

# Pixel2ISDF: Implicit Signed Distance Fields based Human Body Model from Multi-view and Multi-pose Images

Jianchuan Chen<sup>\*1</sup>, Wentao Yi<sup>\*1</sup>, Tiantian Wang<sup>\*2</sup>, Xing Li<sup>\*3</sup>, Liqian Ma<sup>4</sup>,  
Yangyu Fan<sup>3</sup>, and Huchuan Lu<sup>†1</sup>

<sup>1</sup> Dalian University of Technology, China

<sup>2</sup> University of California, Merced

<sup>3</sup> Northwestern Polytechnical University, China

<sup>4</sup> ZMO AI Inc.

{Janaldo, raywit}@mail.dlut.edu.cn, lhchuan@dlut.edu.cn

twang61@ucmerced.edu, fan\_yangyu@nwpu.edu.cn

lixing36@foxmail.com, liqianma.scholar@outlook.com

**Abstract.** In this report, we focus on reconstructing clothed humans in the canonical space given multiple views and poses of a human as the input. To achieve this, we utilize the geometric prior of the SMPLX model in the canonical space to learn the implicit representation for geometry reconstruction. Based on the observation that the topology between the posed mesh and the mesh in the canonical space are consistent, we propose to learn latent codes on the posed mesh by leveraging multiple input images and then assign the latent codes to the mesh in the canonical space. Specifically, we first leverage normal and geometry networks to extract the feature vector for each vertex on the SMPLX mesh. Normal maps are adopted for better generalization to unseen images compared to 2D images. Then, features for each vertex on the posed mesh from multiple images are integrated by MLPs. The integrated features acting as the latent code are anchored to the SMPLX mesh in the canonical space. Finally, latent code for each 3D point is extracted and utilized to calculate the SDF. Our work for reconstructing the human shape on canonical pose achieves 3rd performance on WCPA MVP-Human Body Challenge.

**Keywords:** Implicit Representation, SMPLX, Signed Distance Function, Latent Codes Fusion

## 1 Introduction

High-fidelity human digitization has attracted a lot of interest for its application in VR/AR, image and video editing, telepresence, virtual try-on, etc. In this work, we target at reconstructing the high-quality 3D clothed humans in the

---

<sup>\*</sup> J. Chen, W. Yi, T. Wang and X. Li - Equal contributions.

<sup>†</sup> Corresponding author.

canonical space given multiple views and multiple poses of a human performer as the input.

Our network utilizes multiple images as the input and learns the implicit representation for the given points in the 3D space. Inspired by the advantages of implicit representation such as arbitrary topology and continuous representation, we adopt this representation to reconstruct high-fidelity clothed 3D humans. To learn geometry in the canonical space, we utilize the SMPLX mesh [15] in the canonical space as the geometric guidance. Due to the correspondence between the posed mesh and the mesh in the canonical space, we propose to first learn the latent codes in the posed mesh and then assign the latent codes to the canonical mesh based on the correspondence. By utilizing the posed mesh, image information in the 2D space can be included by projecting the posed mesh to the image space.

Given the multi-view images as input, normal and geometry networks are utilized to extract the features for the vertices on the SMPLX mesh [15]. We utilize a normal map as the intermediate prediction which helps generate sharp reconstructed geometry [21].

To integrate the features from different views or poses, we utilize a fusion network to generate a weighted summation of multiple features. Specifically, we first concatenate the features with the means and variances of features from all inputs followed by a Multi-layer Perceptron (MLP) predicting the weight and transformed features. Then, the weighted features will be integrated into the latent code through another MLP.

The latent code learned by the neural network is anchored to the SMPLX mesh in the canonical space, which serves as the geometry guidance to reconstruct the 3D geometry. Because of the sparsity of the vertices, we utilize a SparseConvNet to generate effective features for any 3D point following [19]. Finally, we use the trilinear interpolation to extract the latent code followed by an MLP to produce the SDF which models the signed distance of a point to the nearest surface.

