

SimPose: Effectively Learning DensePose and Surface Normals of People from Simulated Data

Tyler Zhu, Per Karlsson, and Christoph Bregler

Google Research
{tylerzhu, perk, bregler}@google.com

Abstract. With a proliferation of generic domain-adaptation approaches, we report a simple yet effective technique for learning difficult per-pixel 2.5D and 3D regression representations of articulated people. We obtained strong sim-to-real domain generalization for the 2.5D DensePose estimation task and the 3D human surface normal estimation task. On the multi-person DensePose MSCOCO benchmark, our approach outperforms the state-of-the-art methods which are trained on real images that are densely labelled. This is an important result since obtaining human manifold’s intrinsic uv coordinates on real images is time consuming and prone to labeling noise. Additionally, we present our model’s 3D surface normal predictions on the MSCOCO dataset that lacks any real 3D surface normal labels. The key to our approach is to mitigate the “Inter-domain Covariate Shift” with a carefully selected training batch from a mixture of domain samples, a deep batch-normalized residual network, and a modified multi-task learning objective. Our approach is complementary to existing domain-adaptation techniques and can be applied to other dense per-pixel pose estimation problems.

Keywords: Person pose estimation, simulated data, dense pose estimation, 3D surface normal, multi-task objective, sim-real mixture

1 Introduction

Robustly estimating multi-person human body pose and shape remains an important research problem. Humans, especially well trained artists, can reconstruct densely three dimensional multi-person sculptures with high accuracy given a RGB reference image. Only recently [2, 18, 19, 40, 42, 48, 53, 59] it was demonstrated that this challenging task can also be done using convolutional neural networks with modest in-the-wild accuracy. Improving the accuracy may lead to new applications such as telepresence in virtual reality, large-scale sports video analysis, and 3D MoCap on consumer devices. The challenges presented in natural images of people include mixed clothing and apparel, individual fat-to-muscle ratio, diverse skeletal structure, high-DoF body articulation, soft-tissue deformation, and ubiquitous human-object human-human occlusion.

We explore the possibility of training convolutional neural networks to learn the 2.5D and 3D human representations directly from renderings of animated

dressed 3D human figures (non-statistical), instead of designing an increasingly more sophisticated labeling software or building yet another statistical 3D human shape model with simplifications (e.g., ignoring hair or shoes). However it’s well-known that “deep neural networks easily fit random labels” [60]. As our simulated 3D human figures are far from being photorealistic, it’s easy for the model to overfit the simulated dataset without generalizing to natural images. Therefore the challenge lies in carefully designing a practical mechanism that allows training the difficult 2.5D and 3D human representations on the simulated people domain and generalize well to the in-the-wild domain where there are more details present, e.g., clothing variance (e.g. shoes), occlusion, diverse human skin tone, and hair.



Fig. 1: Visualization of SimPose’s 3D surface normal predictions on MSCOCO dataset.

Deep Convolutional Neural Networks (ConvNets) have been successfully applied to human pose estimation [8,23,44,45], person instance segmentation [9,20,37], and dense pose estimation [18] (the representation consists of human mesh’s continuous intrinsic uv coordinates). To reach the best task accuracy, they all require collecting real human annotated labels for the corresponding task [18,28]. As the research frontier for human pose and shape estimation expands into predicting new 2.5D and 3D human representations, it becomes self-evident that labelling large-scale precise groundtruth 2.5D and 3D annotations becomes significantly more time-consuming or perhaps even impossible.

Several generic domain-adaptation approaches have been proposed to leverage computer graphics generated precise annotations for in-the-wild tasks: 1) Using Generative Adversarial Network (GAN) framework [7] to jointly train a generator styled prediction network and a discriminator to classify whether the prediction is made on real or simulated domain. 2) Using Domain Randomization [43] to enhance the input image and thus expand the diversity and horizon of the input space of the model. Our approach is complementary to them. For instance, the GRU-domain-classifier [17] or the output space discriminator [46] can be added as additional auxiliary tasks in our multi-task objective, and more

uniform, sometimes corrupting, domain randomization can be added on top of our mild data augmentation.

In contrast, we create a new multi-person simulated dataset with accurate dense human pose annotations, and present a simple yet effective technique which mitigates the “Inter-domain Covariate Shift”. The closest prior works [48, 49] to ours in the human pose estimation field utilize a statistical 3D human shape model SMPL [30] for training (building SMPL requires more than 1K 3D scans and geometric simplifications e.g., ignoring hair, shoes). Our approach uses just 17 non-statistic 3D human figures and outperforms SOTA models on DensePose MSCOCO benchmark without additional few-shot real domain fine tuning.

