

DenseRaC: Joint 3D Pose and Shape Estimation by Dense Render-and-Compare

Yuanlu Xu^{1,2}

Song-Chun Zhu²

Tony Tung¹

¹Facebook Reality Labs, Sausalito, USA

²University of California, Los Angeles, USA

merayxu@gmail.com, sczhu@stat.ucla.edu, tony.tung@fb.com

Abstract

We present DenseRaC, a novel end-to-end framework for jointly estimating 3D human pose and body shape from a monocular RGB image. Our two-step framework takes the body pixel-to-surface correspondence map (i.e., IUUV map) as proxy representation and then performs estimation of parameterized human pose and shape. Specifically, given an estimated IUUV map, we develop a deep neural network optimizing 3D body reconstruction losses and further integrating a render-and-compare scheme to minimize differences between the input and the rendered output, i.e., dense body landmarks, body part masks, and adversarial priors. To boost learning, we further construct a large-scale synthetic dataset (MOCA) utilizing web-crawled Mocap sequences, 3D scans and animations. The generated data covers diversified camera views, human actions and body shapes, and is paired with full ground truth. Our model jointly learns to represent the 3D human body from hybrid datasets, mitigating the problem of unpaired training data. Our experiments show that DenseRaC obtains superior performance against state of the art on public benchmarks of various human-related tasks.

1. Introduction

Though much progress has been made in human pose estimation, body segmentation and action recognition, it remains underexplored to leverage such estimations into the 3D world, due to the difficulty in data acquisition, ambiguities from monocular inputs and nuisances in natural images (e.g., illumination, occlusion, texture). Existing learning-based methods [22, 39, 55] heavily rely on sparse 2D/3D landmarks (i.e., skeleton joints), body part masks or silhouettes. However, it is ambiguous to recover 3D human pose and body shape from such limited information.

In this paper we propose DenseRaC, a new framework for 3D human pose and body shape estimation from monocular RGB image, as illustrated in Fig. 2:

- The task is solved in a two-step framework, first by estimating pixel-to-surface correspondences (i.e., IUUV images) from the RGB inputs, and then by leveraging the estimated IUUV images into 3D human pose and body shape.
- A parametric human pose and body shape representation




Figure 1. DenseRaC estimates 3D human poses and body shapes given people-in-the-wild images. The proposed framework handles scenarios with multiple people, all genders, and various clothing in real time. Here, we show results on Internet images [1].

is integrated into the forward pass and backward propagation, inspired by recent work [22, 39].

- An IUUV image based dense render-and-compare scheme is incorporated into the framework. We minimize 3D reconstruction errors as well as discrepancies between inputs and rendered images from estimated outputs.

We learn the proposed model with both unpaired and paired data, compatible with different levels of supervisions. The end-to-end training minimizes multiple losses defined upon human pose and body shape jointly, including parameter regression, 3D reconstruction, landmark reprojection, body part segmentation, as well as adversarial loss on impossible configurations (see Sec. 3.3).

To boost learning, we further construct a large scale synthetic dataset covering diversified human poses and body shapes. The synthetic data is generated using web-crawled 3D animations and scanned all-gender body shapes for human studies, and rendered from various camera views (see Sec. 4). Learning from synthetic data mitigates the problem of unpaired, partial paired, or inaccurately annotated training data in popular public people-in-the-wild and Mocap benchmarks, as well as improves the model robustness against varied camera views and occlusions.

In our experiments, we evaluate DenseRaC on three tasks: 3D pose estimation, semantic body segmentation and 3D body reconstruction. Qualitative and quantitative experimental results show DenseRaC outperforms existing methods on both public benchmarks and the newly proposed synthetic dataset (see Sec. 5).

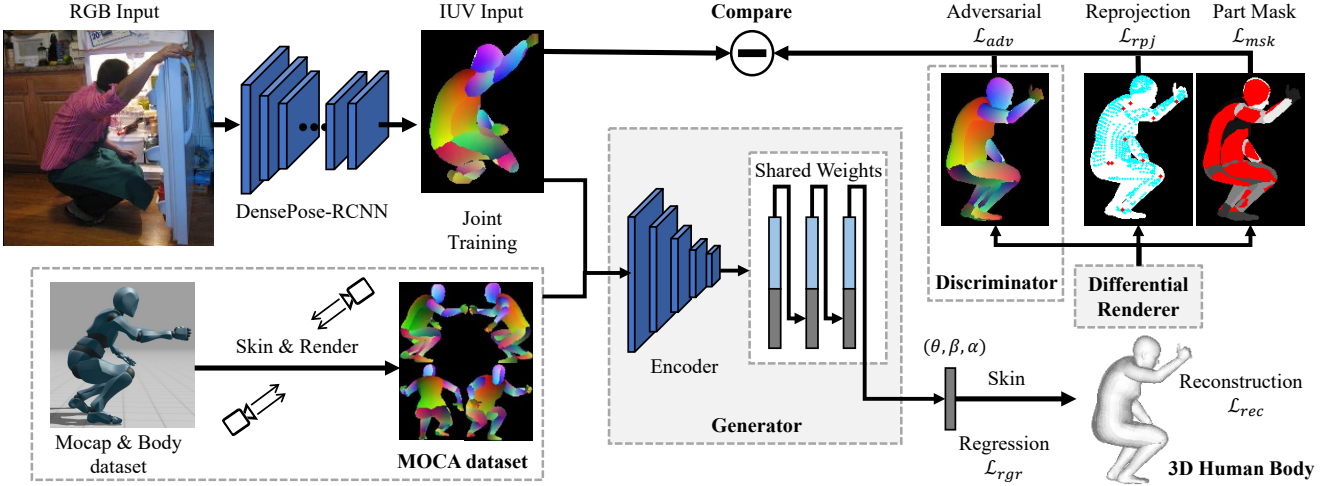


Figure 2. Illustration of DenseRaC. Our two-step framework uses pixel-to-surface correspondences of human body as the intermediate representation, fed with data sources either from estimations on realistic images through DensePose-RCNN or rendered images on synthetic 3D humans. Given IUV images, we develop a deep neural network conducting parametric pose and shape regression and a differentiable renderer performing render-and-compare. The proposed framework optimizes losses of 3D reconstruction and discrepancies between inputs and rendered outputs by end-to-end learning.

To the best of our knowledge, this is the first end-to-end framework introducing a pixel-to-surface correspondence map as the intermediate representation and a corresponding dense render-and-compare scheme for learning 3D human pose and body shapes. We believe DenseRaC shows a great potential for numerous real-world applications in surveillance, entertainment, AR/VR, etc. Some featured results are shown in Fig. 1.