## 2 Related Work

**Implicit Neural Representations.** Recently, the implicit representation encoded by a deep neural network has gained extensive attention, since it can represent continuous fields and can generate the details on the clothes, such as wrinkles and facial expressions. Implicit representation has been applied successfully to shape representation [16, 18, 26]. Here we utilize the signed distance function (SDF) as the implicit representation, which is a continuous function that outputs the point’s signed distance to the closest surface, whose sign encodes whether the point is inside (negative) or outside (positive) of the watertight surface. The underlying surface is implicitly represented by the isosurface of  $SDF(\cdot) = 0$ .

**3D Clothed Human Reconstruction.** Reconstructing humans has been widely studied given the depth maps [4, 22, 27], images [2, 10, 14], or videos [11, 13] as

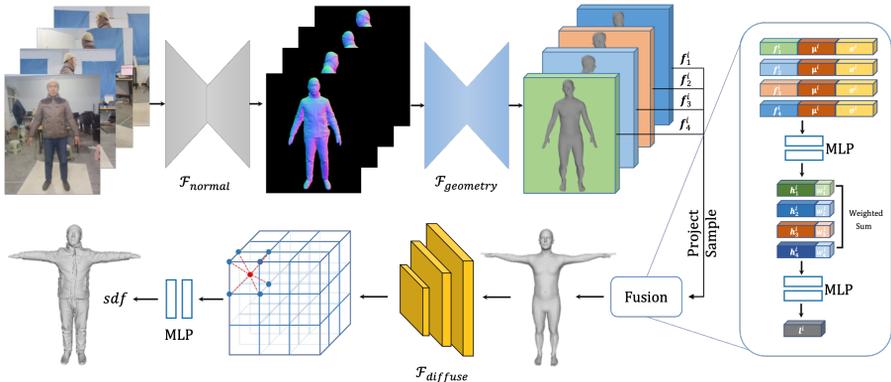


Fig. 1. Overview of our clothed human reconstruction pipeline.

the input. Utilizing the SMPLX mesh, the existing works show promising results with the RGB image as the input. NeuralBody [19] adopts the posed SMPLX mesh to construct the latent code volume that aims to extract the implicit pose code. ICON [24] proposes to refine the SMPLX mesh and normal map iteratively. SelfRecon [9] utilizes the initial canonical SMPLX body mesh to calculate the canonical implicit SDF. Motivated by these methods, we propose to utilize the SMPLX mesh in the canonical space as a geometric prior to reconstruct the clothed humans.

### 3 Methodology

Given the multi-view and multi-pose RGB images of human and SMPLX parameters as the input, we aim to reconstruct the clothed 3D geometry in the canonical space. Here the input images are denoted as  $\{I_k\}_{k=1}^K$ , where  $k$  denotes the image index and  $K$  is the number of images. The corresponding SMPLX parameters are denoted as  $\{\theta_k, \beta_k, s_k, t_k\}_{k=1}^K$ , where  $\theta_k, \beta_k$  are the pose and shape parameters of SMPLX and  $s_k, t_k$  are the camera parameters used for projection.

#### 3.1 Multi-view and Multi-pose Image Feature Encoding

Our feature extraction networks utilizes multi-view and multi-pose images  $\{I_k\}_{k=1}^K$  as input and outputs the geometric feature maps that help to predict the 3D geometry in the canonical space. Specifically, we first adopt the normal network  $\mathcal{F}_{normal}$  to extract the normal maps and then utilize the geometry network  $\mathcal{F}_{geo}$  to generate geometric feature maps  $\{F_k\}_{k=1}^K$  that will be further utilized to extract the pixel-aligned features for the vertices on the posed mesh.

$$F_k = \mathcal{F}_{geometry} (\mathcal{F}_{normal} (I_k)) \quad k = 1, 2, \dots, K \quad (1)$$

In particular, we adopt the pretrained image-to-image translation network from PIFuHD [21] as our normal network, and use the pretrained ResUNet34 [8] backbone as our geometry network.