In the following sections of the paper, we discuss related works and describe our SimPose approach in more detail: our simulated dataset, system implementation, experiment setup, the accuracy comparison with SOTA models on DensePose MSCOCO benchmark, and 3D human surface normal results. Our main contributions in the paper are the following:

- Create **simulated multi-person datasets** and present a **simple yet effective batch mixture training strategy with multi-task objective** to learn 2.5D dense pose uv and 3D surface normal estimation model without using any real dense labels.
- Evaluate our proposed approach on the DensePose MSCOCO benchmark [18]. We attained **favourable results using only simulated human uv labels, better than state-of-the-art models** [18, 34] which are trained using real DensePose labels. We also present our model’s 3D surface normal predictions on MSCOCO dataset which lacks any 3D surface normal labels.
- Show our approach trained using **only 17 non-statistical** 3D human figures can obtain better accuracy on the above benchmark than using SMPL [30] (a statistical 3D human shape model built with more than 1K 3D scans). This is a big reduction in number of 3D human scans required.

2 Related Work

Human pose estimation and its extension has been a popular research topic over the years. Early approaches study the structure of the human skeleton with a Pictorial Structure model (PS model) [16]. The method is later improved by combining it with a probabilistic graphics model [3, 57]. With the advent of deep learning, the methods are being simplified and the outputs are becoming richer and finer, as we discuss below.

ConvNet Human Pose Estimation Toshev and Szegedy [45] applies convolutional neural networks on the 2D human pose estimation task by regression to the (x, y) locations of each body joint in a cascaded manner. Tompson et al. [44] proposes a fully convolutional structure by replacing the training target from two scalar (x/y) coordinates to a 2D probabilistic map (or heatmap).

The robust heatmap based pose representation became popular and has been extended in many different ways. Wei et al. [51] trains a multi-stage fully convolutional network where the prediction from previous stages are fed into the next stage. Newell et al. [36] proposes an encoder-decoder structured fully convolutional network (called HourglassNet) with long-range skip connection to avoid the dilemma between the output heatmap resolution and the network receptive field. Papandreou et al. [38] augments the heatmap with an offset map for a more accurate pixel location. Chen et al. [11] proposes a cascaded pyramid structured network with hard keypoint mining. Recently, Bin et al. [54] shows that state-of-the-art 2D human pose estimation performance can be achieved with a very simple de-convolutional model.

3D Human Pose & Shape Estimation Tasks The progress on 2D human pose estimation motivates the community to move forward to more high-level human understanding tasks and take advantage of the 2D pose results. Bottom-up multi-person pose estimation [8, 35, 37] simultaneously estimates the pose of all persons in the image, by learning the instance-agnostic keypoints and a grouping feature. 3D human pose estimation [32, 62, 63] learns the additional depth dimension for the 2D skeleton, by learning a 3D pose dictionary [62] or adding a 3D bone length constraint [63]. Human parsing [14, 27] segments a human foreground mask into more fine-grained parts. DensePose [18, 34] estimation aims to align each human pixel to a canonical 3D mesh. Kanazawa et al. [25] predicts plausible SMPL coefficients using both 2D reprojection loss and 3D adversarial loss. Kolotouros et al. [26] improves it by incorporating SMPL fitting during training. Xu et al. [56] shows that using pre-trained DensePose-RCNN’s predictions can further improve 3D pose and shape estimation accuracy.

Simulated Synthetic Datasets When collecting training data for supervised learning is not feasible, e.g., for dense per-pixel annotation or 3D annotation, simulated training data can be a more useful solution. Chen et al. [10] synthesizes 3D human pose data by deforming a parametric human model (SCAPE [4]) followed by adding texture and background. Similarly, the SURREAL dataset [49] applies the CMU mocap motion data to the SMPL human model [31] on SUN [55] background, with 3D pose, human parsing, depth, and optical flow annotations. Zheng et al. [61] utilizes a Kinect sensor and DoubleFusion algorithm to capture posed people and synthesizes training images which contain one person per image. Compared to the previous simulated human datasets, our datasets contain multi-person images and two 3D human model sources (non-statistical and statistical) for comparison.

Domain Adaptation As the simulated synthetic data is from a different distribution compared with the real testing data, domain adaptation is often applied to improve the performance. A popular approach is domain adversarial training [47], which trains a domain discriminator and encourages the network to fool this domain discriminator. This results in domain invariant features. Domain randomization [43] generates an aggressively diverse training set which varies all factors instead of the training label, and causes the network to learn the inherited information and ignore the distracting randomness. Other ideas

incorporate task specific constraints. Constrained CNN [39] iteratively optimizes a label prior for weakly supervised semantic segmentation in a new domain, and Zhou et al. [63] enforce bone length constraint for 3D human pose estimation in the wild. In this paper, we propose a simpler alternative that generalizes well.