2. Related Work

The proposed method is mainly related to researches in three fields.

Monocular 3D pose estimation is a longstanding problem in computer vision. Current approaches train deep networks from large-scale training sets to regress 3D human joint transformations [18, 27]. Deep neural net architectures enable direct body location with pose prediction, which is an advantage compared to traditional model-based methods that require good initialization [4, 26]. Several methods predict 3D pose directly given monocular data [52, 41, 50, 38, 32, 16, 19, 47]. On the other hand, many approaches lift 2D human poses [8, 5], used as intermediate representation, and learn a model for 2D-3D pose space mapping [61, 63, 62, 34, 9]. State of the art in this track obtains fascinating performance on popular benchmarks limited to laboratory instrumented environments, and yet shows unsatisfactory results on in-the-wild images. Another common issue is that most existing methods do not incorporate a physically plausible human skeleton model and lack constraints on the estimated results, which results in extra post-processings for graphics related applications.

3D human body reconstruction aims at recovering full 3D meshes of the human body from single RGB images or video sequences, rather than major 3D skeleton joints. For example, Zuffi *et al.* [64] integrated both realistic body

model and part-based graphical models [58, 57, 59] for jointly emphasizing graphics-like models of human body shape and part-based human pose inference. In [30, 4, 26, 53], a skinned body model (SMPL) is used to formulate body shape as a linear function of deformation basis (*i.e.*, with blend shapes). In [51, 42, 22, 39], SMPL is considered as the parametric representation of 3D human body and DNNs are developed to estimate such parameters end-to-end. Guler *et al.* [13, 12] build a FCN for human shape estimation by learning dense image-to-template correspondences. Other work [7, 55, 20] focuses on reconstructing 3D body shapes using RGB or RGBD images and not directly estimates 3D human pose and body shapes. These approaches are also suitable for multiple-view video capture setup [35, 54]. In this paper, we use a SMPL variant as the parametric representation of 3D human body and further develop a pixel-to-surface dense correspondence based render-and-compare framework.

Learning from synthetic humans. Modeling 3D humans in arbitrary scenes requires representative training sets. A number of previous work has considered automatically generating data for assisting 3D models, *e.g.*, upper body [40], full-body silhouettes [2]. [14] artificially renders pedestrians in a scene while leveraging camera parameters and geometrical layout, and further trains a scene-specific pedestrian detector. In [44], real 2D pose samples are reshaped by adding small perturbations, and augmented with different backgrounds. Rogez *et al.* [49], for a given 3D pose, combines local image patches from several images with kinematic constraints to create a new synthetic image. Rahmani *et al.* [46] fits synthetic 3D human models to Mocap skeletons and renders human poses from numerous virtual viewpoints. Varol *et al.* [56] also generate a synthetic human body dataset with random factors (*e.g.*, pose, shape, texture, background, *etc.*). These datasets cannot solely serve




Figure 3. Illustration of mapping from pixel to 3D surface. Our framework estimates an IUV image and dense 3D landmarks from an RGB input, whose pixels refer to 3D points on the body model.

to train models generalized to real data, due the gap between synthesized and realistic images. In this paper, we propose to use pixel-to-surface correspondence maps to bridge the gap. The joint training on hybrid datasets is proved to be effective in improving performance on realistic data. To our best knowledge, we are the first to address joint human pose and body shape estimation using such training modalities.

3. DenseRaC Framework

As illustrated in Fig. 2, the proposed framework estimates 3D human poses and body shapes in two steps: first obtaining pixel-to-surface correspondences (*i.e.*, IUV images) and then leveraging the intermediate results IUV images into 3D surfaces. There are two sources of IUV inputs: i) estimations from RGB inputs using a pre-trained DensePose-RCNN [12], and ii) rendered IUV images from synthetic data.

Our framework employs a compact and expressive 3D human body model, which is parameterized by 3D human pose $\theta \in \mathbb{R}^{58 \times 3}$, body shape $\beta \in \mathbb{R}^{50}$, instead of directly estimating 3D point clouds, voxels or depth maps. The 3D human pose is represented as a tree structure, with 58 relative 3D rotations between parent and child joints while the body shape is represented by 50 shape coefficients, as elaborated in Sec. 3.5.

3.1. Network Architecture

Given IUV inputs, we design a network architecture consisting of three modules:

- A generator with a back-boned base network (*i.e.*, ResNet-50 [15]) to extract expressive feature maps and a regressor which takes the stretched feature maps (*i.e.*, 2048D feature vector) from the base network as inputs and estimates 3D human body parameters $[\theta, \beta]$ and camera parameters $\alpha \in \mathbb{R}^3$ (*i.e.*, 227D concatenated vector). The camera model is assumed to be an orthographic projection, parameterized by scale factor f and camera axis (x, y) . The regressor is composed of 3 fully connected layers with 1024 nodes each. Inspired by [22], we consider the regressor to model an iterative update $\Delta_{\theta, \beta, \alpha}$ to the final output, starting from the parameter mean $[\bar{\theta}, \bar{\beta}, \bar{\alpha}]$. The weights are shared across all three layers, simulating the recursive tree structure within 3D human pose.
- A differentiable renderer creates 2D projections of the reconstructed 3D human body mesh, using the estimated camera parameters (see Sec. 3.3). We implement a differ-

ential rasterizer which creates an IUV image suitable for gradient flow. Following a render-and-compare scheme, we define three losses to measure and minimize the differences between the input IUV image and the rendered IUV image from our model output.

- A discriminator to constrain impossible configurations for unpaired data. We design two shallow networks with two fully connected layers as a discriminator. One is used for discriminating 3D human poses and the other one for body shapes. The number of nodes for pose and shape in sub-networks are set to 512 and 64, respectively.

3.2. IUV as Proxy Representation

As illustrated in Fig. 3, we utilize the IUV image as a proxy representation. An IUV map, similarly to UV map in graphics, defines pixel-to-surface correspondences (one-to-one), from 2D image to 3D surface mesh. Each pixel of an IUV image refers to a body part index I , and (U, V) coordinates that map to a unique point on the body model surface (see Sec. 3.5).