### 3.2 Structured Latent Codes

After obtaining the geometric feature maps  $\{F_k\}_{k=1}^K$ , we extract the pixel-aligned features for each vertex  $\mathbf{v}_k^i$  on the posed mesh  $M(\theta_k, \beta_k)$ . For each vertex  $\mathbf{v}_k^i$ , we first project it to the image space by utilizing the weak-perspective projection  $\Phi$  according to camera parameters scale  $s_k$  and translation  $t_k$ , then adopt the bilinear interpolation operation to extract the pixel-aligned features  $\mathbf{f}_k^i$ .

$$\Phi : \hat{\mathbf{x}} = s\Pi(\mathbf{x}) + t, \quad \mathbf{x} \in \mathbb{R}^3, \hat{\mathbf{x}} \in \mathbb{R}^2 \quad (2)$$

$$\mathbf{f}_k^i = \text{bilinear}(F_k, \Phi(\mathbf{v}_k^i, s_k, t_k)) \quad (3)$$

where  $\Pi$  is the orthogonal projection,  $\mathbf{x}$  is the point in 3D space and  $\hat{\mathbf{x}}$  is the projected points in 2D image space.

To integrate the feature of the  $i$ -th vertex in the canonical space from multiple views/poses, we use a fusion network that takes  $\{\mathbf{f}_k^i\}_{k=1}^K$  as the input and outputs the integrated feature  $\mathbf{l}^i$ , which is illustrated in Figure 1. Specifically, the mean  $\boldsymbol{\mu}^i$  and variance  $\boldsymbol{\sigma}^i$  of features  $\{\mathbf{f}_k^i\}_{k=1}^K$  is calculated and then concatenated with  $\mathbf{f}_k^i$  to serve as the input of a MLP. The MLP predicts the new feature vector and weight  $\{w_k^i\}_{k=1}^K$  for each feature, which generates a weighted sum of features from multiple inputs. The weighted feature is finally forwarded to an MLP for feature integration,  $\mathbf{l}^i$  which serves as the structured latent code for the vertex  $\mathbf{v}_k^i$ .

Different from the latent code in NeuralBody [19] which is initialized randomly for optimizing specific humans, our latent code is the feature vector learned by a network with the normal map as the input, which can generalize to humans unseen from training ones.

The above-mentioned latent codes are generated based on the posed mesh. The posed mesh and the canonical mesh share the same latent codes because of their topology correspondence which forms the set of latent codes by  $\mathcal{Q} = \{\mathbf{l}^1, \mathbf{l}^2, \dots, \mathbf{l}^{N_V}\}$ ,  $\mathbf{l}^i \in \mathbb{R}^d$ . Here  $N_V$  represents the number of vertices.

### 3.3 Implicit Neural Shape Field

The learned latent codes are anchored to a human body model (SMPLX [15]) in the canonical space. SMPLX is parameterized by shape and pose parameters with  $N_V = 10,475$  vertices and  $N_J = 54$  joints. The locations of the latent codes  $\mathcal{Q} = \{\mathbf{l}^1, \mathbf{l}^2, \dots, \mathbf{l}^{N_V}\}$  are transformed for learning the implicit representation by forwarding the latent codes into a neural network

To query the latent code at continuous 3D locations, trilinear interpolation is adopted for each point. However, the latent codes are relatively sparse in the 3D space, and directly calculating the latent codes using trilinear interpolation

will generate zero vectors for most points. To overcome this challenge, we use a SparseConvNet [19] to form a latent feature volume  $\mathcal{V} \in \mathbb{R}^{H \times H \times H \times d}$  which diffuses the codes defined on the mesh surface to the nearby 3D space.

$$\mathcal{V} = \mathcal{F}_{diffuse}(\{\mathbf{l}^1, \mathbf{l}^2, \dots, \mathbf{l}^{N_v}\}) \quad (4)$$

Specifically, to obtain the latent code for each 3D point, trilinear interpolation is employed to query the code at continuous 3D locations.

Here the latent code will be forwarded into a neural network  $\varphi$  to predict the SDF  $\mathcal{F}_{sdf}(\mathbf{x})$  for 3D point  $\mathbf{x}$ .

$$\mathcal{F}_{sdf}(\mathbf{x}) = \varphi(\text{trilinear}(\mathcal{V}, \mathbf{x})) \quad (5)$$

### 3.4 Loss Function

During training, we sample  $N_S$  spatial points  $\mathcal{X}$  surrounding the ground truth canonical mesh. To train the implicit SDF, we deploy a mixed-sampling strategy: 20% for uniform sampling on the whole space and 80% for sampling near the surface. We adopt the mixed-sampling strategy because of the following two reasons. First, sampling uniformly in the 3D space will put more weight on the points outside the mesh during network training, which results in overfitting when sampling around the iso-surface. Second, sampling points far away from the reconstructed surface contribute little to geometry reconstruction, which increases the pressure of network training.