3 Simulated Multi-person Dense Pose Dataset




Fig. 2: Visualization of the 3D human model sources used to create our two simulated people datasets. Renderpeople 3D figures to the left and SMPL with SURREAL textures to the right. For SMPL, we only visualize its two base shapes here, but we sample shapes from its continuous shape space in our experiments. The human manifold’s intrinsic uv coordinates labels are visualized as isocontour lines where the background is the dense value of v .

Collecting dense continuous labels on real images is not only expensive and time-consuming but also hard to label accurately. By creating a simulated synthetic dataset, we are able to get consistent per pixel 2.5D and 3D human pose and shape labels (body part segmentation, body part intrinsic uv coordinates, and 3D surface normals), in comparison to [18] where the annotations are not per-pixel labelled. This section describes how we prepare our human 3D models with the non trivial task of generating correct uv mapping.

3.1 Human 3D Models

We create and compare two simulated datasets with two different sources of human 3D models, one consisting of only 17 rigged high resolution 3D scans of humans acquired from Renderpeople [41] and one using the statistical SMPL body model [31] (built from 1785 scans) together with SURREAL surface textures [49]. Other differences between the sources are the visual quality of the meshes and textures. The Renderpeople models include clothing and hair while the SMPL model does not. The Renderpeople textures are more realistic while the SURREAL textures may have artifacts. 3D reference points are attached to the skeleton of each rigged human model. We align these points with the MSCOCO 2D human keypoint definition for rendering simulated 2D keypoint.

Render People We ensure that all our human 3D models have the same uv and body part segmentation mapping as the references in DensePose labels [18] so that our approach can be evaluated on the MSCOCO DensePose benchmark. This is not easy to achieve as there are no available tools out there. We propose Algorithm 1, which takes about 40 min per Renderpeople model (total 17 3D human model) for a person with previous experience in 3D modeling software.

Algorithm 1 Transfer DensePose UV to Renderpeople meshes

```

for each bodypart submesh  $bp_{ref}$  in reference mesh do
  Mark a subset of the boundary vertices in  $bp_{ref}$  as landmarks  $L_{ref}$ .
  Store UV coordinates  $UV_{ref}$  for each vertex in  $L_{ref}$ .
end for
for each Renderpeople mesh  $rp$  do
  for each bodypart submesh  $bp_{ref}$  in reference mesh do
    Cut a submesh  $bp_{rp}$  from  $rp$  that is similar to  $bp_{ref}$ .
    Mark the same landmark vertices  $L_{rp}$  as in  $L_{ref}$ .
    Copy UV coordinates from  $UV_{ref}$  to  $UV_{rp}$ .
    Linearly interpolate UV coordinates along the boundary vertices in  $bp_{rp}$ .
    Mark boundary vertices in  $bp_{rp}$  as static.
    Unwrap UV on the inner non-static vertices, similar to [13].
  end for
end for

```

SMPL It is trivial to populate the DensePose UV coordinates from [18] since they also use the SMPL model as the reference mesh.

To make it easier to train and compare our approach on the datasets separately, we make sure that each image only has humans from one of the sources. Figure 3 shows example images from both sources. In total, we generated 100,000 images for each source.




Fig. 3: Examples from our simulated datasets. Top row shows the Renderpeople dataset and bottom row shows the SMPL dataset. Only the uv label is shown.




Fig. 4: Examples from our Renderpeople simulated datasets. From left to right: 2D keypoints, instance segmentation, body parts, uv , and 3D surface normal.

4 Our Approach

Given a two-stage multi-person pose estimation meta-architecture, our approach, as illustrated by Figure 5, anchors the domain adaptation of the 2nd-stage model’s new 2.5D and 3D human pose & shape estimation tasks on a set of easier 2D tasks. Inspired by the recent works of multi-task & meta-learning [5, 15, 58], a modified domain-label-based multi-task objective trains a deep normalized network end-to-end on a batch mixture of sim-real samples. We show that within a deep network, batch normalization, together with convolutions, can permit whitening and aligning bi-modal distributions via Batch Mixture Normalization. From an empirical perspective, our approach: **i.** generalizes from only 17 non-statistical 3D human figures, which we show in Table 1, and hence suffices as a functional replacement of using a statistical human shape model built from >1k 3D scans; **ii.** is able to directly domain-adapt continuous intrinsic uv coordinates and 3D surface normals.

4.1 Human Pose Estimation Meta-architectures

Among two widely used meta-architectures for multi-person human pose & shape estimation tasks, one is the bottom-up meta-architecture and the other is the two-stage meta-architecture. In the bottom-up meta-architecture, a neural network is applied fully convolutionally to predict instance-agnostic body parts for all people, as well as additional auxiliary representations for grouping atomic parts into individual person instances (e.g., part affinity field [8] and mid-range offsets [37]). The two-stage meta-architecture relies on a separate detection network or an RPN (region proposal network) whose outputs, ROIs (Region of Interest) of the input image or of the intermediate feature maps, are then cropped, resized, and fed into a 2nd stage pose estimation network.