As also discussed in [39], RGB input contains much more information of the human target than 2D joints, silhouettes, or body part masks that are traditionally used as proxy representation. However information such as appearance, illumination or clothing may not be relevant for inferring the 3D geometry, and even overfits the network to nuisance factors. Similar to [39], we also observe that explicit body part representations are more useful for the task of 3D human pose and body shape estimation, compared to RGB images and plain silhouettes. Better part segmentation produces better 3D reconstruction accuracy, while providing full spatial coverage of the person (compared to joint heatmaps). While further increasing the number of segmentation parts above a certain threshold (12) only incrementally leverage 3D pose prediction accuracy, it nevertheless greatly improves body shape estimation (see Sec. 5). We argue that prior work only estimates average body shape.

Note we further use two sources of IUV images as inputs, *i.e.*, IUV images from realistic images estimated from [12] and IUV images from virtual humans synthesized by our renderer (see Sec. 3.3). The IUV estimation could be obtained by other off-the-shelf models or two-stage/end-to-end training. Both inputs go through our neural network model and are used to estimate 3D human pose and body shape parameters. Thus, there are several benefits for using IUV image representation: i) improving robustness against nuisances of light and texture in natural images, ii) providing richer geometry information on 3D human body (by including body part masks and dense landmarks), iii) unifying realistic and synthetic data for joint learning.

3.3. Dense Render-and-Compare

In this paper, 3D human pose and body shape are represented compactly by a parametric model (see Sec. 3.5). Parametrized 3D human body is inferred and fit to the input image, given also camera parameters. Human body surface is represented as a 3D triangular mesh, and body posing is

obtained by standard linear blend skinning. To fully compare a reconstructed 3D human body to a 2D observation of it, we integrate a differentiable renderer, *i.e.*, a computer graphics technique that creates a 2D image from a 3D object using differentiable operations [31, 23], and develop an end-to-end weakly-supervised training scheme.

Rendering consists of projecting 3D vertices of a mesh onto a 2D image plane and rasterizing it (*i.e.*, sampling the faces). 3D-to-2D projection is obtained by a combination of differentiable transformations [33]. Rasterization is a discrete operation that requires gradient definition to allow back-propagation in a neural network. In [31], the authors approximate derivatives at occlusion boundaries which are discontinuous, while colors are interpolated between vertices (*i.e.*, there is no differentiation with respect to texture). In [23], the authors obtain approximate gradients by blurring image to avoid sudden pixel color change. This produces non-zero gradients and enables gradient-flow between pixel (color) value to vertex position. However, lighting and material properties in natural images are complex to model and integrate into neural networks.

On the contrary, our IUUV representation is invariant to background, lighting conditions and surface texture like clothing (see Sec. 3.2). In addition, UV values on each body part I are continuous with respect to neighbor pixels (see Fig. 3). This actually allows to naturally compute gradients on mesh surface and at boundaries and back-propagate them through network layers.

Our renderer creates IUUV image comparable to the generated output of [12] (see Fig. 4). Self-occlusion is handled by depth buffering. Our rasterizer draws only the surface faces closest to the camera (and facing it) at each pixel. During back propagation, we only pass gradient flows of pixels corresponding to visible regions.

Different from [53, 24] where render-and-compare losses are computed upon silhouettes and 2D depth maps, we compute dense render-and-compare losses \mathcal{L}_{rac} using IUUV values between ground-truth IUUV images and rendered ones (see Sec. 3.4). The differentiable renderer (including IUUV rasterizer) and losses are implemented with differentiable operations using a neural net framework with automatic differentiation [6, 53, 24].

3.4. Loss Terms

Our model integrates a dense render-and-compare module with corresponding loss computations in the backward propagation, hence leveraging previous methods [42, 22, 39, 55]. The loss function is defined as

$$\mathcal{L} = \mathcal{L}_{rac} + 1\mathcal{L}_{rec} + 1\mathcal{L}_{rgr}, \quad (1)$$

where 1 indicates the existence of such annotation, \mathcal{L}_{rac} , \mathcal{L}_{rec} and \mathcal{L}_{rgr} denote render-and-compare loss, 3D reconstruction loss and parameter regression loss, respectively.

• **Render-and-Compare Loss \mathcal{L}_{rac}** is evaluated under three measurements, that is,

$$\mathcal{L}_{rac} = \mathcal{L}_{rpj} + \mathcal{L}_{msk} + \mathcal{L}_{adv}, \quad (2)$$

where \mathcal{L}_{rpj} , \mathcal{L}_{msk} and \mathcal{L}_{adv} denote landmark reprojection loss, part mask loss and adversarial loss, respectively.

Landmark Reprojection Loss \mathcal{L}_{rpj} measures displacement between ground truth and estimated dense 2D landmarks:

$$\mathcal{L}_{rpj} = \sum_i^N \mathbf{1}_i \|p_i - \hat{p}_i\|_1, \quad (3)$$

where $\mathbf{1}_i$ indicates the visibility (1 if visible, 0 otherwise) for i -th 2D landmark (N in total), $p_i \in \mathbb{R}^2$ and $\hat{p}_i \in \mathbb{R}^2$ represent i -th 2D landmark from ground truth and 3D mesh reprojection, respectively. To correctly localize the landmarks from ground truth (*i.e.*, IUUV image estimated from DensePose [12]), we formulate this problem as a point-to-point greedy match and solve the correspondence problem by k-Nearest Neighbor (k-NN) search. Specifically, we first create a k-D tree for IUUV values of 3D body mesh vertices. For any input IUUV image, we search for 1-NN of each visible pixel and obtain a matched pair with the closest 3D body mesh vertex within a distance threshold τ . Empirically, $\tau \in [0.01, 0.1]$ yields 100-300 matching pairs considered as near-optimal one-to-one 2D/3D dense landmarks correspondences. This serves as a weakly-supervised scaffold to densely fit 3D human body to the re-projected 2D image. Note the matching is computed offline and serves as a pre-processing step on IUUV inputs, as shown in Fig. 5.

Part Mask Loss \mathcal{L}_{msk} provides semantic information for the location of body part:

$$\mathcal{L}_{msk} = \sum_k (1 - \text{IoU}(I_k, \hat{I}_k)), \quad \text{IoU}(I_k, \hat{I}_k) = \frac{|I_k \cap \hat{I}_k|}{|I_k \cup \hat{I}_k|}, \quad (4)$$

where k is body part index and $\text{IoU}(\cdot, \cdot)$ represents intersection over union of two masks. We keep the same body segments (12 parts) I and (U, V) mapping as specified in [12].

