

A Neural Network for Detailed Human Depth Estimation from a Single Image

Sicong Tang^{1,*} Feitong Tan^{1,*} Kelvin Cheng¹ Zhaoyang Li¹ Siyu Zhu² Ping Tan¹
¹ Simon Fraser University ² Alibaba A.I Labs

{sta105, feitongt, kelvinz, zla143, pingtan}@sfu.ca, siting.zsy@alibaba-inc.com

Abstract

This paper presents a neural network to estimate a detailed depth map of the foreground human in a single RGB image. The result captures geometry details such as cloth wrinkles, which are important in visualization applications. To achieve this goal, we separate the depth map into a smooth base shape and a residual detail shape and design a network with two branches to regress them respectively. We design a training strategy to ensure both base and detail shapes can be faithfully learned by the corresponding network branches. Furthermore, we introduce a novel network layer to fuse a rough depth map and surface normals to further improve the final result. Quantitative comparison with fused ‘ground truth’ captured by real depth cameras and qualitative examples on unconstrained Internet images demonstrate the strength of the proposed method.

1. Introduction

Understanding human images is an important problem in computer vision with many applications ranging from human-computer-interaction and surveillance to telecommunication. Many works [25, 22, 34, 33, 21, 18] have been developed to recover 2D or 3D skeleton joints from a RGB image. Since the skeleton only captures sparse information of the human body, DensePose [2] estimates a dense UV map (i.e. a correspondence map between the input image and a 3D template model). But this UV map can not recover 3D shape without additional 3D pose information, which limits its application.

On the other hand, there are many works [3, 19, 7, 31, 13, 17, 4, 10, 27] to recover a dense 3D deformable model of the human body from a single image, e.g. the SCAPE [3] and SMPL [19] models, which are learned from a large dataset of scanned body shapes. While generating 3D models, these methods only inference the naked body shape without capturing the clothes details.

This paper aims at recovering a detailed depth map for

the foreground human object from a single RGB image. This problem has been studied in the earlier work [36] with synthetic human images. Another recent work [35] recovers a volumetric 3D model of the imaged person. Results from both methods are too coarse for many applications. In comparison, we design a neural network to estimate highly detailed depth maps that are fine enough to capture cloth wrinkles, which might potentially be exploited for telepresence applications like the Microsoft Holoportation [8].

Our network is designed with two novel insights. Firstly, we argue it is important to separate the depth to a smooth base shape and a residual detail shape and regress them respectively. The base shape captures the large overall geometry layout, while the detail shape captures small bumps such as cloth wrinkles. The value range of the base shape is at the scale of one meter, while that of the detail shape is at a few centimeters. Thus, we design a network with two branches for the base and detail shapes respectively to facilitate the training process. Specifically, we propose a 2-stage training strategy to ensure the effectiveness of this separation. These two branches are trained respectively in the 1st stage and then finetuned together in the 2nd stage. Secondly, we follow the intuition in [40] to estimate surface normals to facilitate depth map estimation. Specifically, we generalize the algorithm in [23] that fuses surface normals and a coarse depth to an iterative formulation. In this way, we build a parameter-free network layer to fuse the estimated normals and a coarse depth map for improved results.

Our final network captures visually appealing detailed depth images from a single RGB image. The evaluation on our own captured real data and some unconstrained online images demonstrate its effectiveness. We will publish our dataset and source code with the paper to facilitate further research.

2. Related works

3D Human Pose Estimation. With the recent development of deep convolutional neural networks (CNNs), there are significant improvements on 3D human pose estimation [21, 33, 18, 22]. Despite the differences in network architectures, many works [25, 28, 22, 33, 35] use a like-

*These authors contributed equally to this work.

likelihood heatmap to represent the distribution of each joint’s location and show better performance than directly regressing the joint location. Instead of taking the maximum from a heatmap, Sun *et al.* [33] compute the expected coordinates from a heatmap to reduce the artifacts due to quantization. The recent work DensePose[2] is even able to recover dense UV coordinate for each pixel on human body. Unlike our method, most of these methods only recover sparse 3D joint positions. While DensePose provides dense result, its result is not in 3D but rather a 2D UV coordinate map. We adopt a pose estimation network as an intermediate layer and use its results to guide the dense depth recovery.

Body Shape Estimation. The 3D shape of a human body can be parameterized by the SCAPE or SMPL models [3, 19] with two sets of independent parameters, controlling the skeleton pose and body shape respectively. Both models are derived from a large set of scanned 3D human shapes. Given these parametric human models, many methods [3, 19, 7, 31, 13, 17, 27] recover dense human body shape from a single RGB image by estimating the shape and pose parameters. Meanwhile, there are also some non-paramterized methods [36, 35] which directly regress discretized body shape representation from a RGB image. The above methods only recover the 3D shape of the naked human body and geometry details like the clothes are not modeled, which make them not suitable for visualization tasks. While the method [1] can predict the SMPL model with clothe wrinkle, it needs to be fed a video of a moving person with designed pose. To overcome this limitation, our network aims at recovering shape details from a single image.

Generic Dense Depth Estimation. Depth estimation from a single image has gained increasing attention in the computer vision community. Most works like [37, 38, 20, 15, 39, 41, 9, 16] are proposed for indoor and outdoor scenes. We focus on depth estimation of humans, which allows us to build much stronger shape prior than these generic depth estimation methods. Specifically, our network first estimates the skeleton joints and a body part segmentation to facilitate the depth estimation.