We adopt the two-stage meta-architecture. The crop-and-resize operation is used as spatial attention of the second-stage network to restrict the divergence between real and simulation distributions within the ROIs. More specifically, our two-stage meta-architecture is the following: the 1st stage is a Faster RCNN person detector with ResNet101 [21] backbone (following [38], we make the same surgery to the original Faster RCNN including using dilated convolutions and training only on MSCOCO [29] person bounding box alone); The 2nd stage pose estimation network is shown in Figure 5. Before applying the corresponding

tasks’ losses, the outputs from the 2nd stage model are upsampled to be of the same size as the input image crops using the differentiable bilinear interpolation.

4.2 Human Pose Estimation Tasks

We adopt the term “task” that is used unanimously in the meta-learning, multi-task learning, and reinforcement learning literature. We cluster the human pose & shape estimation tasks into two categories: **i.** the 2D tasks and **ii.** the 2.5D/3D tasks. The 2D tasks are well-studied and stabilized, and they are: 2D body key-point localization and 2D person instance segmentation. The 2.5D/3D tasks are new, and we select the following two tasks: human mesh intrinsic uv coordinates regression and 3D surface normal vector regression. Below we describe each task and the task formulation in details.




Fig. 5: The ResNet101 network takes image crops batch mixture from both real and simulated domains as input and generates predictions for both 2.5D 3D tasks and 2D tasks which consist of sparse 2D body keypoints and instance segmentation prediction tasks. We compute the 2D tasks’ losses for both the real and the simulated domains and compute the 2.5D and 3D tasks’ losses only for the simulated domain.

2D Human Keypoints Estimation Task

The sparse 2D human keypoint estimation task requires precisely localizing the 2D subpixel positions of a set of keypoints on the human body (e.g., 17 keypoints defined by MSCOCO: “left shoulder”, “right elbow”, etc). We follow [38] and let the 2nd stage network predict heatmaps and offset fields. We use Hough-voting [38] to get more concentrated score maps. The final 2D coordinates of each keypoint are extracted via the spatial argmax operation on the score maps. The heatmap predictions are supervised using per-pixel per-keypoint sigmoid cross entropy loss and the offset field is supervised using per-pixel per-keypoint Huber loss.

2D Person Instance Segmentation Estimation Task

The task of instance segmentation requires classifying each pixel into foreground pixels that belong to the input crop’s center person, and background pixels that don’t. We follow [9] and let the 2nd stage network predict per-pixel probability of whether it belongs to the center person or not, and use standard per-pixel sigmoid cross entropy loss to supervise the prediction.

2.5D Body Part Intrinsic UV Coordinates Regression Task

For the 2.5D uv coordinates regression task, following the work of [18], we predict K body part segmentation masks and $2K$ body uv maps in the DensePose convention (e.g., “right face”, “front torso”, “upper right leg”, and etc). We use per-pixel per-part sigmoid cross entropy loss to supervise the body part segmentation and used per-pixel per-part smooth l_1 loss to supervise the uv predictions. We only backprop the l_1 loss through the groundtruth body part segmentation region. For the uv maps we subtract 0.5 from the groundtruth value to center the range of regression.

3D Person Surface Normal Regression Task

For the 3D person surface normal task, the 2nd stage network directly regresses per-pixel 3D coordinates of the unit surface normal vectors. During post-training inference, we l_2 renormalize the 2nd-stage networks’s per-pixel 3D surface normal predictions: $\hat{\mathbf{n}} = \frac{\mathbf{n}}{\|\mathbf{n}\|}$. We use per-pixel smooth l_1 loss to supervise the surface normal predictions. We only backprop the loss L through the groundtruth person instance segmentation region S using a discrete variance of:

$$\frac{\partial L(\theta)}{\partial \theta} = \frac{\partial}{\partial \theta} \sum_i^m \iint_{S_i} \|\mathbf{n} - \mathbf{f}_\theta(I_i)\| dx dy \quad (1)$$

where θ is all the trainable parameters of the network, m is the batch size, I_i is the i -th image in the batch, (x, y) are the coordinates in the image space, $\mathbf{n} = (n_x, n_y, n_z)$ is the groundtruth 3D surface normal vector at (x, y) , $\mathbf{f}_\theta(\cdot)$ is the network predicted surface normal vector at (x, y) , S_i is the support of the groundtruth person segmentation in image I_i , and $\|\cdot\|$ is the smooth l_1 norm.

4.3 Multi-task Learning and Batch Mixture Normalization

Learning to estimate people’s 2.5D intrinsic uv coordinates and 3D surface normals from renderings of 17 simulated people alone, and generalizing to natural image domain of people, requires a deep neural network to learn domain-invariant representations of the human body. We utilize a shared backbone trained end-to-end to achieve this. The network needs to learn domain-invariant representations of a 3D human body at its higher layers that are robust to the inevitably diverged representations at its lower layers. We refer to this underlying issue as the Inter-domain Covariate Shifts.