Adversarial Loss \mathcal{L}_{adv} constrains configuration plausibility. Unlike [22] using unpaired or Mosh-based [29] weakly-supervised SMPL annotations, we use ground-truth 3D human poses and body shapes from our synthetic dataset, which contains much larger action variations than most Mocap sequences (see Sec. 4). We believe such long-tail poses are crucial for the adversarial loss in finding the decision boundary. Hereby, we account for millions of synthetic samples as both paired ground truth and unpaired adversarial prior for realistic datasets. We follow the GAN loss definitions in [10] and train our generator and discriminator jointly.

• **3D Reconstruction Loss \mathcal{L}_{rec}** measures the deformation of reconstructed 3D human body, compared with ground truth:

$$\mathcal{L}_{rec} = \sum_i^K \|P_i - \hat{P}_i\|_2, \quad (5)$$

where P_i and \hat{P}_i represent 3D keypoint positions from input and generated 3D mesh, respectively.

• **Parameter Regression Loss \mathcal{L}_{rgr}** measures mean square errors between estimated parameters $[\theta, \beta, \alpha]$ and ground truth $[\hat{\theta}, \hat{\beta}, \hat{\alpha}]$:

$$\mathcal{L}_{rgr} = \|[R_\theta, \beta, \alpha] - [R_{\hat{\theta}}, \hat{\beta}, \hat{\alpha}]\|_2, \quad (6)$$




Figure 4. IUUV images from MOCA generated by rasterizing 3D bodies obtained with 3D poses from Mixamo and body shapes from CAESAR. MOCA contains 2M+ images with fully paired ground truth.

where R_θ denotes the rotation matrix of θ . Notably, pose parameters are first transformed in rotation matrices. Losses are computed upon such matrices and gradients are automatically back-propagated. This helps avoid the singularity problem of XYZ-Euler based 3D rotation and requires no extra constraints on the rotation matrix, that is,

$$RR^T = \text{diag}(1, \dots, 1), \quad \det(R) = 1, \quad (7)$$

where $\det(\cdot)$ denotes the matrix determinant.

3.5. Human Body Model

We use a body shape model similar to SMPL [30, 4]. The statistical body model is obtained by PCA on pose-normalized 3D models of real humans, obtained by non-rigid registration of a body template to 3D scans of the CAESAR dataset¹, which represents anthropometric variability of 4,400 men and women. The body template mesh has 7,324 vertices, 14,644 triangular faces and a skeletal rig with body and hand joints.

Our model is trained with all 3D scans in the dataset, resulting in a statistical model that can describe bodies from unseen in-the-wild images regardless of gender. An arbitrary body shape can then be described by a set of shape coefficients (*i.e.*, shape parameters or shape blend shapes) using a linear representation. Truncating shape coefficients to 50 principal components enables reconstruction of all-gender body shapes without noticeable distortions: *e.g.*, the SMPL-Male with 10 coefficients does not reconstruct well female shape (RMSE=9.9mm), while an all-gender model does (RMSE=6.3/3.4mm with 10/50 coeffs respectively).

Considering potential applications in AR/VR, 3D animations and better utilization of annotations, we enrich the standard SMPL 24-joint skeleton with 28 joints for modeling fingers and 5 more joints on spine and head for better flexibility. We further add a root node for global translation and rotation, leading to a skeleton with 58 joints.

4. MOCA Synthetic Dataset

The literature has provided several datasets to evaluate human 3D pose (*e.g.*, H3.6M [18], MPI-INF-3DHP [36]), but only few for joint 3D pose and body shape (*e.g.*, SURREAL [56] and UP-3D [26]). However, SURREAL is dedicated to body segmentation and depth estimation and only has a rough skeleton (24 major body joints), while UP-3D has *weakly-supervised* shapes (from SMPL fitted to LSP and MPII), arguably imprecise [55].

Hence, we propose **MOCA**, a large-scale synthetic dataset with 2,089,104 images containing ground-truth body shapes and 3D poses, as shown in Fig. 4. For various human poses and actions, we seek to a popular collection center of 3D human animations (*i.e.*, Mixamo²), whose sources mainly come from Mocap systems and artist designs. We implement a web crawler to fetch high fidelity animations. Notably, Mixamo supports tuning parameters (*e.g.*, limb length, energy, overdrive) for each action sequence to generate variants. As we observe certain parameter settings may introduce artifacts, we thus keep the default setting for all sequences. We collect a set of 2,446 3D animation sequences with 261,138 frames at 30 fps, covering wide action categories of sports, combat, daily and social activities. We extract a finer 3D skeleton with fingers and facial bones using Maya and re-map those joints onto our body model.

We then generate 2,781 bodies using the 3D scans from CAESAR dataset and compute corresponding (PCA) shape coefficients. By combining 3D pose θ and shape β , we pose body models to specific pose&shape configurations by standard linear blend skinning.

The complete combination of all 3D poses and body shapes produces an enormous amount of 3D human body samples. Currently, we randomly select 8 body shapes for each action sequence. We further add a random camera view for each sequence, and render them as IUUV image sequences using our IUUV rasterizer (see Sec. 3.3), obtaining a dataset with 2,089,104 frames in total and fully paired ground truth of body shape, 3D pose and the camera view. For training/testing set partition, we set the ratio as 90%/10%. We synthesize the training set with the first 2,201 Mixamo action sequences and 2,502 CAESAR body shapes and leave the rest 246 action sequences and 279 body shapes only visible to the testing set.

5. Experiments

We evaluate DenseRaC on several public large-scale benchmarks for three tasks: 3D pose estimation, body shape estimation and body semantic segmentation. We further assess human 3D reconstruction results (*i.e.*, mesh-level reconstruction, joint&shape parameter estimation) on the proposed large synthetic dataset MOCA that contains ground-truth 3D pose and body shape. Our experiments compare favorably to the state of the art. Estimated 3D poses and

¹<http://store.sae.org/caesar/>

²<http://www.mixamo.com>




Figure 5. Pre-processed training samples from public benchmarks. Left: original image, right: estimated IUUV image, ground-truth keypoint annotations (yellow) and dense landmarks (red).

body shapes are stable on videos (see additional materials). Our qualitative results also show natural hand poses (*e.g.*, opened, clenched).

5.1. Datasets

We use five public human benchmarks plus our synthesized MOCA for model training and evaluation, *i.e.*, LSP [21], MPII [3], COCO [28], H3.6M [17, 18] and MPI-INF-3DHP [36]. We adopt standard training/validation/testing partitions on all datasets and calibrate loss terms using cross-validation. When a certain dataset is used for evaluation, all data from other datasets will be used in training.