3. Overview

The overall structure of the proposed network is shown in Figure 1. The input is a 256×256 3-channel RGB image containing a human as the foreground. The network first computes the heatmaps of the 3D skeleton joints and a body part segmentation through two Hourglass networks [25], which are referred as *Skeleton-Net* and *Segmentation-Net* respectively in this paper. We then concatenate the outputs of these two modules with the input RGB image and feed them to the *Depth-Net* to compute the initial depth maps, which consists of a base shape and a detail shape.

In a separate branch, another Hourglass network, re-

ferred as *Normal-Net*, computes a surface normal map of the human body from the input RGB image and the segmentation mask generated by the Segmentation-Net. We then compose the base shape and detail shape, and fuse the composed shape and normal map through a parameter-free shape refinement module to produce the final shape.

During training, we first pre-train the Skeleton-Net, Segmentation-Net, and Depth-Net on synthetic data [36] respectively. Meanwhile, the Normal-Net is pre-trained on the deforming fibre dataset [5]. Then we finetune the complete network on the real image dataset captured by ourselves with a depth camera, while keeping the parameters of Skeleton-Net and Segmentation-Net fixed.

4. Segmentation and Skeleton Networks

Inspired by the BodyNet [35], 3D joints and body part segmentation are highly correlated with the final estimation of human shapes. We therefore apply two Hourglass networks [25] to estimate the heatmaps of 3D joints and a body part segmentation from the input RGB image. As demonstrated in the ablation studies, this intermediate supervision of 3D joints and body part segmentation is essential for the depth estimation, especially for the base shape.

Here, a human body contains 16 joints and 14 body parts. For each joint, our Skeleton-Net predicts a heatmap indicating the probability of its position [29]. The 3D joints are defined in the camera coordinate system, where the xy -axes are aligned with the image axes, and the z axis is the camera principal direction. We discretize the z coordinate between $[-0.6, 0.6]$ meters into 19 bins and set the depth of the pelvis joint as 0. The x and y coordinates are discretized into 64 bins over the image plane. Therefore, the network estimate a heatmap of size $64 \times 64 \times 19$ for each joint, resulting in a skeleton representation as a $64 \times 64 \times 19 \times 16$ heatmap.

Unlike [35], we discard the 2D joint estimation sub-network and predict the 3D joints directly, which makes our network more compact. In order to achieve good accuracy with this compact network, we adopt the integral regression [33] to train the Skeleton-Net.

For body part segmentation, the Segmentation-Net predicts the probability heatmap for the 14 body parts and the background, which results in a $64 \times 64 \times 15$ heatmap. Following the previous work of human part segmentation [26], we adopt the spatial cross-entropy loss in training.

5. Depth Estimation Network

To better estimate a detailed depth map with cloth wrinkles, we divide the depth map of a human body into a smooth base shape and a residual detail shape: the base shape captures the main geometry layout of the human body, while the detail shape is responsible for describing

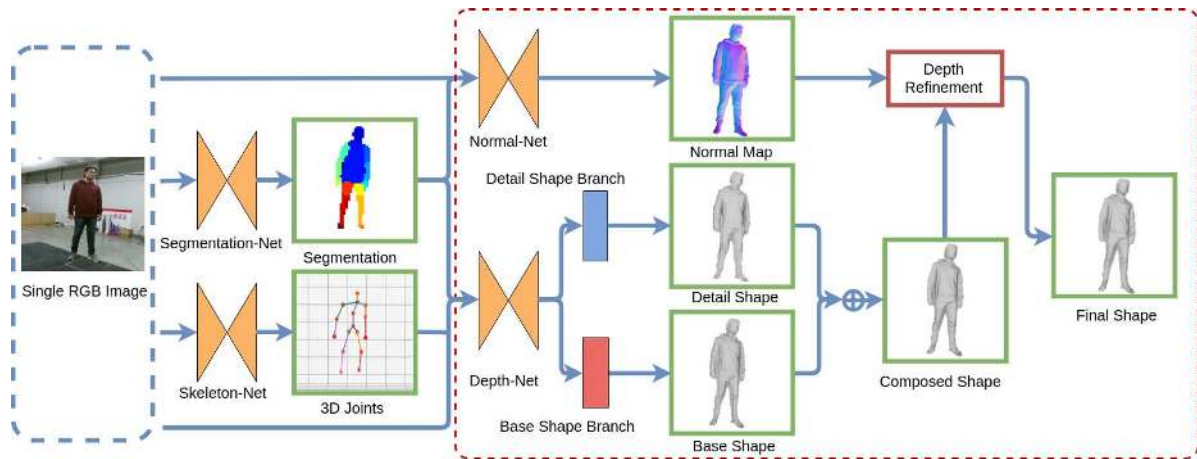


Figure 1. The structure of our proposed network. The Skeleton-Net and Segmentation-Net generate the heatmaps of 3D skeleton joints and body part segmentation respectively. Their results are further fused with the input image to compute the base shape and detail shape via the Depth-Net. In a separate branch, the Normal-Net estimates a surface normal map. The composed shape and normal map are further fused in the depth refinement module to produce the final result.