Overall, we enforce 3D geometric loss  $L_{sdf}$  and normal constraint loss  $L_{normal}$ .

$$\mathcal{L} = \lambda_{sdf} \mathcal{L}_{sdf} + \lambda_{normal} \mathcal{L}_{normal} \quad (6)$$

**3D Geometric Loss.** Given a sampling point  $\mathbf{x} \in \mathcal{X}$ , we employ the L2 loss between the predicted  $\mathcal{F}_{sdf}(\mathbf{x})$  and the ground truth  $\mathcal{G}_{sdf}(\mathbf{x})$ , which are truncated by a threshold  $\delta$ ,

$$\mathcal{L}_{sdf} = \frac{1}{N_S} \sum_{\mathbf{x} \in \mathcal{X}} \|\mathbb{C}(\mathcal{F}_{sdf}(\mathbf{x}), \delta) - \mathbb{C}(\mathcal{G}_{sdf}(\mathbf{x}), \delta)\|_2. \quad (7)$$

Here  $\mathbb{C}(\cdot, \delta) = \max(-\delta, \min(\cdot, \delta))$ .

**Normal Constraint Loss.** Beyond the geometric loss, to make the predicted surface smoother, we deploy Eikonal loss [6] to encourage the gradients of the sampling points to be close to 1.

$$\mathcal{L}_{normal} = \frac{1}{N_S} \sum_{\mathbf{x} \in \mathcal{X}} \left| \|\nabla_{\mathbf{x}} \mathcal{F}_{sdf}(\mathbf{x})\|_2 - 1 \right|_2. \quad (8)$$

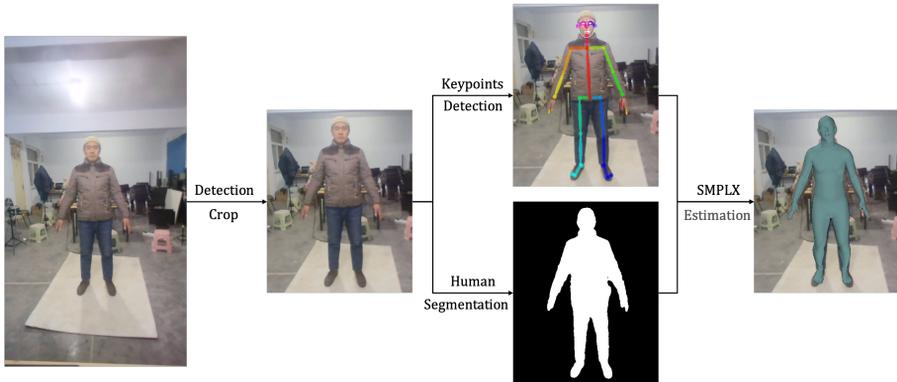


Fig. 2. The pipeline for data processing.

## 4 Experiments

### 4.1 Datasets

We use the WCPA [1] dataset as training and testing datasets, which consists of 200 subjects for training and 50 subjects for testing. Each subject contains 15 actions, and each action contains 8 RGB images from different angles (0, 45, 90, 135, 180, 225, 270, 315). Each image is an  $1280 \times 720$  jpeg file. The ground truth of the canonical pose is a high-resolution 3D mesh with detailed information about clothes, faces, etc. For the training phase, we randomly select 4 RGB images with different views and poses as inputs to learn 3D human models.

