim to reconstruct multi-person geometry. Our pipeline is composed of three stages. (Top): In the generation stage, we extract coarse meshes for the multi-person from the input image. (Bottom): In the contact estimation stage, we detect the region of contact between individuals in the image. Finally, in the contact-based refinement stage, we generate detailed multi-person geometry by leveraging the information obtained from the previous two stages.

ple from global scene space. However, the results are limited to expressing coarse geometry and the method often fails when it faces people with heavy occlusion.

**Surface Contact Estimation.** Surface contacts provide a strong cue to understand the physical state between two different surfaces. Thus, there has been a growing interest in modeling such surface contacts in the process of 3D human reconstruction for human-to-object [4, 38], human-to-scene [14, 32], and human-to-human [9]. Fieraru et al. [9] introduced a representation of part-based contact signatures at a 3D body model [26] to predict accurate 3D body poses from the scenes of multiple people. However, the above methods do not consider the surface geometry of clothed humans and are brittle under severe occlusions.

### 3. Method

The overall pipeline is illustrated in Fig. 2. Our method consists of 3 main stages which will be explained in detail in the remaining sections.

#### 3.1. Pre-processing

**Human Segmentation.** We segment each human instance from an image  $I$  by using an existing method [24]. From this, we could obtain the  $i$ -th person’s segmentation mask  $\mathbf{m}_i$ .

**Normal Map.** We use an off-the-shelf normal detection method [30] to estimate the normal map  $\mathbf{N}_i$  from  $I_i$ . While

it generates front and back normal maps from  $I_i$ , we used only the front normal map.

**SMPL Parameters.** We first initialize the 2D and 3D joints ( $\mathbf{J}^{2D}$ ,  $\mathbf{J}^{3D}$ ) from a single image by using an existing monocular keypoints prediction method [5]. The inverse kinematic method [23] is applied to calculate the SMPL [26] pose parameter  $\theta$  and shape parameter  $\beta$  from the estimated 3D joints  $\mathbf{J}^{3D}$ .

**Camera Parameter and 3D Translation.** We estimate camera parameters  $\mathbf{C}_{\text{param}}$ , i.e., scale  $s$ , x-translation  $t_x$ , y-translation  $t_y$ , via the singular value decomposition (SVD) method to align SMPL mesh  $\mathbf{M}_{\text{smpl}}$  from the image space:

$$\mathbf{J}_{\text{smpl}}^{2D} \begin{bmatrix} s \\ t_x \\ t_y \end{bmatrix} = \mathbf{J}^{2D}. \quad (1)$$

The SMPL 2D joints  $\mathbf{J}_{\text{smpl}}^{2D}$  are obtained from the SMPL 3D joints  $\mathbf{J}_{\text{smpl}}^{3D}$  based on the orthogonal projection and  $\mathbf{J}_{\text{smpl}}^{3D}$  are differentially obtained from the mesh  $\mathbf{M}_{\text{smpl}}$  using the SMPL skeletal regressor provided by [26]. Additionally, we calculate 3D location  $\mathbf{T}^{3D}$  with  $\mathbf{C}_{\text{param}}$  and  $\mathbf{J}_{\text{smpl}}^{3D}$  similar to [33, 34].

**Semantic Part Segmentation and 2D Keypoint Heatmap.** We generate a semantic part segmentation mask  $\mathbf{m}_i^{\text{part}}$  by differentially rendering 75 regions of SMPL 3D mesh  $\mathbf{M}_{\text{smpl}}(\theta, \beta)$  with different colors. We generate 2D keypoint heatmap  $\mathbf{h}_i$  from 2D joints  $\mathbf{J}^{2D}$ , by drawing  $\mathbf{J}^{2D}$  Gaussian in the positions of 2D joints.