For all training and testing samples, we crop out each person in the whole image using ground-truth bounding boxes. All samples are resized to ~ 150 - 180 pixel height with preserved aspect ratio, and further adjusted to 224×224 with padding/cropping respectively. We then run IUUV image estimation [12] on all samples. Considering a sample \mathcal{I} may contain multiple people and false alarms, we compute a saliency score $s = \frac{|m|}{\|m_c - \mathcal{I}_c\|_2}$ for each detected person mask m , where m_c and \mathcal{I}_c represent the center of the person mask and the image, respectively. We then pick the person mask with the largest saliency score and suppress the other detection responses.

For the training set, we further run pixel-to-surface matching (as described in Sec. 3.3) to create dense correspondences. We discard samples with less than 200 corresponding pairs, as IUUV image estimation usually failed under such situation. As illustrated in Fig. 5, pre-processing suppresses nuisances in the training samples quite well. During training, all training samples will further be augmented with a random jittering of translation, scaling and reflection to improve the model robustness. We also randomly black out a rectangle image region for the synthetic samples to simulate occlusion in realistic scenarios.

To unite the skeleton structure across all datasets, we use the same 14 joints as in LSP for joint related computation while maintaining our 58-joint skeleton in the backend.

5.2. Implementation Details

In these experiments, the whole framework is implemented with TensorFlow and runs on a DGX workstation

H3.6M	Protocol #1	Protocol #2	Protocol #3
	MPJPE	MPJPE	MPJPE
Martinez <i>et al.</i> (ICCV'17) [34]	62.9	47.7	84.8
Fang <i>et al.</i> (AAAI'18) [9]	60.3	45.7	72.8
Rhodin <i>et al.</i> (CVPR'18) [48]	66.8	-	-
Yang <i>et al.</i> (CVPR'18) [60]	58.6	37.7	-
Hossain <i>et al.</i> (ECCV'18) [16]	51.9	42.0	-
Lassner <i>et al.</i> (CVPR'17) [26]	80.7	-	-
HMR (CVPR'18) [22]	88.0	56.8	77.3
Pavlakos <i>et al.</i> (CVPR'18) [42]	-	75.9	-
NBF (3DV'18) [39]	-	59.9	-
DenseRaC baseline	82.4	53.9	77.0
+ render-and-compare	79.5	51.4	75.9
+ synthetic data	76.8	48.0	74.1

MPI-INF-3DHP	Protocol #1			Protocol #2		
	PCK	AUC	MPJPE	PCK	AUC	MPJPE
Mehta <i>et al.</i> (3DV'17) [36]	75.7	39.3	117.6	-	-	-
Mehta <i>et al.</i> (TOG'17) [37]	76.6	40.4	124.7	83.9	47.3	98.0
HMR (CVPR'18) [22]	72.9	36.5	124.2	86.3	47.8	89.8
DenseRaC baseline	73.1	36.7	123.1	86.8	47.8	88.7
+ render-and-compare	74.7	38.6	124.9	87.5	48.3	86.7
+ synthetic data	76.9	41.1	114.2	89.0	49.1	83.5

Table 1. Quantitative comparisons of mean per joint position error (MPJPE), PCK and AUC between the estimated 3D pose and ground truth on **H3.6M** under *Protocol #1, #2, #3* and **MPI-INF-3DHP** under *Protocol #1, #2*. - indicates results not reported. Lower MPJPE, higher PCK and AUC indicate better performance. Best scores are marked in **bold**.

with 2 Intel E5 CPUs, 512GB memory and 8 Titan V100 GPUs. Data synthesis and pre-processing (*i.e.*, IUUV image estimation) are implemented with multi-gpu data parallelism. The multi-gpu renderer processes around 300 fps and takes 2 days to generate 2 million MOCA samples (total size 2.7TB). Data pre-processing on realistic datasets takes 12 hours to prepare 0.8 million samples.

For learning, only a single GPU is used due to difficulty in gradient transfer and a potential performance drop. We use batch size 128, learning rate 10^{-5} for the generator and 10^{-4} for the discriminator, and Adam as the optimizer. Our full model is jointly trained on all datasets for 30 epochs. Empirically, for one batch, the forward pass takes around 15ms and the backward propagation takes (~ 130 ms) with IUUV image render-and-compare (~ 55 ms) as the overhead. The total training process takes around a week to complete. For inference, IUUV images are first estimated at around 15 fps and then the forward pass of our model is called, taking 120 fps and thus achieves *real time*.

5.3. 3D Pose Estimation

We first evaluate our method for the task of 3D pose estimation on **H3.6M** [18] and **MPI-INF-3DHP** [36] datasets.

For **H3.6M**, we use three evaluation protocols used to measure the performance: i) **Protocol #1** uses 5 subjects (S1, S5, S6, S7 and S8) for training and 2 subjects (S9 and S11) for testing. Sequences are down-sampled to 10 fps and all 4 cameras and trials are used for evaluation. MSE is measured between estimated and ground-truth 3D joints. ii)