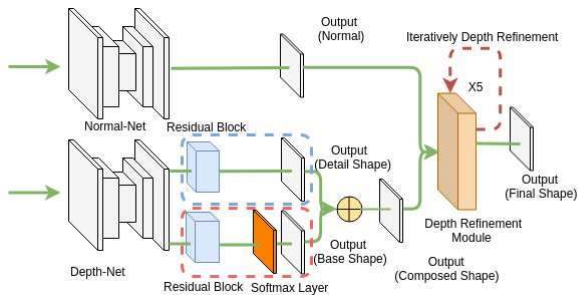


Figure 2. Architecture of Depth-Net together with base shape branch and detail shape branch, Normal-Net and Depth Refinement Module. The branches in blue and red dashed rectangles correspond to the detail and base shape branch respectively.

local geometry details such as cloth wrinkles.

As shown in Figure 2 which corresponds to the part in the red dashed rectangle in Figure 1, the Depth-Net is composed of a U-Net [32] and a two-branch architecture. The concatenation of the RGB image and bilinearly-upsampled heatmaps (64×64 to 256×256) of 3D joints and segmentation is fed into this network, and the two branches, namely base and detail shape branch, output a base shape and detail shape separately. Because the human layout is approximately one-meter range with low frequency in image plane and the detailed cloth wrinkles is just several centimeters with higher frequency, the two branches concentrate on these two different distributions respectively.

To effectively train the Depth-Net, we set the median of the ground-truth depth as 0 and decouple this zero-median depth image into a base shape and detail shape. Specifically, we apply the bilateral filter to the depth image to smooth out the details and obtain the base shape. We denote this base shape as $F(D_{gt})$, where D_{gt} is the ground-truth depth

image and $F(\cdot)$ is the operation of the bilateral filter. In our work, the depth sigma is set as 0.10 meters and the space sigma is set as 75 pixels for the bilateral filter. The ground-truth of the detail shape is computed as a residual R_{gt} :

$$R_{gt} = D_{gt} - F(D_{gt}). \quad (1)$$

For the base shape, we discretize the depth range between $[-0.6, 0.6]$ meters into 19 bins for each pixel. The softmax layer which follows a residual block in the base branch generates a $256 \times 256 \times 19$ heatmap indicating the probability of the depth bin. Afterwards, a 256×256 depth map can be calculated from the heatmap by an integral operation [33]. Meanwhile, in the detail branch, a residual depth map of detail shape which has a higher frequency is regressed directly. At last, we add the base shape and detail shape together to obtain the composed shape.

In order to guide the base and detail branch to focus on their target domain (base shape and detail shape), we train our Depth-Net following a two-stage strategy. In the first stage, the base and detail branch are pre-trained separately to obtain well-conditioned initial values. In the second stage, we perform end-to-end training on three combined weights with the supervision of the intermediate base and detail shape branches.

5.1. Training stage 1

Once we have the ground-truth base and detail shape, we pre-train these two branches independently with the following loss functions:

$$\begin{aligned} L_{base} &= H(D_{base} - F(D_{gt}), \alpha_1), \\ L_{detail} &= H(D_{detail} - R_{gt}, \alpha_2), \end{aligned} \quad (2)$$

where D_{base} and D_{detail} are the base and detail depth to be regressed respectively. $H(x, \alpha)$ is the Huber loss function,

α_1 and α_2 are set as 0.2 meters and 0.05 meters. Here, $H(x, \alpha)$ is defined as:

$$H(x, \alpha) = \begin{cases} 0.5x^2, & x \leq \alpha, \\ 0.2(|x| - \alpha), & x > \alpha. \end{cases} \quad (3)$$

This pre-training helps the two branches focus on different aspect of the shape estimation, where the base shape captures the main geometry layout and the detail shape adds on high-frequency wrinkles.

5.2. Training stage 2

In this stage, we jointly train these two branches by using the combined loss L below:

$$L = \beta_1 L_{base} + \beta_2 L_{detail} + \beta_3 L_{composed}, \quad (4)$$

where $\beta_1, \beta_2, \beta_3$ are set as 1, 1, 15. Here, the composed loss $L_{composed}$ is formulated as:

$$L_{composed} = T(D_{base} + D_{detail} - D_{gt}, \alpha_3), \quad (5)$$

where α_3 is set to 0.05 meters in our experiments. $T(x, \alpha)$ is the truncated L_1 loss and it is defined as:

$$T(x, \alpha) = \begin{cases} x, & x \leq \alpha, \\ \alpha, & x > \alpha. \end{cases} \quad (6)$$

The stage 2 improves the consistency between the combined shape and the ground truth, and the truncated L_1 loss is used to define the composed loss $L_{composed}$ which clips the loss value to a bounded range. This truncated loss helps to avoid the training being biased by large shape errors due to imprecise poses, which could overwhelm the errors due to missing cloth wrinkles. As we will see in experiments, this loss helps the detail shape branch to capture details.

6. Normal Network and Depth Refinement

As observed in [40], regressing surface normal is often more reliable than regressing depth directly. We include a network to regress the surface normal at every pixel and use this information to refine the composed depth.

6.1. Normal Network

Here, a Hourglass network takes a RGB image concatenated with a segmentation mask from the Segmentation-Net as input and outputs a normal map.