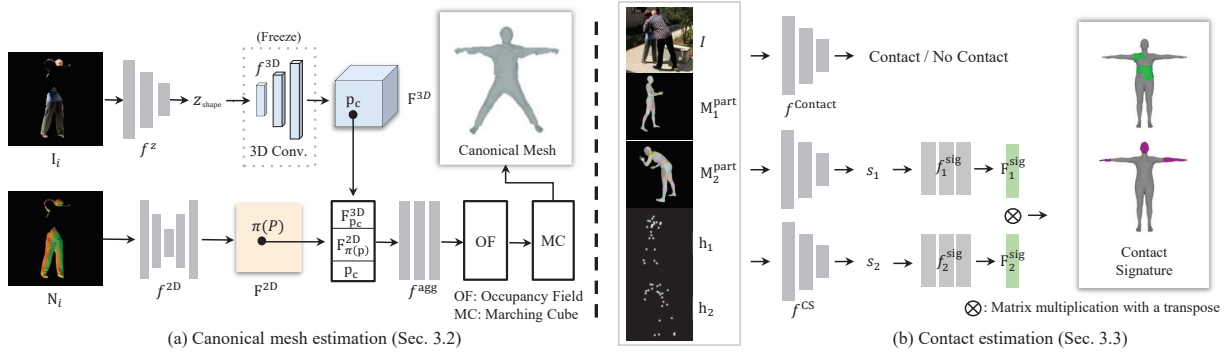


Figure 3. The details of the pipeline for each module in our system.

### 3.2. Detailed Clothed Human Generation

In this section, our goal is to generate the human surface, reflecting details such as clothes and hair. As the direct reconstruction of holistic 3D meshes of detailed humans is non-trivial, we divide the skinning methods into two disjoint steps: First, we estimate the detailed 3D human mesh in the canonical space. Second, we transform the canonical 3D mesh into the posed mesh  $\mathbf{M}$  corresponding to SMPL [26] pose  $\theta$  and shape  $\beta$  parameters using the forward skinning module. In the remainder of this sub-section, we detailed each sub-module.

**Canonical Mesh Estimation Module.** Our canonical mesh estimation module is designed to generate the canonical mesh  $\tilde{\mathbf{M}}$ . This module is composed of four sub-networks: a latent feature extractor  $f^z$ , a 3D voxel feature generator  $f^{3D}$ , a 2D image-aligned feature extractor  $f^{2D}$  and an aggregation network  $f^{agg}$  as shown in the left side of Fig. 3.

The latent feature extractor  $f^z$  extracts the latent vector  $z_{shape}$  that reflects global feature of human styles from an input image  $\mathbf{I}$ . The 3D voxel feature generator  $f^{3D}$  converts it into the feature vector  $\mathbf{F}^{3D}$ . The 2D image-aligned feature extractor  $f^{2D}$  extracts the 2D image-aligned feature  $\mathbf{F}^{2D}$  from a normal map  $\mathbf{N}$ .

For any 3D point  $\mathbf{p}$ , we transform it to  $\mathbf{p}_c$  in canonical space [6]. We obtain the associated canonical features  $\mathbf{F}_{\mathbf{p}_c}^{3D}$  by tri-linear interpolation to the  $\mathbf{F}^{3D}$  according to  $\mathbf{p}_c$ .  $\pi(\mathbf{p})$  is the perspective camera projection from the 3D point  $\mathbf{p}$  to the image plane and the local 2D feature  $\mathbf{F}_{\pi(\mathbf{p})}^{2D}$  is obtained from the  $\mathbf{F}^{2D}$  by bi-linear interpolation at the projected pixel coordinates  $\pi(\mathbf{p})$ . Then, we concatenate  $\mathbf{F}_{\mathbf{p}_c}^{3D}$ ,  $\mathbf{F}_{\pi(\mathbf{p})}^{2D}$  and  $\mathbf{p}_c$  to make the geometry- and image-aligned feature. Finally, it is fed to the aggregation network  $f^{agg}$  to obtain the high fidelity occupancy field  $o(\mathbf{p})$ . Similar to [29], the occupancy field  $o(\mathbf{p}) \in [0, 1]$  indicates whether  $\mathbf{p}$  is inside mesh surface or not. Using the Marching\_Cube algorithm [27], we extract the final canonical mesh  $\tilde{\mathbf{M}}$  from obtained occupancy fields  $\{o(\mathbf{p})\}$ .

We initialize our  $f^{3D}$  using the pre-trained model in [6], while further training  $f^z$ ,  $f^{2D}$  and  $f^{agg}$  to upgrade their

method in a way that (1) infers the style code from the 2D images; (2) extracts geometric details from the normal map  $\mathbf{N}$ ; and (3) combines information from both normal map and RGB images.