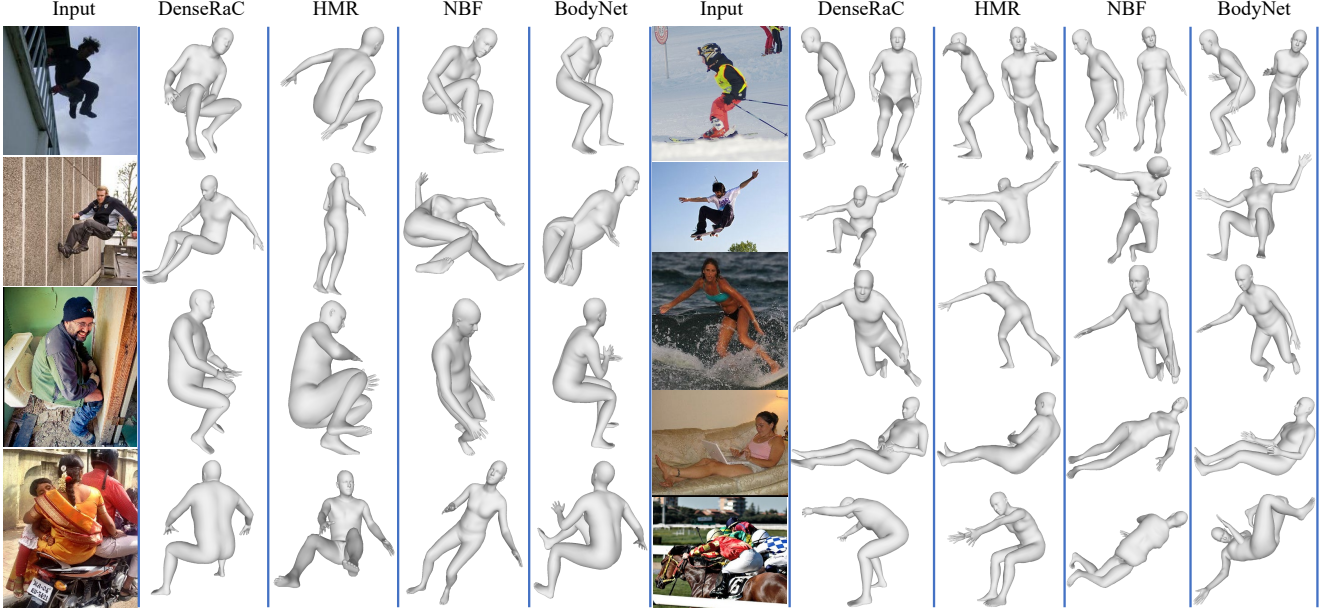


Figure 6. Qualitative comparisons of results estimated from DenseRaC versus state of the art [22, 39, 55]. DenseRaC estimates 3D human poses and body shapes closest to the reality. Note that all examples come from the test set. Best viewed in color.

UP-3D	Body Part		Fg/Bg	
	Accuracy	F1	Accuracy	F1
SMPL on DpCut (ECCV'16) [4]	87.7	0.64	91.9	0.88
SMPL, UP-P91 (ICCV'17) [26]	87.3	0.61	91.0	0.86
HMR (CVPR'18) [22]	87.1	0.60	91.7	0.87
BodyNet (ECCV'18) [55]	-	-	92.8	0.84
DenseRaC	87.9	0.64	92.4	0.88

MOCA	Body Part		Fg/Bg	
	Accuracy	F1	Accuracy	F1
HMR (CVPR'18) [22]	86.6	0.19	92.1	0.60
DenseRaC	89.3	0.27	96.4	0.68

Table 2. Quantitative comparisons of foreground and part segmentation on **UP-3D** and **MOCA** datasets. Accuracy unit is in %. - indicates results not reported. Best scores are marked in **bold**.

Protocol #2 selects the same subjects for training and testing as Protocol #1, while evaluation is only conducted on sequences captured from the frontal camera (*i.e.*, “cam 3”) from trial 1 on all frames. Predictions are post-processed via rigid transformations (*i.e.*, per-frame Procrustes analysis) before comparison. iii) **Protocol #3** uses the same subjects, frame rates and trials for training and testing in Protocol #1 except that camera views are further partitioned. The first three cameras (*i.e.*, “cam 0, 1, 2”) are used for training and the last camera (*i.e.*, “cam 3”) for testing.

For **MPI-INF-3DHP**, we use all sequences from S1-S7 as training set and sequences from S8 as testing set. We regard **Protocol #1** as the default comparison and **Protocol #2** as applying rigid transformations before comparison.

We compare our method with both task-oriented 3D pose state of the art [50, 62, 34, 36, 37, 9, 48, 60, 16] and four parametric body model based estimators [26, 22, 42, 39]. We set up two baselines to validate the effectiveness of two key components in the proposed framework: render-and-compare and joint learning with synthetic data. In

“DenseRaC baseline”, we use SMPL model and the same losses as [22], only switch input sources from RGB images to IUUV images. Variant “+ render-and-compare” denotes adding the proposed dense render-and-compare scheme losses into the framework and part masks. Variant “+ synthetic data” switches to our human body model and further uses augmented synthetic data for joint learning.

As reported in Table 1, we can observe each component in DenseRaC contributes to the final performance and leads DenseRaC to outperform state-of-the-art parametric body model estimators by a large margin. Also notice DenseRaC is comparable with latest task-oriented 3D pose estimators.

5.4. Human Body Segmentation

Given rendered images from outputs, we further employ semantic segmentation as another task to measure how similar the reconstructed 3D human body looks to the person in the input image. We evaluate the tasks of human body segmentation and test our approach on the LSP subset of UP-3D [26] and MOCA datasets. For UP-3D, we post-process our 24 body part masks by merging into the annotated 6 body part masks (*i.e.*, head, torso, left and right leg, and left and right arm) and evaluate on body part and foreground segmentation, while we evaluate both body part segmentation (ignoring 4 subtle body parts, *i.e.*, hands and feet) and foreground segmentation on MOCA. We measure segmentation accuracy and mean F1 score of the results and report metrics and comparisons in Table 5.3. It can be observed that our method achieves comparable or better performance with state of the art [4, 26, 22, 55] on all datasets.

5.5. 3D Human Body Reconstruction

Notice 3D pose estimation and body semantic segmentation are tasks focusing on evaluating partial knowledge of

Methods	MPJPE	MPVPE	MSE $_{\theta,\beta}$
HMR (CVPR'18) [22] unpaired	110.2	-	-
HMR (CVPR'18) [22] paired	91.9	-	-
DenseRaC, \mathcal{L}_{rpj}^J	133.0	174.5	18.227
DenseRaC, $\mathcal{L}_{rpj}^J + \mathcal{L}_{adv}$	131.5	173.6	17.820
DenseRaC, $\mathcal{L}_{rpj}^J + \mathcal{L}_{adv} + \mathcal{L}_{msk}$	122.8	161.5	16.305
DenseRaC, $\mathcal{L}_{rpj}^J + \mathcal{L}_{adv} + \mathcal{L}_{msk}$	107.9	142.3	13.608
DenseRaC, $\mathcal{L}_{rpj}^J + \mathcal{L}_{adv} + \mathcal{L}_{rec}^J$	88.6	121.1	11.901
DenseRaC, $\mathcal{L}_{rpj}^J + \mathcal{L}_{adv} + \mathcal{L}_{msk} + \mathcal{L}_{rec}^J$	86.5	119.8	10.496
DenseRaC, $\mathcal{L}_{rpj}^J + \mathcal{L}_{adv} + \mathcal{L}_{rec}$	82.9	111.0	8.943
DenseRaC, $\mathcal{L}_{rpj}^J + \mathcal{L}_{adv} + \mathcal{L}_{msk} + \mathcal{L}_{rec}$	82.4	110.7	8.722
DenseRaC, $\mathcal{L}_{rpj}^J + \mathcal{L}_{msk} + \mathcal{L}_{rec} + \mathcal{L}_{rgr}$	80.4	105.4	8.164
DenseRaC, full	80.3	105.2	8.151