Batch Normalization [22] was designed to reduce the “Internal Covariate Shift” [6, 12] for supervised image classification tasks. It assumes similar activation statistics within the batch and whitens the input activation distribution by its batch mean μ and batch variance σ . When training a shared backbone using a batch mixture of images sampled from the two domains of a) simulated

people and b) natural images of people, the two input activation distributions almost certainly don't share the same mean μ and variance σ . Below we show one non-trivial realization of multi-modal whitening using: 1. Batch Normalization, 2. Fuzzy AND ($\wedge$) OR ($\vee$) implemented as convolution followed by activation, and 3. implicitly learnt domain classifier within the network. We call it Batch Mixture Normalization (BMN), Algorithm 2. During training, we believe the deep batch-normalized shared backbone can whiten bi-modal activations from the sim-real batch mixture using a fuzzy variance of BMN, when it's the optimal thing to do:

Algorithm 2 Batch Mixture Normalization

Require: $s^k(\cdot)$: a k-th layer featuremap (useful only on sim images)
Require: $r^k(\cdot)$: a k-th layer featuremap (useful only on real images)
Require: $z^k(\cdot)$: a k-th layer featuremap implicitly learnt as domain classifier
Require: Fuzzy AND ($\wedge$) OR ($\vee$) implemented as convolution followed by activation

Sample a batch of m image $\{I_i\}_i^m \sim \mathcal{D}_{sim} \cap \mathcal{D}_{real}$
 Evaluate k layers to get: $\{z\}_i^m, \{s\}_i^m, \{r\}_i^m = z^+(\{I_i\}_i^m), s^+(\{I_i\}_i^m), r^+(\{I_i\}_i^m)$
 Mask domain outliers (the k+1 layer): $[\dots, \{s\}, \{r\}, \dots] = [\dots, \boxed{\{z \wedge s\}, \{\bar{z} \wedge r\}} \dots]$
 Batch Normalize (the k+2 layer): $[\dots, \{s\}, \{r\}, \dots] = \text{BatchNorm}([\dots, \{s\}, \{r\}, \dots])$
 Bring into the same channel (the k+3 layer): $[\dots, \{y\}, \dots] = [\dots, \boxed{\{s \vee r\}}, \dots]$
 ▷ For succeeding layers ($\geq k + 4$), feature $\{y\}_i^m$ is the whitened and aligned version of $s^k(\{I_i\}_i^m)$ and $r^k(\{I_i\}_i^m)$.

We incorporate domain labels in the multi-task objective to cope with the constraint that 2.5D human intrinsic uv and 3D human body surface normal labels are only available on the simulated domain. Intuitively each task demands that the network learns more about the human body structure and shape in a complementary way. Due to mixing and normalizing, the network's forward and backward passes depend on the composition of the sim-real batch mixture.

$$L(\{I_i\}, \theta) = \sum_{l=1}^m L^{2D}(f_\theta(\{I_i\})) + \sum_{\substack{j=1 \\ I_j \in \mathcal{D}_{sim}}}^m [L^{2.5D}(f_\theta(\{I_i\})) + L^{3D}(f_\theta(\{I_i\}))] \quad (2)$$

where L^{2D} is the loss term for the 2D keypoint estimation task and the 2D instance segmentation task, $L^{2.5D}$ is the loss term for the human intrinsic uv coordinates regression task, and L^{3D} is the loss term of the 3D human body surface normal regression task. $f_\theta(\cdot)$ is the shared human pose and shape backbone parameterized by θ , m is the batch size, and $\{I_i\}_i^m$ is the batch mixture of simulated image crops and real image crops. In practice, we adopt ResNet101 as our backbone and apply a set of 1×1 convolutions at the last layer of ResNet101 to

predict: keypoint heatmaps, keypoint offset field, person instance segmentation, body part segmentations, body part intrinsic uv coordinates, and 3D human surface normals (see Section 4.2). We apply a per task weight to balance the loss for each task. To avoid exhaustive grid-search and task gradient interference [58], we adopt a greedy hyper-parameter search for individual task weights: 1) search for 2D task weights: $w_{heatmap}$, $w_{offsets}$, $w_{segment}$. 2) fix the 2D task weights and sequentially search for the weights for uv and normal tasks: w_{parts} , w_{uv} , and w_{normal} . We used the following weights for each task: $w_{heatmap} = 4.0$, $w_{offsets} = 1.0$, $w_{segment} = 2.0$, $w_{parts} = 0.5$, $w_{uv} = 0.25$, and $w_{normal} = 1.00$. During training we apply an additional gradient multiplier (10x) to the 1×1 convolution weights used to predict the heatmap, offset field, and person instance segmentation.