To train  $f^z$ , we use the mean squared error (MSE) loss as follows:

$$L_z(f^z) = \|z_{shape} - z_{shape}^{gDNA}\|_2^2, \quad (2)$$

where  $z_{shape}^{gDNA}$  denotes the fitted latent code by gDNA [6]. To train  $f^{2D}$  and  $f^{agg}$ , we use the binary cross entropy (BCE) loss [47],

$$L_o = \sum_{\mathbf{p} \in S} \gamma \cdot o^{GT}(\mathbf{p}) \cdot \log(o(\mathbf{p})) + (1 - \gamma)(1 - o^{GT}(\mathbf{p})) \log(1 - o(\mathbf{p})), \quad (3)$$

where  $S$  is the set of sampled points and  $\gamma$  denotes the ratio of points outside the surface in  $S$ , which helps to train the network more stably.  $o^{GT}(\mathbf{p})$  is the ground-truth occupancy at the sampled point  $\mathbf{p}$ .

**Forward Skinning.** We use the pre-trained network [6] to deform canonical mesh  $\tilde{\mathbf{M}}$  towards the posed mesh  $\mathbf{M}$  with SMPL [26] pose  $\theta$  and shape  $\beta$  parameters. It generates matrix  $W \in \mathbb{R}^{N_v \times 24}$  that transforms  $N_v$  canonical mesh vertices into skinning weight fields for 24 bones, conditioned on  $z_{shape}$  and  $\beta$ . As a result, given target  $\theta$  and  $W$ , canonical mesh vertices  $v$  on  $\tilde{\mathbf{M}}$  can be transformed to the posed space  $v'$  on the posed mesh  $\mathbf{M}$ .

### 3.3. Contact Map-based Refinement

**Contact Estimation Stage.** To predict contact estimation from images, we design a contact estimation network that takes an image, semantic part segmentation and 2D key-point heatmap as inputs; and outputs contact signatures [9], i.e., a representation that indicates whether there exist surface contacts between the body parts of two people from the 3D mesh model, as shown in Fig. 2. We concatenate  $\mathbf{I}$ ,  $\mathbf{m}_1^{part}$ ,  $\mathbf{m}_2^{part}$ ,  $\mathbf{h}_1$ , and  $\mathbf{h}_2$  to feed them to the contact discriminator  $f^{Contact}$  and the contact segmentation estimator  $f^{CS}$ .

The contact discriminator  $f^{\text{Contact}}$  is designed to classify whether there exist contacts between two interacting persons. The contact segmentation estimator  $f^{\text{CS}}$  is designed to predict the contact segmentation  $\mathbf{s}_1 \in [0, 1]^{75 \times 1}$  and  $\mathbf{s}_2 \in [0, 1]^{75 \times 1}$  for two persons. These are 75-dimensional vectors whose values indicate whether each region out of the overall 75 regions is contacted or not. These two intermediate vectors are further fed to MLP layers  $f^{\text{sig}}$  whose output vectors are used to infer the final contact signature  $\mathbf{C} \in [0, 1]^{75 \times 75}$  based on their dot products followed by a nonlinearity, i.e., sigmoid.

The overview of the network for this stage is shown on the right side of Fig. 3. To better distinguish the contact signature features of two interacting people, we employ the two-path architecture for each person and jointly estimating segmentation and signature is synergistic as described in [9].

To train  $f^{\text{Contact}}$ ,  $f^{\text{CS}}$  and  $f^{\text{sig}}$ , we use the balanced binary cross entropy (BBCE) loss as follows:

$$L_{BBCE} = -(B\bar{y} \log(y) + (1 - \bar{y}) \log(1 - y)) \quad (4)$$

where  $\bar{y}$ ,  $y$  and  $B$  denote the ground-truth, estimated value and balance term.