Table 3. Quantitative comparisons of MPJPE, MPVPE, Pose&Shape Parameter Mean Square Error $MSE_{\theta,\beta}$ on MOCA dataset. Lower values are better. See text for detailed explanations.

the reconstructed 3D human body, We further evaluate the reconstructed 3D human body using two metrics: Mean Per Mesh Vertex Position Error (MPVPE) and regression error on MOCA dataset. These two metrics consider the 3D human body as a whole and provide more guidance about how well the reconstructed 3D human body is. For comparison, we re-train HMR which takes IUUV images as input and uses 2D/3D joint supervisions (*i.e.*, only 14 2D/3D joints in LSP format) and their original unpaired data (Mosh [29] on H3.6M and external Mocap) for the adversarial prior. As reported in Table 3, DenseRaC still significantly outperforms the competitive method.

Ablative Studies. We set up variants of DenseRaC to validate effectiveness of each loss terms. We also define two loss variants \mathcal{L}_{rpj}^J and \mathcal{L}_{rec}^J representing 14-joint-only keypoint reprojection and 3D reconstruction losses, respectively. From the results, we could reach the following conclusions: i) All loss terms contribute to the final performance; ii) Losses used for dense render-and-compare provide richer information than those from sparse joints, greatly reduce impossible 3D body configurations; iii) When task oriented loss terms are given (*i.e.*, \mathcal{L}_{rec} and \mathcal{L}_{rgr}), the contribution from the dense render-and-compare scheme seems to be suppressed, yet such finer supervisions help DenseRaC reach a much better local optimum.

Empirical Studies. We present qualitative results and comparisons to have a better understanding of merits of our method. As shown in Fig. 6, DenseRaC outperforms other competitive methods and reconstructs more plausible and natural 3D human bodies. Notably, HMR, which relies on sparse landmarks, sometimes reconstructs plausible 3D human body appearance, but confuses body front and back. Both NBF and BodyNet are sensitive to occlusions and heavy clothing. When fitting SMPL to such erroneously reconstructed volumes, BodyNet tends to produce highly non-human body shapes³. For all three methods, the estimated human bodies are arguably in an average body shape and insensitive to genders. We also search failure cases on validation set, as shown in Fig. 7. DenseRaC suffers from




Figure 7. Current limitations: heavy occlusions (first row), incorrect IUUV estimations (second row) and under-represented body shapes like children (third row). Each triplet shows the original image, IUUV from [12] (our model input), and our model output.




Figure 8. Comparisons for cascaded and end-to-end frameworks on the application of virtual dressing.

errors in IUUV estimations (*e.g.*, occlusions, long-tail data), and is limited by the orthographic projection assumption and SMPL-based human body representation.

We also explored virtual dressing, namely draping virtual clothing on 3D human body, using our beneath-clothing estimation. As shown in Fig. 1 (top right) and Fig. 8, a cascaded framework for adding physical simulations of clothing is possible [11, 25] and more visually acceptable than end-to-end volumetric reconstruction of BodyNet.

6. Conclusion

We propose DenseRaC, a new end-to-end framework for reconstructing 3D human body from monocular RGB images in the wild. DenseRaC utilizes the pixel-to-surface correspondence map as proxy representation and incorporates a dense render-and-compare scheme to minimize the gap between rendered outputs and inputs. We further boost the model training with large scale synthetic data (MOCA), mitigating the problem of unpaired training data. The proposed framework obtains superior performance and we will explore handling occlusion and interaction (*e.g.*, by multi-view fusion [45], temporal smoothing [43]) next.

Acknowledgements. We would like to thank Tengyu Liu and Elan Markowitz for helping with data collection, Tuur Jan M Stuyck and Aaron Ferguson for cloth simulation, Natalia Neverova and colleagues at FRL, FAIR and UCLA for their support and advice.

³We uses results from 3D skeleton fitting for BodyNet, as volume fitting usually performs much worse.

References

- [1] Images and videos available at *youtube.com*, *onlinedoctor.superdrug.com*, *shutterstock.com*. 1
- [2] Ankur Agarwal and Bill Triggs. Recovering 3D human pose from monocular images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(1):44–58, 2006. 2
- [3] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014. 6
- [4] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, 2016. 2, 5, 7
- [5] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [6] Peter Dayan, Geoffrey Hinton, Radford Neal, and Richard Zemel. The helmholtz machine. *Neural Computing*, 1995. 4
- [7] Endri Dibra, Himanshu Jain, Cengiz Oztireli, Remo Ziegler, and Markus Gross. Hs-nets: Estimating human body shape from silhouettes with convolutional neural networks. In *International Conference on 3D Vision*, 2016. 2
- [8] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *IEEE International Conference on Computer Vision*, 2017. 2
- [9] Hao-Shu Fang, Yuanlu Xu, Wenguan Wang, Xiaobai Liu, and Song-Chun Zhu. Learning pose grammar to encode human body configuration for 3d pose estimation. In *AAAI Conference on Artificial Intelligence*, 2018. 2, 6, 7
- [10] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Annual Conference on Neural Information Processing Systems*, 2014. 4
- [11] Peng Guan, Loretta Reiss, David A. Hirshberg, Alexander Weiss, and Michael J. Black. Drape: Dressing any person. In *ACM SIGGRAPH*, 2012. 8
- [12] Riza Alp Guler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2, 3, 4, 6, 8
- [13] Riza Alp Guler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, and Iasonas Kokkinos. Densereg: Fully convolutional dense shape regression in-the-wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [14] Hironori Hattori, Vishnu Naresh Boddeti, Kris M Kitani, and Takeo Kanade. Learning scene-specific pedestrian detectors without real data. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [15] Kaifeng He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 3
- [16] Mir Rayat Imtiaz Hossain and James J. Little. Exploiting temporal information for 3d human pose estimation. In *European Conference on Computer Vision*, 2018. 2, 6, 7
- [17] Catalin Ionescu, Fuxin Li, and Cristian Sminchisescu. Latent structured models for human pose estimation. In *IEEE International Conference on Computer Vision*, 2011. 6
- [18] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014. 2, 5, 6
- [19] Umar Iqbal, Pavlo Molchanov, Thomas Breuel, Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5d heatmap regression. In *European Conference on Computer Vision*, 2018. 2
- [20] Zhongping Ji, Xiao Qi, Yigang Wang, Gang Xu, Peng Du, and Qing Wu. Shape-from-mask: A deep learning based human body shape reconstruction from binary mask images. *arXiv preprint arXiv:1806.08485*, 2018. 2
- [21] Sam Johnson and Mark Everingham. Learning effective human pose estimation from inaccurate annotation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2011. 6
- [22] Angjoo Kanazawa, Michael J. Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2, 3, 4, 6, 7, 8
- [23] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 4
- [24] Abhijit Kundu, Yin Li, and James Rehg. 3d-rcnn: Instance-level 3d object reconstruction via render-and-compare. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 4
- [25] Zorah Laehner, Daniel Cremers, and Tony Tung. Deepwinkles: Accurate and realistic clothing modeling. In *European Conference on Computer Vision*, 2018. 8
- [26] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2, 5, 6, 7
- [27] Sijin Li, Weichen Zhang, and Antoni B Chan. Maximum-margin structured learning with deep networks for 3d human pose estimation. In *IEEE International Conference on Computer Vision*, 2015. 2
- [28] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollar, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 2014. 6
- [29] Matthew Loper, Naureen Mahmood, and Michael Black. Mosh: Motion and shape capture from sparse markers. In *SIGGRAPH Asia*, 2014. 4, 8
- [30] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics*, 34(6):248, 2015. 2, 5
- [31] Matthew M. Loper and Michael J. Black. Opendr: An approximate differentiable renderer. In *European Conference on Computer Vision*, 2014. 4
- [32] Diogo C. Luvizon, David Picard, and Hedi Tabia. 2d/3d pose estimation and action recognition using multitask deep learn-