Compared with [52], our main differences are: 1) we have a new multi-person 3D surface normal prediction (Section 4.2) that Detectron2 doesn't have; 2) we obtained competitive DensePose accuracy, learnt from synthetic labels (Section 5.2); 3) our modified multi-task objective induces a richer label set and gradients for the simulated part (add. 3D surface normal) of the mini-batch, vs a reduced label set and gradients for the real part of the mini-batch; 4) we optimize only the 2nd-stage pose estimator instead of jointly optimizing RPN and multi-heads.

5 Evaluation

Table 1: UV Performance on DensePose COCO **minival** split (higher is better).

	AP	AP ₅₀	AP ₇₅	AP _M	AP _L	AR	AR ₅₀	A ₇₅	AR _M	AR _L
Prior Works (use labelled IUUV)										
DP-RCNN-cascade [18]	51.6	83.9	55.2	41.9	53.4	60.4	88.9	65.3	43.3	61.6
DP-RCNN-cascade + masks [18]	52.8	85.5	56.1	40.3	54.6	62.0	89.7	67.0	42.4	63.3
DP-RCNN-cascade + keypoints [18]	55.8	87.5	61.2	48.4	57.1	63.9	91.0	69.7	50.3	64.8
Slim DP: DP-RCNN (ResNeXt101) [34]	55.5	89.1	60.8	50.7	56.8	63.2	92.6	69.6	51.8	64.0
Slim DP: DP-RCNN + Hourglass [34]	57.3	88.4	63.9	57.6	58.2	65.8	92.6	73.0	59.6	66.2
Ours (use <i>only</i> simulated IUUV)										
SimPose (3D model source: Renderpeople)	57.3	88.4	67.3	60.1	59.3	66.4	95.1	77.8	62.4	66.7
SimPose (3D model source: SMPL)	56.2	87.9	65.3	61.0	58.0	65.2	95.1	75.2	63.2	65.3

5.1 Experimental Setup

Our SimPose system is implemented in the TensorFlow framework [1]. The 2nd-stage model is trained using the proposed approach. We use 4 P100 GPUs on one single machine and use synchronous stochastic gradient descent optimizer (learning rate is set to 0.005, momentum value is set to 0.9, batch size is set to 8 for each GPU). We train the model for 240K steps. The ResNet101 backbone has been pre-trained for 1M steps on the same sparse 2D 17 human keypoints and instance segmentation annotations mentioned above (no real MSCOCO DensePose labels are used) from an ImageNet classification checkpoint. The 1st-stage

Faster RCNN person detector is trained on the MSCOCO person bounding box labels using asynchronous distributed training with 9 K40 GPUs. Stochastic gradient descent with momentum is used as optimizer (learning rate is set to 0.0003, momentum value is set to 0.9, learning rate is decayed by 10x at the 800K step of the total 1M training steps). The ResNet101 backbone has been pretrained on the ImageNet classification task.

We use the Unity game engine to generate our simulated people datasets. We render up to 12 humans per image. Each image is 800x800 pixels with 60 degrees field of view and with motion blur disabled. All shaders use the default PBR [33] shading in Unity with low metallic and smoothness values. The shadows use a high resolution shadow map with four shadow cascades. The position of each human on the ground plane is random given the constraint that it can not intersect another human. We pose the human models with a random pose from one of 2,000 different Mixamo animations, similar to [50]. The 17 Renderpeople models do not have blend shapes or different textures. Instead we randomly augment the hue [24] of the Renderpeople clothing to add variation.




Fig. 6: SimPose’s UV predictions on DensePose MSCOCO minival split.

5.2 Evaluation of SimPose’s UV Prediction

Table 1 shows our SimPose approach’s accuracy on the DensePose MSCOCO minival split (1.5K images). Despite the fact that our model has not used any real DensePose labels, it achieves 57.3 average precision measured by the GPS (Geodesic Point Similarity) metrics on the challenging multi-person DensePose MSCOCO benchmark, which is better than the DensePose-RCNN [18] model’s average precision of 55.8, and better than the state-of-the-art Slim DensePose [34] on most of the breakdown AP & AR metrics. Both DensePose-RCNN and Slim DensePose have been trained using real DensePose labels.

Furthermore, in Table 1 we also compare the accuracy of our system trained from Renderpeople and SMPL separately. We found using only 17 Renderpeople 3D human models and our proposed approach, the system achieves better performance than using SMPL (a statistical 3D human shape model built from more than 1K 3D human scans). We conduct ablation studies of the simulated data mixing ratio (Table 2) and the UV task weight (Table 3) for the 2nd stage model. In Figure 6, we visualize our SimPose system’s UV predictions on the DensePose MSCOCO minival split.