This network is trained with the ground-truth normal computed from the ground-truth depth map D_{gt} . To compute the ground-truth normal N_{gt} , we take the nearby 3D points at each pixel to estimate its normal direction by the standard linear least square fitting. The loss function is the mean angular difference between the ground-truth and the regressed normal.

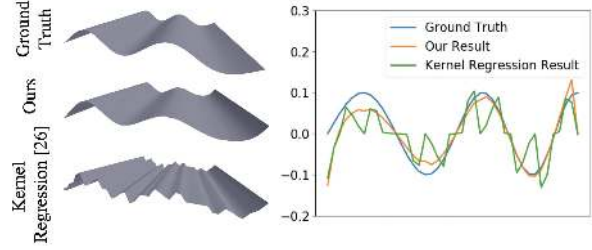


Figure 3. Comparison of our depth refinement with [30] on a toy example of Sine curve. Left: ground-truth, results from our method and the [30] (from top to bottom). Right: sectional view of these results.

6.2. Depth Refinement

We fuse the composed depth and surface normal here to improve the depth quality. Similar to [24], we formulate the problem with two constraints. Firstly, the tangent vector of the final shape should be perpendicular to the input surface normal at each pixel. Secondly, the final shape should be close to the initial shape. Rather than solving a large linear system for a global optimization which is impractical for a neural network, we introduce an iterative solution.

At each iteration, we update the depth assuming its neighboring depth is fixed. Concretely, we define (N_{ix}, N_{iy}, N_{iz}) as the normal of pixel i in x, y, z directions, and (X_i^n, Y_i^n, Z_i^n) as the position of pixel i after the n -th iteration. At the $n+1$ -th iteration, we update Z_i^{n+1} for each pixel i with the depth of neighboring pixels fixed at Z_j^n . Here, $j \in \mathcal{N}_i$ is a neighboring pixel of i and there are 4 neighbors for each pixel in cardinal directions. The update function is defined as:

$$Z_i^{n+1} = \lambda Z_i^0 + (1 - \lambda) \frac{\sum_{j \in \mathcal{N}_i} (Z_{ij}^n + Z_{ji}^n)}{8}, \quad (7)$$

where Z_{ij}^n is the depth of i that makes the edge ij and N_j perpendicular, and Z_{ji}^n is the depth of i that makes ij and N_i perpendicular. Specifically, they can be computed as:

$$\begin{aligned} Z_{ij}^n &= \frac{N_{jx}(X_j^n - X_i^n) + N_{jy}(Y_j^n - Y_i^n) + N_{jz}Z_j^n}{N_{jz}}, \\ Z_{ji}^n &= \frac{N_{ix}(X_j^n - X_i^n) + N_{iy}(Y_j^n - Y_i^n) + N_{iz}Z_i^n}{N_{iz}}. \end{aligned} \quad (8)$$

Here, λ is the hyper-parameter (fixed at 0.4).

The above shape refinement is iterated for 5 times in our network to simulate the iterative solution of the original energy equation in [24]. Figure 3 compares our method with the ‘Kernel Regression’ layer [30] on a toy example, which is also designed to fuse the surface normal and depth. Figure 4 shows a comparison with the work [24] on real data and our method also produces more convincing result.

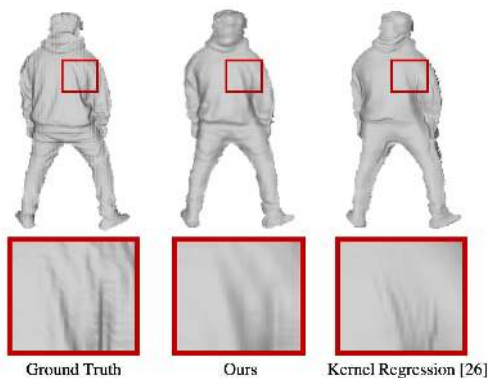


Figure 4. Comparison of our depth refinement with the ‘Kernel Regression’ in [30] on a real data. From left to right, there are the ground truth shape, results of our method and the ‘Kernel Regression’ respectively.

7. Experiment

To demonstrate the effectiveness of our method, we evaluate it using ablation studies and both qualitative and quantitative comparisons with other relevant works [36, 35], a surface-from-normals method [14] and a general depth estimation network [16]. To test the performance of human shape estimation with fine-grain geometrical details, we build up our own dataset for evaluation.

Implementation Details. All input RGB images are cropped to center the person with size 256×256 , assuming that the bounding box of person is given. The RM-Sprop [11] algorithm with a fixed learning rate of 1×10^{-5} is used. We first train our Segmentation-Net, Skeleton-Net and Depth-Net on SURREAL [36], a large-scale synthetic human body dataset without geometrical details. At this stage, the batch size is set to 6 for these three networks, and for Depth-Net we only add base shape loss to train base shape branch since the synthetic data does not have much geometrical details. The Normal-Net is pre-trained on a deforming fibre dataset [5]. After the base shape branch of Depth-Net converges, which takes 10 epochs, 12 hours on a GTX 2080 GPU, we fix the weight of Skeleton-Net and Segmentation-Net and fine-tune the Depth-Net and Normal-Net jointly on our own captured data with a batch size of 1. It takes another 12 epochs, 10 hours for stage 1 and another 8 epochs, 6 hours on stage 2. During inference, our network takes 75.5ms for the whole pipeline, and 61.1ms without iterative depth refinement on a RTX 2080.