**Refinement Stage.** We further refine the scene-space 3D body poses considering people’s contacts and their geometry surface. To this end, we select two meshes  $\{\mathbf{M}_i, \mathbf{M}_j\}$  where  $i \neq j$ , and refine the SMPL pose  $\theta$ , shape  $\beta$  and 3D location  $\mathbf{T}^{3D}$  of two meshes  $\mathbf{M}$ . Note that, if  $f^{\text{Contact}}$  predicts that there is “no contact”, we skip this stage. We optimize  $\mathbf{T}^{3D}$  of the meshes using the contact loss  $L_{\text{contact}}$  proposed in [9] as the following formula:

$$\mathbf{T}^{3D*} = \arg \min_{\mathbf{T}^{3D}} L_{\text{contact}}(\mathbf{T}^{3D}) \quad (5)$$

where  $L_{\text{contact}} = \sum_{r_1=1}^{75} \sum_{r_2=1}^{75} \mathbf{C}(r_1, r_2) D(r_1, r_2)$ .  $r_1$  and  $r_2$  are 75 regions of each person.  $\mathbf{C}$  is the estimated contact signature between two regions of interacting persons. It represents 1 if there is contact between  $r_1$  and  $r_2$ , otherwise 0.  $D(r_1, r_2)$  denotes the distance between two regions which is formulated as:

$$D(r_1, r_2) = \sum_{f_1 \in r_1} \min_{f_2 \in r_2} d(f_1, f_2) + \sum_{f_2 \in r_2} \min_{f_1 \in r_1} d(f_1, f_2) \quad (6)$$

where  $f_1$  and  $f_2$  are the faces which belong to two regions  $r_1$  and  $r_2$ , respectively.  $d$  is the Euclidean distance between the centers of the two faces  $f_1$  and  $f_2$ .

Finally, we globally refine SMPL pose ( $\theta$ ) and shape ( $\beta$ ) parameters as well as  $\mathbf{T}^{3D}$  via the contact loss  $L_{\text{contact}}$  combined with the penetration loss  $L_{\text{penet}}$  [25], Gaussian mixture model prior loss  $L_{\text{gmm}}$  [21] and regularizer  $L_{\text{reg}}$  as follows:

$$\theta^*, \beta^*, \mathbf{T}^{3D*} = \arg \min_{\theta, \beta, \mathbf{T}^{3D}} L_{\text{opt}}(\theta, \beta, \mathbf{T}^{3D}) \quad (7)$$

where

$$\begin{aligned} L_{\text{opt}}(\theta, \beta, \mathbf{T}^{3D}) &= \lambda_c L_{\text{contact}}(\theta, \beta, \mathbf{T}^{3D}) \\ &+ \lambda_p L_{\text{penet}}(\theta, \beta, \mathbf{T}^{3D}) \\ &+ \lambda_g L_{\text{gmm}}(\theta) \\ &+ \lambda_r L_{\text{reg}}(\theta, \beta) \end{aligned} \quad (8)$$

where  $L_{\text{penet}} = \sum_{v \in V_{1,\text{penet}}} \text{SDF}_2(v) + \sum_{v \in V_{2,\text{penet}}} \text{SDF}_1(v)$  prevents collisions between two people’s meshes. The first term of  $L_{\text{penet}}$  means the sum of signed distance fields ( $\text{SDF}_2$ ) that measures the distance between the second person’s mesh surface and the first person’s mesh vertices  $V_{1,\text{penet}}$ , which are penetrating the surface of the second person’s mesh. The second term of  $L_{\text{penet}}$  represents the opposite situation.  $\text{SDF}_1$  denotes the signed distance fields of the first person’s mesh.  $V_2$  is the second human vertex set inside the first person’s mesh. Gaussian mixture model prior loss  $L_{\text{gmm}}$  reduces the probability of predicting the pose that deviates from distribution and regularizer  $L_{\text{reg}} = \|\theta - \theta_{\text{init}}\|_2^2 + \|\beta - \beta_{\text{init}}\|_2^2$  ensures the refined SMPL pose and shape parameters do not diverge much from the initial parameters  $\{\theta_{\text{init}}, \beta_{\text{init}}\}$ .  $\lambda_c, \lambda_p, \lambda_r, \lambda_g$  balances the weights over the losses.

**Updating the Posed Mesh.** After obtaining the refined parameters  $\theta^*, \beta^*, \mathbf{T}^{3D*}$ , we re-obtain our posed mesh  $\mathbf{M}^*$  using the forward skinning with new SMPL parameter  $\beta^*$  and  $\theta^*$ . We also translate the re-posed mesh  $\mathbf{M}^*$  to scene space by adding the  $\mathbf{T}^{3D*}$  to the posed mesh.