Table 2: Ablation Study of Simulated Data Mixing Ratio. Higher mixing ratio of simulated data is detrimental to model’s 2D tasks accuracy on MSCOCO and has diminishing returns for the 2.5D and 3D tasks on the Renderpeople validation set.

Task	Metric	Percentage of Sim Data				
		100%	75%	50%	25%	0%
Normal	ADD (lower is better)	19.3°	19.5°	19.9°	21.5°	74.6°
UV	L2 (lower is better)	0.075	0.075	0.077	0.082	0.386
Segmentation	IOU (higher is better)	0.26	0.68	0.69	0.69	0.70
Keypoint	OKS (higher is better)	0.32	0.74	0.76	0.75	0.77

Table 3: Ablation Study of UV Task Weight. Higher value doesn’t affect 2D tasks on COCO and has diminishing returns for UV accuracy on Renderpeople validation set.

Task	Metric	UV Task Weight				
		0.00	0.25	0.50	1.00	2.00
UV	L2 (lower is better)	0.384	0.077	0.076	0.074	0.071
Segmentation	IOU (higher is better)	0.698	0.692	0.697	0.695	0.697
Keypoint	OKS (higher is better)	0.761	0.756	0.760	0.760	0.756

5.3 Evaluation of SimPose’s Surface Normal Prediction

In Figure 1, 7, and 8, we qualitatively visualize SimPose’s 3D surface normal predictions on a withheld MSCOCO validation set which the model hasn’t been trained on. Our model generalizes well to challenging scenarios: crowded multi-person scenes, occlusions, various poses and skin-tones. We evaluate the 3D surface normal predictions on the Renderpeople validation set using ADD (Average Degree Difference). We conduct ablation studies of the simulated data mixing ratio (Table 2) and the 3D normal task weight (Table 4).

Table 4: Ablation Study of 3D Normal Task Weight. Higher value doesn’t affect 2D tasks on COCO and has diminishing returns on the Renderpeople validation set.

Task	Metric	3D Normal Task Weight				
		0.00	0.25	0.50	1.00	2.00
Normal	ADD (lower is better)	76.1°	22.1°	21.1°	19.9°	18.6°
UV	L2 (lower is better)	0.077	0.077	0.076	0.077	0.077
Segmentation	IOU (higher is better)	0.699	0.697	0.700	0.696	0.698
Keypoint	OKS (higher is better)	0.763	0.759	0.760	0.756	0.760



Fig. 7: SimPose’s 3D surface normal predictions on MSCOCO dataset.



Fig. 8: Comparison of normal predictions trained from different 3D model sources.

6 Conclusion

We have shown that our SimPose approach achieves more accurate in-the-wild multi-person dense pose prediction without using any real DensePose labels for training. It also learns 3D human surface normal estimation from only simulated labels and can predict 3D human surface normal on the in-the-wild MSCOCO dataset that currently lacks any surface normal labels. Our SimPose approach achieves this using only 17 non-statistical 3D human figures. We hope the creation process of our simulated dataset and the proposed training scheme opens doors for training other accurate 2.5D and 3D human pose and shape estimation models without manually collecting real world annotations, and still generalizes for accurate multi-person prediction in-the-wild usage.

Acknowledgements We would like to thank Nori Kanazawa for reviewing our implementation, Fabian Pedregosa and Avneesh Sud for proofreading our paper.

References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., et al.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), software available from tensorflow.org
2. Alldieck, T., Magnor, M., Bhatnagar, B.L., Theobalt, C., Pons-Moll, G.: Learning to reconstruct people in clothing from a single RGB camera. In: CVPR (2019)
3. Andriluka, M., Roth, S., Schiele, B.: Pictorial structures revisited: People detection and articulated pose estimation (2009)
4. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. In: ACM transactions on graphics (TOG) (2005)
5. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: NeurIPS (2018)
6. Bjorck, N., Gomes, C.P., Selman, B., Weinberger, K.Q.: Understanding batch normalization. In: NeurIPS (2018)
7. Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: CVPR (2017)
8. Cao, Z., Wei, S.E., Simon, T., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: CVPR (2017)
9. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. TPAMI (2017)
10. Chen, W., Wang, H., Li, Y., Su, H., Wang, Z., Tu, C., Lischinski, D., Cohen-Or, D., Chen, B.: Synthesizing training images for boosting human 3d pose estimation. In: 3DV (2016)
11. Chen, Y., Wang, Z., Peng, Y., Zhang, Z., Yu, G., Sun, J.: Cascaded pyramid network for multi-person pose estimation. arXiv preprint arXiv:1711.07319 (2017)
12. De, S., Smith, S.L.: Batch normalization biases deep residual networks towards shallow paths. ArXiv (2020)
13. Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., Stuetzle, W.: Multiresolution analysis of arbitrary meshes. In: Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH (1995)
14. Fang, H.S., Lu, G., Fang, X., Xie, J., Tai, Y.W., Lu, C.: Weakly and semi supervised human body part parsing via pose-guided knowledge transfer. In: CVPR (2018)
15. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML (2017)
16. Fischler, M.A., Elschlager, R.: The representation and matching of pictorial structures. In: IEEE TOC (1973)
17. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.S.: Domain-adversarial training of neural networks. JMLR (2015)
18. Güler, R.A., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. In: CVPR (2018)
19. Habermann, M., Xu, W., Zollhöfer, M., Pons-Moll, G., Theobalt, C.: Livecap: Real-time human performance capture from monocular video. ACM Trans. Graph. (2019)
20. He, K., Gkioxari, G., Dollár, P., Girshick, R.B.: Mask r-cnn. In: ICCV (2017)