Dataset. We collect a RGBD dataset for real persons. Here the dataset contains 26 different people performing simple actions captured by a Microsoft Kinect2 camera.

For the training data, we capture approximate 800 frames for each person, leading to over 20,000 training depth images in total. For quantitative evaluation, we use depth cameras to capture video clips of a person with a fixed pose and

Methods	Accuracy			MAE
	1.25cm	2.5cm	5.0cm	
Ours (Final Shape)	30.06	51.57	75.76	3.208
Ours (Base + Detail)	29.24	50.93	75.52	3.282
Ours (Base Shape)	28.03	50.10	75.32	3.396
Ours (Off-the-Shelf)	28.57	50.70	76.54	3.546
SURREAL [36]	21.32	37.52	50.06	3.976
BodyNet [35]	17.14	32.59	56.98	4.366
Laina et al. [16]	19.84	36.48	60.94	4.902
Kovese et al. [14]	15.51	29.87	55.39	5.789

Table 1. Performance of depth estimation on the test set. ‘Ours (Base)’ stands for the base shape without adding detail wrinkles. ‘Ours (Base + Detail)’ refers to the composed shape before the depth refinement.

employ the InfiniTAM [12] to fuse captured sequences. The high-quality depth maps are rendered according to the fused mesh and camera poses with Blender [6]. Our testing data contains 5 different persons, each person is captured with 12 different poses and 3 different clothing styles.

Note that we only use the fused depth maps for evaluation, the training data are raw depth maps since it is infeasible to fuse all the meshes with thousands of poses for rendering the depth maps.

7.1. Quantitative Results

Figure 5 shows our results compared with the fused ground-truth depth. We can see that our method can successfully capture cloth wrinkles and produce visually appealing 3D mesh from testing real images, despite our model is trained on the noisy raw depth images.

Comparison with [36, 16, 35, 14]. There are only a few works that can compute a depth map of human body from a single image. We compare with the two most recent works [36, 35] and a representative general depth estimation framework [16], and since we use normal map to refine human depth in our framework, we also evaluate a surface-from-normals method [14] with the normals from our Normal-Net. At last, to show the generalizability of our network, we replace our segmentation and 3D pose estimation module with off-the-shelf networks [32, 36] and evaluate the performance of Depth-Net. To make the comparison fair, we fine-tune [36, 16] on our dataset. Unfortunately, the BodyNet [35] needs a volumetric shape representation and its loss function contains the multiview constraints, thus it can not be fine-tuned on our data. Here, the pixel accuracy as percentage of pixels with depth errors smaller than some specified threshold is employed as the evaluation metric. It shows in Table 1 that the final shape after refinement always produces the highest accuracy. Here we notice that our network still works well with off-the-shelf segmentation and 3D pose estimation methods, and deducing the correct human shape just from normal is difficult due to noisy normal



Figure 5. Some results on the testing data. From left to right, these images are: the single input RGB image, the ground truth shape and our result. It can be seen that our method is able to recover the main layout as well as certain geometry details. Note that our results are trained on the noisy raw depth images captured by the Kinect2 camera, however, our network is still able to give polished results.

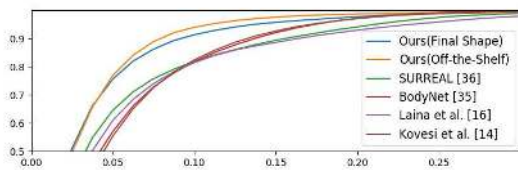


Figure 6. Cumulative Distribution Function of depth error of our method and comparison methods [36, 35, 16].

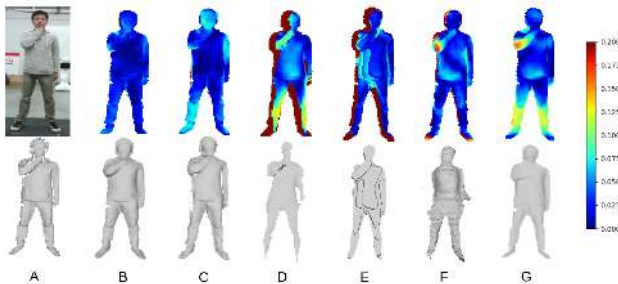


Figure 7. Qualitative comparisons. The first row shows the heatmaps for depth errors, while the second row shows the recovered mesh. Left to right columns: A. Ground truth, B. Ours (Final Shape), C. Ours (Off-the-Shelf), D. SURREAL [36], E. BodyNet [35], F. Laina et al. [16] and G. Kovesi et al. [14] respectively.

estimation and depth discontinuities. We also use the Mean Absolute Error (MAE) as a more global metric to prove that our method captures not only details but also overall shapes. Furthermore, we plot the Cumulative Distribution Function (CDF) of the shape errors by different methods in Figure 6, which illustrates that our method outperforms others with different shape scales.

Figure 7 shows a more intuitive visualization for the comparison with [36], [35], [16] and [14]. At the first row we show the heatmaps for depth errors. The method in SURREAL [36] produces incorrect human body segmentation, which leads to large errors at the boundary. The BodyNet [35] has significant quantization errors due to the coarse volumetric representation. [16] generates very rough depth maps with large structure error because of lacking intermediate supervision of 3D joints and segmentation. The result of [14] shows it can not handle depth discontinuous cases such as when putting hands in front of the torso.