## 4. Experiments

### 4.1. Datasets

**Generation.** We conduct our experiments using THuman2.0 [41], MultiHuman [45], and 3DPW [36] datasets. THuman2.0 [41] is publicly available, and provides 500 high-fidelity 3D human scans with different clothing styles and poses. The 3D scans have the corresponding texture map and SMPL parameters. We use this dataset to train our  $f^z$ ,  $f^{2D}$ , and  $f^{\text{agg}}$ . MultiHuman [45] consists of 150 high-quality 3D human scans with the corresponding texture map. We use it for evaluation. It provides 3D scans with occlusion by humans or objects. Each sample contains 1 to 3 people and each person wears different clothing. This dataset is divided into 5 categories by the level of occlusions and the number of people. In addition, we use 3DPW [36] for evaluation. It provides SMPL [26] pose parameters and clothed mesh.

**Contact estimation.** We use FlickrCI3D [9] to train our contact estimators. It provides the paired data of images and the labels of contact validity and its signature. Among them, we filter out noisy ground truth and select 5K paired data for training.

Categories Methods	single		occluded single		two natural-inter		two closely-inter		three	
	CD	P2S	CD	P2S	CD	P2S	CD	P2S	CD	P2S
PIFu [29]	2.3240	2.5473	3.3008	3.1365	3.6404	3.5573	4.0191	4.2739	3.3058	3.4899
PIFuHD [30]	1.9745	2.0258	4.2392	4.5324	3.3252	3.0838	3.3901	3.3945	3.2487	3.1508
ICON [40]	3.6610	4.6765	4.6721	5.0283	5.0958	6.1449	5.0285	6.2282	4.1542	5.0633
DMC <sup>†</sup> [45]	2.5878	4.0971	3.1396	4.6999	2.5335	4.3184	3.2481	5.3352	3.8079	5.7033
Ours <sup>†</sup> w/o N	1.4069	1.9364	1.6574	<b>1.8731</b>	1.1996	1.9692	1.8038	2.7844	1.7642	2.7325
Ours <sup>†</sup>	<b>1.3730</b>	<b>1.8441</b>	<b>1.6329</b>	1.8898	1.1933	1.9578	1.8382	2.8614	1.7608	2.7112
Ours <sup>†</sup> w R.	-	-	-	-	<b>1.1723</b>	<b>1.9065</b>	<b>1.3358</b>	<b>2.0669</b>	<b>1.5399</b>	<b>2.2636</b>

Table 1. The comparison of our method to the state-of-the-art approaches on MultiHuman [45]. <sup>†</sup> denotes that model uses the estimated SMPL parameters from [5]. Ours w/o N denotes the results without using a normal map input N. Ours is without refinement based on the contact signature. Ours w R. is our full method. Refinement is not available to “single” and “single occluded” because they have only one person.

Methods	Office Call		Hug		ShakeHands	
	CD	P2S	CD	P2S	CD	P2S
PIFu [29]	5.26	3.61	5.77	3.81	4.60	2.80
PIFuHD [30]	4.93	3.19	6.30	4.43	4.30	3.10
ICON [40]	4.80	4.07	5.45	4.24	4.05	3.14
DMC <sup>†</sup> [45]	4.30	3.68	6.67	6.30	4.36	3.95
Ours <sup>†</sup> w R.	<b>2.16</b>	<b>1.79</b>	<b>3.22</b>	<b>2.84</b>	<b>2.93</b>	<b>2.60</b>

Table 2. We compare our proposed method with state-of-the-art methods on 3DPW [36]. <sup>†</sup> denotes that model uses the estimated SMPL parameters from [5].

## 4.2. Metrics

Following the measuring protocols from existing works [29, 45], we use two metrics to measure the reconstruction quality.

**Chamfer Distance (CD).** It measures the bidirectional point distance between the predicted and ground truth meshes based on their closest-set Euclidean distance where the smaller Chamfer distance means more accurate reconstruction results.

**Point to Surface (P2S).** It measures the directional surface distance from the prediction to the ground truth based on their closest-set Euclidean distance. The smaller number is the better.