### 4.2 Image Preprocessing

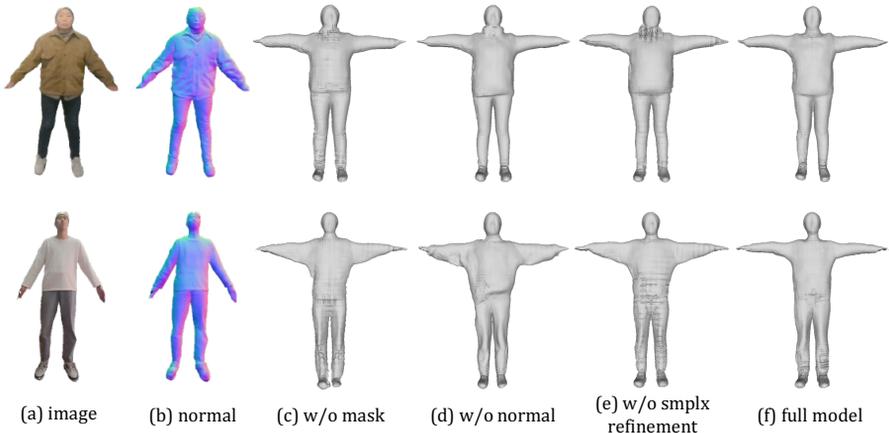
**Image Cropping.** For each image, we first apply VarifocalNet [28] to detect the bounding boxes to localize the humans. Next, we crop the input images with the resolution of  $512 \times 512$  according to the bounding boxes. When the cropped images exceed the bounds of the input images, the cropped images will be padded with zeros.

**Mask Generation.** For each image, we first apply DensePose [7] to obtain part segmentations. Then we set the parts on the human as the foreground human mask and set the values of the background image pixels as zero. The masked images are served as the input of our network.

The mask is then refined using the MatteFormer [17], which can generate better boundary details. Following [23], the trimap adopted in [17] is generated based on the mask using the erosion and dilation operation.

### 4.3 SMPLX Estimation and Optimization

It is very challenging to get accurate SMPLX from 2D RGB images due to the inherent depth ambiguity. First, we use Openpose [3] to detect the 2D keypoints of



**Fig. 3. The visualization results on the test dataset.**

the person in the image and then use ExPose [5] to estimate the SMPLX parameters and camera parameters. However, due to extreme illumination conditions and complex backgrounds, the SMPLX obtained in this way is not sufficiently accurate, and we need to refine the SMPLX parameters further. We utilize the 2D keypoints and masks to optimize SMPLX parameters  $\theta, \beta$ . For the 2D keypoints loss, given the SMPLX 3D joints location  $J(\theta, \beta)$ , we project them to the 2D image using the weak-perspective camera parameters  $s, t$ . For the mask loss, we utilize PyTorch3D [20] to render the 2D mask given the posed mesh  $M(\theta, \beta)$ . Then the mask loss is calculated based on the rendered mask and the pseudo ground truth mask  $\mathcal{M}_{gt}$ .

$$\theta^*, \beta^* = \min_{\theta, \beta} \|\Phi(J(\theta, \beta), s, t) - \mathcal{J}_{gt}\| + \lambda \|\Psi(M(\theta, \beta), s, t) - \mathcal{M}_{gt}\| \quad (9)$$

#### 4.4 Implementation Details

During training, we randomly choose  $K = 4$  images from total of  $8 * 15$  images as input to extract the 2D feature map. Taking into account the memory limitation, we set the latent feature volume  $\mathcal{V} \in \mathbb{R}^{224 \times 224 \times 224 \times 64}$ , where each latent code  $\mathbf{l}^i \in \mathbb{R}^{64}$ . During training, we randomly sample  $N_t = 10,000$  points around the complete mesh. To stably train our network, we initialize the SDF to approximate a unit sphere [25]. We adopt the Adam optimizer [12] and set the learning rate  $lr = 5e-4$ , and it spends about 40 hours on 2 Nvidia GeForce RTX 3090 24GB GPUs. For inference, the surface mesh is extracted by the zero-level set of SDF by running marching cubes on a  $256^3$  grid.