7.2. Ablation Studies

In this section, we verify the effectiveness of the individual components of our method. To this end, we trained another 5 networks in the following settings and compared their results with ours.

Without Skeleton and Segmentation Cues: We discard Skeleton-Net and Segmentation-Net and only feed RGB image to Depth-Net to predict human body depth while the other conditions keep the same.

Without Depth Separation: We replace the two-branch architecture of the Depth-Net with only one branch. We train this network for the same epochs with the Huber loss defined as:

$$L = H(D_{pred} - D_{gt}, \alpha_4). \quad (9)$$

where α_4 is set as 0.20 meters in this setting.

Only Stage 1 Training: We keep the two-branches architecture and trained it only on stage 1 for the same total epochs.

Only Stage 2 Training: The network is the same, while we train it directly on stage 2 without well-initialized weight of

Methods	Accuracy			MAE
	1.25cm	2.5cm	5.0cm	
Ours	29.24	50.93	75.52	3.282
W/o seg&skeleton	25.74	46.19	71.04	4.382
W/o separation	28.00	49.42	72.97	3.480
Only stage 1	26.64	48.14	72.61	3.592
Only stage 2	27.89	50.31	74.87	3.332
W/o truncated loss	28.03	49.84	74.23	3.410

Table 2. Performance of depth prediction of our method and other five settings in Section 7.2

the base and detail branches.

Huber Loss on Composed Shape: We follow the two stages training strategy on the same network but use Huber loss instead of truncated L1 loss to define the composed Loss in stage 2:

$$L_{composed} = H(D_{base} + D_{detail} - D_{gt}, \alpha_5). \quad (10)$$

where α_5 is 0.20m in this setting.

We tested the five different settings mentioned above. Figure 8, 9, 10, 11, 12, show some qualitative comparisons of the results from these different settings. Specifically, Figure 8 shows that in the setting without Segmentation-Net and Skeleton-Net, the Depth-Net will lose the high-level human body information such as 3D joints and body part segmentation, hence the results show some structural issues, like broken meshes on some examples. Figure 9 clearly demonstrates that the network without a two-branch architecture is not able to recover small-scale geometry details. From the results of Figure 10 and Figure 11 we can see that the recovered surface under these two settings are very coarse. Because without using truncated L1 loss which clips the composed error in stage 2 to improve the consistency of two branches, the large layout error may overwhelm the detail error and leads to unstable results from two branches. Figure 12 shows without stage 1 guiding two branches focusing on their target distribution, the detail branch is not working on recovering the small wrinkle specifically. In summary, it is clear that our method produces the best shape details, main layout and smooth surface, which demonstrates the effectiveness of separating the base shape and detail shape and two-stage training with truncated L1 loss on the composed shape.

Table 2 further provides the quantitative comparison of these settings on our testing data with fused ground truth depth [12]. Note that compared results are the composed shape without refinement. The proposed method consistently outperforms the other settings.

7.3. Qualitative Results

To demonstrate our network can be generalized to unconstrained data, Figure 13 shows our results on some unconstrained Internet images. Our method also successfully

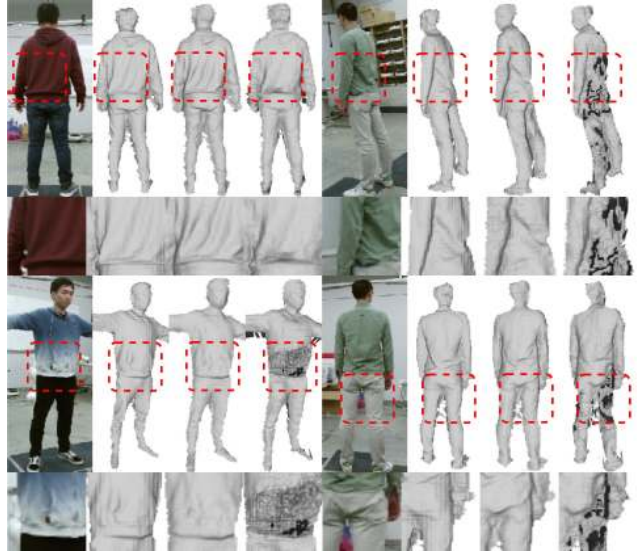


Figure 8. Comparison of our proposed method and ‘W/o Skeleton and Segmentation Cues’. From left to right, they are the image, ground truth and the results from our method and the setting without Segmentation-Net and Skeleton-Net cues. It is clear that without high-level information to guide the depth estimation, the result might have large shape errors.