## 4.3. Baselines

We compare our model with other state-of-the-art methods [29, 30, 40, 47]. PIFu [29] and PIFuHD [30] use the pixel-aligned implicit function so that they can generate the pixel-aligned clothed human mesh. ICON [40] uses the front and back normal map and fuses their features with a signed distance field to generate the detailed clothed human mesh. However, these methods can not reconstruct invisible parts of humans occluded by objects or other people. DeepMultiCap [45] fuses multi-view features by the proposed attention module. While it requires multi-view images as inputs for the best performance, we use a single image for the fair comparison. For the quantitative evaluation of PIFu, PI-

FuHD, and ICON, we manually transform their results from the image to the scene space using the ground-truth meshes: this means that the errors from 3D scene-space pose estimation are factored out from their methods. For qualitative results in Fig 7, we visualize their methods without such advantages since predicting the scene-space position of humans is not a part of their algorithms.

## 4.4. Results

The comparison results of SOTA and our method are shown in Tab. 1 and Fig 7. Overall, our method achieves state-of-the-art performance on MultiHuman [45] where in Tab. 1. Based on Tab. 1 and Fig 7, pixel-aligned methods [29, 30, 40] show the weak reconstruction results for the occluded body parts including a significant amount of artifacts such as missing information. While DeepMultiCap [45] produces a mesh that is less noisy compared to other baselines, its geometry reconstruction results are still coarse, and its performance is often affected by the quality of global coordinates estimation of SMPL meshes. Unlike other approaches, our method is able to reconstruct clothed meshes from highly occluded people. Thanks to the global pose refinement with contact signatures, our method is robust to the occlusion, and the active utilization of the upgraded human geometry prior enables the complete and detailed reconstruction of multiple people under occlusion. In addition, Tab. 2 shows the comparison on in-the-wild data [36]. It highlights that our method is generalizable and is able to reconstruct human meshes with details even from unconstrained data. More qualitative results are presented in Supple.

## 4.5. Ablation Study

We conduct an ablation study to confirm the effect of our module. In Tab. 1, we study the performance of our implicit model  $f^{\text{agg}}$  without a normal map input (Ours w/o N); and the 3D reconstruction results with the refined body pose and shape (Ours w R.). Refinement with surface contact priors

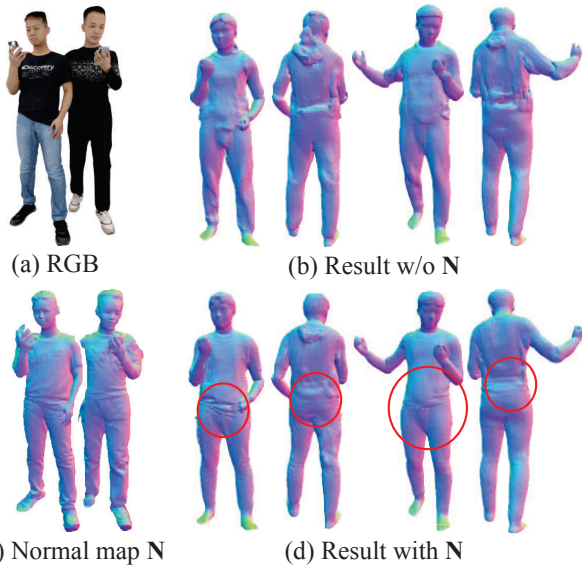


Figure 4. Ablation study on normal map  $N$ . (a) and (c) are the input RGB image and normal map input, respectively. (b) and (d) show the result without and with normal map, respectively. Normal map  $N$  provides more detailed surface information for reconstructing mesh.




Figure 5. Ablation study on the elements of  $L_{opt}$ . First row shows the results in front view. Second row shows the results in top view. (a) is RGB image input. (b) shows initial posed meshes. (c) shows the results without using  $L_{penet}$ . (d) shows the results without using pose and shape prior loss,  $L_{reg}$  and  $L_{gmm}$ . (e) shows the results of our full method.

highly improves the performance from the scenes of multi-person (two or three people) under occlusion. Although utilizing a normal map  $N$  slightly improves the performance, it can help to express the details of the reconstructed surface as shown in Fig 4.

We further study the effect of the loss used in our refinement pipeline in  $L_{opt}$ . We use the estimated SMPL parameters as initial SMPL parameters from [5]. As shown in Fig 5,  $L_{penet}$  prevents penetration between two people.  $L_{gmm}$  and  $L_{reg}$  prevent divergence of SMPL [26] parameters. Our full  $L_{opt}$  refines the SMPL parameters and 3D location more realistically based on contact signature.