- ing. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [33] Steve Marschner and Peter Shirley. Fundamentals of computer graphics. In *CRC Press*, 2015. 4
- [34] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *IEEE International Conference on Computer Vision*, 2017. 2, 6, 7
- [35] Takashi Matsuyama, Shohei Nobuhara, Takeshi Takai, and Tony Tung. 3d video and its applications. In *Springer*, 2012. 2
- [36] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *International Conference on 3D Vision*, 2017. 5, 6, 7
- [37] 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. In *ACM Transactions on Graphics*, 2017. 6, 7
- [38] Bruce Xiaohan Nie, Ping Wei, and Song-Chun Zhu. Monocular 3d human pose estimation by predicting depth on joints. In *IEEE International Conference on Computer Vision*, 2017. 2
- [39] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter V. Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model-based human pose and shape estimation. In *International Conference on 3D Vision*, 2018. 1, 2, 3, 4, 6, 7
- [40] Gregory Shakhnarovich Paul, Paul Viola, and Trevor Darrell. Fast pose estimation with parameter-sensitive hashing. In *IEEE International Conference on Computer Vision*, 2003. 2
- [41] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [42] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3d human pose and shape from a single color image. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2, 4, 6, 7
- [43] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 8
- [44] Leonid Pishchulin, Arjun Jain, Mykhaylo Andriluka, Thorsten Thormahlen, and Bernt Schiele. Articulated people detection and pose estimation: Reshaping the future. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012. 2
- [45] Hang Qi, Yuanlu Xu, Tao Yuan, Tianfu Wu, and Song-Chun Zhu. Scene-centric joint parsing of cross-view videos. In *AAAI Conference on Artificial Intelligence*, 2018. 8
- [46] Hossein Rahmani and Ajmal Mian. 3d action recognition from novel viewpoints. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [47] Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised geometry-aware representation for 3d human pose estimation. In *European Conference on Computer Vision*, 2018. 2
- [48] Helge Rhodin, Jorg Sporri, Isinsu Katircioglu, Victor Constantin, Frederic Meyer, Erich Muller, Mathieu Salzmann, and Pascal Fua. Learning monocular 3d human pose estimation from multi-view images. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 6, 7
- [49] Grégory Rogez and Cordelia Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. In *Annual Conference on Neural Information Processing Systems*, 2016. 2
- [50] Xiao Sun, Jiaxiang Shang, Shuang Liang, and Yichen Wei. Compositional human pose regression. In *IEEE International Conference on Computer Vision*, 2017. 2, 7
- [51] Vince Tan, Ignas Budvytis, and Roberto Cipolla. Indirect deep structured learning for 3d human body shape and pose prediction. In *British Machine Vision Conference*, 2017. 2
- [52] Bugra Tekin, Isinsu Katircioglu, Mathieu Salzmann, Vincent Lepetit, and Pascal Fua. Structured prediction of 3D human pose with deep neural networks. In *British Machine Vision Conference*, 2016. 2
- [53] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *NIPS*, 2017. 2, 4
- [54] Tony Tung, Shohei Nobuhara, and Takashi Matsuyama. Complete multi-view reconstruction of dynamic scenes from probabilistic fusion of narrow and wide baseline stereo. In *IEEE International Conference on Computer Vision*, 2009. 2
- [55] Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *European Conference on Computer Vision*, 2018. 1, 2, 4, 5, 7
- [56] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2, 5
- [57] Wenguan Wang, Yuanlu Xu, Jianbing Shen, and Song-Chun Zhu. Attentive fashion grammar network for fashion landmark detection and clothing category classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [58] Yuanlu Xu, Liang Lin, Wei-Shi Zheng, and Xiaobai Liu. Human re-identification by matching compositional template with cluster sampling. In *IEEE International Conference on Computer Vision*, 2013. 2
- [59] Yuanlu Xu, Lei Qin, Xiaobai Liu, Jianwen Xie, and Song-Chun Zhu. A causal and-or graph model for visibility fluent reasoning in tracking interacting objects. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [60] Wei Yang, Wanli Ouyang, Xiaolong Wang, Jimmy Ren, Hongsheng Li, and Xiaogang Wang. 3d human pose estimation in the wild by adversarial learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 6, 7
- [61] Hashim Yasin, Umar Iqbal, Bjorn Kruger, Andreas Weber, and Juergen Gall. A dual-source approach for 3d pose estimation from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [62] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3D human pose estimation in the wild:

a weakly-supervised approach. In *IEEE International Conference on Computer Vision*, 2017. 2, 7

- [63] Xiaowei Zhou, Menglong Zhu, Spyridon Leonardos, Konstantinos G Derpanis, and Kostas Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [64] Silvia Zuffi and Michael J. Black. The stitched puppet: A graphical model of 3d human shape and pose. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2