#### 4.5 Results

Our method achieved very good results in the challenge, demonstrating the superiority of our method for 3D human reconstruction from multi-view and multi-pose images. To further analyze the effectiveness of our method, we performed

different ablation experiments. According to the results as shown in the table 1 and the figure 3, removing the background allows the model to better reconstruct the human body. The two-stage approach of predicting the normal vector map as input has a stronger generalization line on unseen data compared to directly using the image as input. What is more important is that the precision of SMPLX has a significant impact on the reconstruction performance, demonstrating the necessity of SMPLX optimization.

	w/o mask	w/o normal map	w/o SMPLX refinement	full model
Chamfer Distance ↓	1.1277	1.0827	1.1285	<b>0.9985</b>

**Table 1.** Quantitative metrics of different strategies on the test dataset.

## 5 Conclusion

Modeling 3D humans accurately and robustly from challenging multi-views and multi-posed RGB images is a challenging problem, due to the varieties of body poses, viewpoints, light conditions, and other environmental factors. Our key idea is to overcome these challenges by constructing structured latent codes as the inputs for implicit representation. The latent codes integrate the features for vertices in the canonical pose from different poses or views. To this end, we have contributed a deep-learning based framework to incorporate the parametric SMPLX model and non-parametric implicit function for reconstructing a 3D human model from multi-images.

## References

1. Eeccv 2022 wcpa challenge (2022), <https://tianchi.aliyun.com/competition/entrance/531958/information> 6
2. Bogo, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In: European conference on computer vision. pp. 561–578. Springer (2016) 2
3. Cao, Z., Hidalgo, G., Simon, T., Wei, S., Sheikh, Y.: Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Trans. Pattern Anal. Mach. Intell.* pp. 172–186 (2021) 6
4. Chibane, J., Alldieck, T., Pons-Moll, G.: Implicit functions in feature space for 3d shape reconstruction and completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6970–6981 (2020) 2
5. Choutas, V., Pavlakos, G., Bolkart, T., Tzionas, D., Black, M.J.: Monocular expressive body regression through body-driven attention. In: Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part X. vol. 12355, pp. 20–40 (2020) 7
6. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. arXiv preprint arXiv:2002.10099 (2020) 5
7. Güler, R.A., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7297–7306 (2018) 6
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) 4
9. Jiang, B., Hong, Y., Bao, H., Zhang, J.: Selfrecon: Self reconstruction your digital avatar from monocular video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5605–5615 (2022) 3
10. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7122–7131 (2018) 2
11. Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3d human dynamics from video. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5614–5623 (2019) 2
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) 7
13. Kocabas, M., Athanasiou, N., Black, M.J.: Vibe: Video inference for human body pose and shape estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5253–5263 (2020) 2
14. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2252–2261 (2019) 2
15. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* **34**(6), 1–16 (2015) 2, 4
16. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4460–4470 (2019) 2

17. Park, G., Son, S., Yoo, J., Kim, S., Kwak, N.: Matteformer: Transformer-based image matting via prior-tokens. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11696–11706 (June 2022) [6](#)
18. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019) [2](#)
19. Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., Zhou, X.: Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In: CVPR (2021) [2](#), [3](#), [4](#), [5](#)
20. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. arXiv preprint arXiv:2007.08501 (2020) [7](#)
21. Saito, S., Simon, T., Saragih, J., Joo, H.: 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. pp. 84–93 (2020) [2](#), [4](#)
22. Wang, L., Zhao, X., Yu, T., Wang, S., Liu, Y.: Normalgan: Learning detailed 3d human from a single rgb-d image. In: European Conference on Computer Vision. pp. 430–446. Springer (2020) [2](#)
23. Wang, T., Liu, S., Tian, Y., Li, K., Yang, M.H.: Video matting via consistency-regularized graph neural networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4902–4911 (2021) [6](#)
24. Xiu, Y., Yang, J., Tzionas, D., Black, M.J.: Icon: Implicit clothed humans obtained from normals. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). vol. 2 (2022) [3](#)
25. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* **33**, 2492–2502 (2020) [7](#)
26. Yifan, W., Wu, S., Oztireli, C., Sorkine-Hornung, O.: Iso-points: Optimizing neural implicit surfaces with hybrid representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 374–383 (2021) [2](#)
27. 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: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7287–7296 (2018) [2](#)
28. Zhang, H., Wang, Y., Dayoub, F., Sunderhauf, N.: Varifocalnet: An iou-aware dense object detector. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8514–8523 (2021) [6](#)