We perform an ablation study in Fig. 6 where the reconstruction error (i.e. chamfer distance) is highly suppressed by our refinement module with surface contact and interpenetration priors. Our refinement module is robust to pose




Figure 6. The reconstruction accuracy (CD) with noisy and refined SMPL parameters. The noisy pose is synthetically generated by adding the fraction of Gaussian noise which perturbs ground-truth SMPL parameters.

noise. In the supplemental, we include more ablative and comparative studies.

## 5. Conclusion

We introduce a novel system for scene-space 3D reconstruction of interacting multi-person in clothing from a single image. We address the core challenge of occlusion by utilizing the human priors for complete 3D human geometry and surface contacts. For the geometry, we upgrade the existing generative 3D features by newly designing an implicit network that combines these features and a surface normal map to produce fine-detailed and complete 3D geometry. For the surface contacts, we design an image-based detector that signals the surface contact information between people in 3D. Using these priors, we globally refine 3D body poses to reconstruct accurate and penetration-free human models in scene space. The comparative evaluation and ablation study demonstrate that our method has strong and accurate performance with detailed geometry reconstruction even under heavy occlusion.

**Limitation.** When the occlusion is seriously severe, e.g., only a head part is visible, our method will not be working well due to the complete failure of the initial 3D pose estimation. The geometric diversity is not completely reflective of *in-the-wild* distribution due to the fundamental domain constraints of 3D human geometric prior, e.g., gDNA [6]. In our future work, we would like to explore the modeling of a better 3D geometric prior and its adaptation to the *in-the-wild* environment.

**Acknowledgements.** This work was supported by IITP grants (No. 2021-0-01778 Development of human image synthesis and discrimination technology below the perceptual threshold 10%; No. 2020-0-01336 Artificial intelligence graduate school program (UNIST) 10%; No. 2021-0-02068 Artificial intelligence innovation hub 10%; No. 2022-0-00264 Comprehensive video understanding and generation with knowledge-based deep logic neural network 20%) and the NRF grant (No. RS-2023-00252630 10%), all funded by the Korean government (MSIT). This work was also supported by Korea Institute of Marine Science & Technology Promotion(KIMST) funded by Ministry of Oceans and Fisheries (RS-2022-KS221674) 20% and received support from AI Center, CJ Corporation. (20%).




Figure 7. Comparison of qualitative result with state-of-the-art methods, including (b) PIFu [29], (c) PIFuHD [30], (d) ICON [40], (e) DMC [45]. DMC uses the estimated SMPL parameters from [5] and our method uses refined SMPL parameters. Our method generates complete human and more detailed human mesh than other methods.