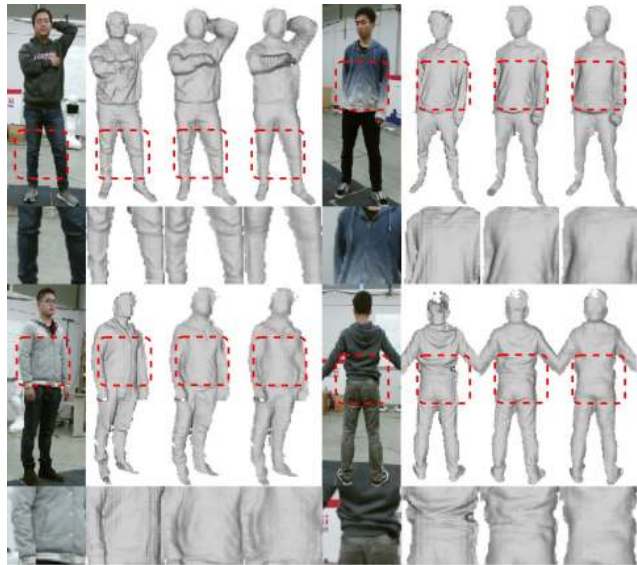


Figure 9. Comparison of our proposed method and ‘No Depth Separation’. From left to right, they are the image, ground truth and the result from our method and the setting with only one depth branch. We can see that the results without a two-branch architecture are rough and do not have many geometry details.

recovers certain shape details on these images. We further visualize the estimated surface normal map, which encodes the cloth wrinkles.

In our demo video, we demonstrate the performance of

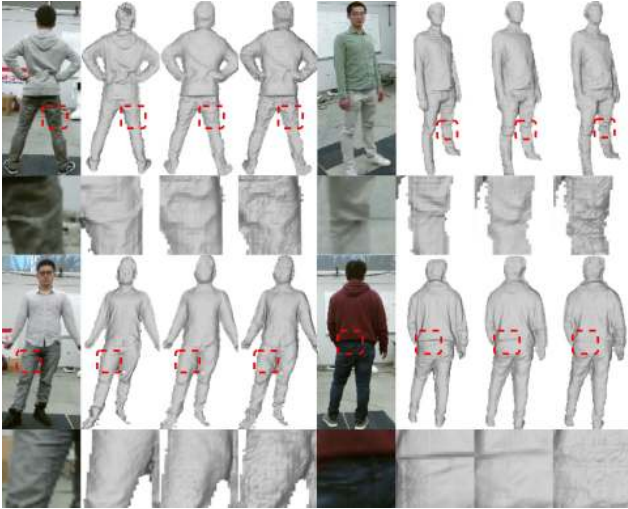


Figure 10. Comparison of our proposed method and ‘Only Stage 1 Training’. From left to right, they are the image, ground truth and the result from our method and the setting with only stage 1 training. From the surface we can see the results without stage 2 will generate wrong wrinkles on the clothes.

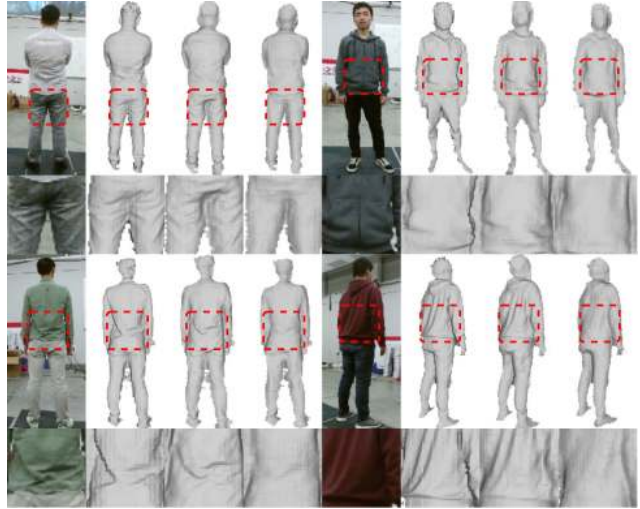


Figure 12. Comparison of our proposed method and ‘Only Stage 2 Training’. From left to right, they are the image, ground truth and the result from our method and the setting with only stage 1 training. By zooming in the results we can see the setting without stage 1 lose majority of geometry details.

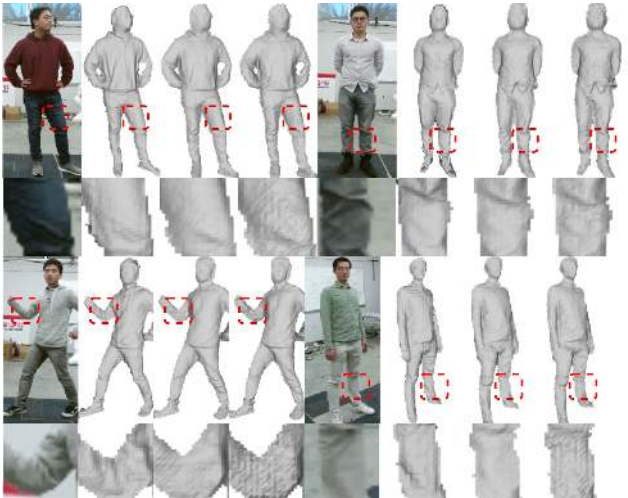


Figure 11. Comparison of our proposed method and ‘Huber Loss on Composed Shape’. From left to right, they are the image, ground truth and the result from our method and the setting with Huber loss on the composed depth. We can see the results without using our truncated L1 loss are unstable and not smooth enough.

our method on some video clips, which are processed in a frame-by-frame fashion. The result shows that our method can even generate temporally coherent results without explicitly modeling it.

8. Conclusion

This paper proposes a neural network to estimate a detailed depth map for the human body in a single input RGB



Figure 13. Some results on unconstrained online images. From left to right, for each example, we show the input image, estimated surface normal and final result.

image. The recovered result can capture fine cloth wrinkles and produce temporally coherent depths for video inputs. It might be used in visualization applications such as the Microsoft Holoportation. This result is achieved by separating and estimating the base shape and detail shape respectively with a novel truncated L1 loss. We also introduce a novel parameter free shape refinement layer to further improve the final result with surface normals. Quantitative evaluation on lab data and qualitative examples on unconstrained Internet data demonstrate the success of the proposed method.

References

- [1] 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
- [2] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018. 1, 2
- [3] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: Shape completion and animation of people. In *ACM Trans. on Graph.*, volume 24, pages 408–416. ACM, 2005. 1, 2
- [4] Alexandru O Balan, Leonid Sigal, Michael J Black, James E Davis, and Horst W Haussecker. Detailed human shape and pose from images. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, Minneapolis, June 2007. 1
- [5] Jan Bednarik, Pascal Fua, and Mathieu Salzmann. Learning to reconstruct texture-less deformable surfaces from a single view. In *2018 International Conference on 3D Vision (3DV)*, pages 606–615. IEEE, 2018. 2, 5
- [6] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Blender Institute, Amsterdam, 5
- [7] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it SMPL: Automatic estimation of 3d human pose and shape from a single image. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 561–578. Springer, 2016. 1, 2
- [8] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4D: Real-time performance capture of challenging scenes. *ACM Trans. on Graph.*, 35(4):114, 2016. 1
- [9] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. 2
- [10] Peng Guan, Alexander Weiss, Alexandru O. Balan, and Michael Black. Estimating human shape and pose from a single image. In *Proc. of International Conference on Computer Vision (ICCV)*, pages 1381–1388, 09 2009. 1
- [11] Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. Neural networks for machine learning, lecture 6a overview of mini-batch gradient descent. 5
- [12] Olaf Kähler, Victor Adrian Prisacariu, Carl Yuheng Ren, Xin Sun, Philip Torr, and David Murray. Very High Frame Rate Volumetric Integration of Depth Images on Mobile Device. *IEEE Transactions on Visualization and Computer Graphics*, 22(11), 2015. 5, 7
- [13] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [14] Peter Kovési. Shapelets correlated with surface normals produce surfaces. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 994–1001. IEEE, 2005. 5, 6
- [15] Arun CS Kumar, Suchendra M Bhandarkar, and P Mukta. DepthNet: A recurrent neural network architecture for monocular depth prediction. In *1st International Workshop on Deep Learning for Visual SLAM, (CVPR)*, volume 2, 2018. 2
- [16] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3D Vision (3DV)*. IEEE, 2016. 2, 5, 6
- [17] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, and Peter V. Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1, 2
- [18] Sijin Li, Weichen Zhang, and Antoni B Chan. Maximum-margin structured learning with deep networks for 3D human pose estimation. In *Proc. of International Conference on Computer Vision (ICCV)*, pages 2848–2856, 2015. 1
- [19] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *Proc. of SIGGRAPH (ACM Trans. on Graph.)*, 34(6):248:1–248:16, Oct. 2015. 1, 2
- [20] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5667–5675, 2018. 2
- [21] Julieta Martinez, Rayat Hossain, Javier Romero, and James J. Little. A simple yet effective baseline for 3D human pose estimation. In *Proc. of International Conference on Computer Vision (ICCV)*, 2017. 1
- [22] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. VNect: Real-time 3d human pose estimation with a single rgb camera. volume 36, 2017. 1
- [23] Diego Nehab, Szymon Rusinkiewicz, James Davis, and Ravi Ramamoorthi. Efficiently combining positions and normals for precise 3D geometry. *ACM Trans. on Graph.*, 24(3), Aug. 2005. 1
- [24] Diego Nehab, Szymon Rusinkiewicz, James Davis, and Ravi Ramamoorthi. Efficiently combining positions and normals for precise 3d geometry. In *ACM Trans. on Graph.*, volume 24, pages 536–543. ACM, 2005. 4
- [25] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 483–499. Springer, 2016. 1, 2
- [26] Gabriel L Oliveira, Abhinav Valada, Claas Bollen, Wolfram Burgard, and Thomas Brox. Deep learning for human part discovery in images. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1634–1641. IEEE, 2016. 2

- [27] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *2018 International Conference on 3D Vision (3DV)*, pages 484–494. IEEE, 2018. 1, 2
- [28] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [29] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7025–7034, 2017. 2
- [30] Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, and Jiaya Jia. GeoNet: Geometric neural network for joint depth and surface normal estimation. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 283–291, 2018. 4, 5
- [31] Helge Rhodin, Nadia Robertini, Dan Casas, Christian Richardt, Hans-Peter Seidel, and Christian Theobalt. General automatic human shape and motion capture using volumetric contour cues. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 509–526. Springer, 2016. 1, 2
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 3, 5
- [33] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 529–545, 2018. 1, 2, 3
- [34] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 1
- [35] Gül Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. BodyNet: Volumetric inference of 3d human body shapes. In *Proc. of European Conference on Computer Vision (ECCV)*, 2018. 1, 2, 5, 6
- [36] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 5, 6
- [37] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2022–2030, 2018. 2
- [38] Zhichao Yin and Jianping Shi. GeoNet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, 2018. 2
- [39] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 340–349, 2018. 2
- [40] Yinda Zhang and Thomas Funkhouser. Deep depth completion of a single rgb-d image. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1, 4
- [41] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2