## References

- [1] 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 *CVPR*, 2019.
- [2] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *ICCV*, 2019.
- [3] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *CVPR*, 2022.
- [4] Bharat Lal Bhatnagar, Xianghui Xie, Ilya A Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Behave: Dataset and method for tracking human object interactions. In *CVPR*, 2022.
- [5] Junuk Cha, Muhammad Saqlain, GeonU Kim, Mingyu Shin, and Seungryul Baek. Multi-person 3d pose and shape estimation via inverse kinematics and refinement. In *ECCV*, 2022.
- [6] Xu Chen, Tianjian Jiang, Jie Song, Jinlong Yang, Michael J Black, Andreas Geiger, and Otmar Hilliges. gdna: Towards generative detailed neural avatars. In *CVPR*, 2022.
- [7] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021.
- [8] Hongsuk Choi, Gyeongsik Moon, JoonKyu Park, and Kyoung Mu Lee. Learning to estimate robust 3d human mesh from in-the-wild crowded scenes. In *CVPR*, 2022.
- [9] Mihai Fieraru, Mihai Zanfir, Elisabeta Oneata, Alin-Ionut Popa, Vlad Olaru, and Cristian Sminchisescu. Three-dimensional reconstruction of human interactions. In *CVPR*, 2020.
- [10] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J Black. Resolving 3d human pose ambiguities with 3d scene constraints. In *ICCV*, 2019.
- [11] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. In *NeurIPS*, 2020.
- [12] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. *NIPS*, 2020.
- [13] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. Arch++: Animation-ready clothed human reconstruction revisited. In *ICCV*, 2021.
- [14] Chun-Hao P Huang, Hongwei Yi, Markus Höschle, Matvey Safroshkin, Tsvetelina Alexiadis, Senya Polikovsky, Daniel Scharstein, and Michael J Black. Capturing and inferring dense full-body human-scene contact. In *CVPR*, 2022.
- [15] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *CVPR*, 2020.
- [16] Uyoung Jeong, Seungryul Baek, Hyung Jin Chang, and Kwang In Kim. Boir: Box-supervised instance representation for multi-person pose estimation. In *BMVC*, 2023.
- [17] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, 2020.
- [18] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018.
- [19] Hyomin Kim, Hyeonsoo Nam, Jungeon Kim, Jaesik Park, and Seungyong Lee. Laplacianfusion: Detailed 3d clothed-human body reconstruction. *TOG*, 2022.
- [20] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *ICCV*, 2021.
- [21] 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 *ICCV*, 2019.
- [22] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *3DV*, 2019.
- [23] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *CVPR*, 2021.
- [24] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [25] Ren Li, Benoît Guillard, Edoardo Remelli, and Pascal Fua. Dig: Draping implicit garment over the human body. In *ACCV*, 2022.
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *TOG*, 2015.
- [27] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 1987.
- [28] Armin Mustafa, Akin Caliskan, Lourdes Agapito, and Adrian Hilton. Multi-person implicit reconstruction from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [29] 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 *ICCV*, 2019.
- [30] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020.
- [31] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, 2021.
- [32] Soshi Shimada, Vladislav Golyanik, Zhi Li, Patrick Pérez, Weipeng Xu, and Christian Theobalt. Hulc: 3d human motion capture with pose manifold sampling and dense contact guidance. In *ECCV*, 2022.
- [33] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J Black, and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *ICCV*, 2021.
- [34] Yu Sun, Wu Liu, Qian Bao, Yili Fu, Tao Mei, and Michael J Black. Putting people in their place: Monocular regression of 3d people in depth. In *CVPR*, 2022.

- [35] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-gif: Neural generalized implicit functions for animating people in clothing. In *ICCV*, 2021.
- [36] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *ECCV*, 2018.
- [37] Shaofei Wang, Marko Mihajlovic, Qianli Ma, Andreas Geiger, and Siyu Tang. Metaavatar: Learning animatable clothed human models from few depth images. *NIPS*, 2021.
- [38] Xianghui Xie, Bharat Lal Bhatnagar, and Gerard Pons-Moll. Chore: Contact, human and object reconstruction from a single rgb image. In *ECCV*, 2022.
- [39] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. Econ: Explicit clothed humans optimized via normal integration. In *CVPR*, 2023.
- [40] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. Icon: implicit clothed humans obtained from normals. In *CVPR*, 2022.
- [41] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgb sensors. In *CVPR*, 2021.
- [42] Andrei Zanfir, Elisabeta Marinoiu, and Cristian Sminchisescu. Monocular 3d pose and shape estimation of multiple people in natural scenes-the importance of multiple scene constraints. In *CVPR*, 2018.
- [43] Andrei Zanfir, Elisabeta Marinoiu, Mihai Zanfir, Alin-Ionut Popa, and Cristian Sminchisescu. Deep network for the integrated 3d sensing of multiple people in natural images. *NIPS*, 31, 2018.
- [44] Jianfeng Zhang, Dongdong Yu, Jun Hao Liew, Xuecheng Nie, and Jiashi Feng. Body meshes as points. In *CVPR*, 2021.
- [45] Yang Zheng, Ruizhi Shao, Yuxiang Zhang, Tao Yu, Zerong Zheng, Qionghai Dai, and Yebin Liu. Deepmulticap: Performance capture of multiple characters using sparse multiview cameras. In *ICCV*, 2021.
- [46] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE transactions on pattern analysis and machine intelligence*, 2021.
- [47] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *ICCV*, 2019.
- [48] Hao Zhu, Xinxin Zuo, Sen Wang, Xun Cao, and Ruigang Yang. Detailed human shape estimation from a single image by hierarchical mesh deformation. In *CVPR*, 2019.
- [49] Pierre Zins, Yuanlu Xu, Edmond Boyer, Stefanie Wuhler, and Tony Tung. Data-driven 3d reconstruction of dressed humans from sparse views. In *2021 International Conference on 3D Vision (3DV)*, 2021.