21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
22. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML (2015)
23. Jain, A., Tompson, J., Andriluka, M., Taylor, G., Bregler, C.: Learning human pose estimation features with convolutional networks. In: ICLR (2014)
24. Joblove, G.H., Greenberg, D.: Color spaces for computer graphics. In: Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH (1978)
25. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: CVPR (2018)
26. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In: ICCV (2019)
27. Liang, X., Shen, X., Feng, J., Lin, L., Yan, S.: Semantic object parsing with graph lstm. In: ECCV (2016)
28. Lin, T.Y., Cui, Y., Patterson, G., Ronchi, M.R., Bourdev, L., Girshick, R., Dollár, P.: Coco 2016 keypoint challenge. In: ECCV (2016)
29. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
30. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIGGRAPH Asia) (2015)
31. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. ACM Transactions on Graphics (TOG) (2015)
32. Martinez, J., Hossain, R., Romero, J., Little, J.J.: A simple yet effective baseline for 3d human pose estimation. ICCV (2017)
33. McAuley, S., Hill, S., Hoffman, N., Gotanda, Y., Smits, B., Burley, B., Martinez, A.: Practical physically-based shading in film and game production. In: ACM SIGGRAPH 2012 Courses. SIGGRAPH '12 (2012)
34. Neverova, N., Thewlis, J., Guler, R.A., Kokkinos, I., Vedaldi, A.: Slim densepose: Thrifty learning from sparse annotations and motion cues. In: CVPR (2019)
35. Newell, A., Huang, Z., Deng, J.: Associative embedding: End-to-end learning for joint detection and grouping. In: NeurIPS (2017)
36. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: ECCV (2016)
37. Papandreou, G., Zhu, T., Chen, L.C., Gidaris, S., Tompson, J., Murphy, K.: Personlab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. In: ECCV (2018)
38. Papandreou, G., Zhu, T., Kanazawa, N., Toshev, A., Tompson, J., Bregler, C., Murphy, K.: Towards accurate multi-person pose estimation in the wild. In: CVPR (2017)
39. Pathak, D., Krahenbuhl, P., Darrell, T.: Constrained convolutional neural networks for weakly supervised segmentation. In: ICCV (2015)
40. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: CVPR (2019)
41. RenderPeople: Renderpeople: 3d people for renderings. <https://renderpeople.com>
42. Saito, S., , Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: ICCV (2019)

43. Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P.: Domain randomization for transferring deep neural networks from simulation to the real world. In: IROS (2017)
44. Tompson, J., Jain, A., LeCun, Y., Bregler, C.: Joint training of a convolutional network and a graphical model for human pose estimation. In: NIPS (2014)
45. Toshev, A., Szegedy, C.: Deeppose: Human pose estimation via deep neural networks. In: CVPR (2014)
46. Tsai, Y.H., Sohn, K., Schultze, S., Chandraker, M.K.: Domain adaptation for structured output via discriminative patch representations. In: ICCV (2019)
47. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: CVPR (2017)
48. Varol, G., Ceylan, D., Russell, B., Yang, J., Yumer, E., Laptev, I., Schmid, C.: BodyNet: Volumetric inference of 3D human body shapes. In: ECCV (2018)
49. Varol, G., Romero, J., Martin, X., Mahmood, N., Black, M.J., Laptev, I., Schmid, C.: Learning from synthetic humans. In: CVPR (2017)
50. Villegas, R., Yang, J., Ceylan, D., Lee, H.: Neural kinematic networks for unsupervised motion retargeting. In: CVPR (2018)
51. Wei, S.E., Ramakrishna, V., Kanade, T., Sheikh, Y.: Convolutional pose machines. In: CVPR (2016)
52. Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. <https://github.com/facebookresearch/detectron2> (2019)
53. Xiang, D., Joo, H., Sheikh, Y.: Monocular total capture: Posing face, body, and hands in the wild. In: CVPR (2019)
54. Xiao, B., Wu, H., Wei, Y.: Simple baselines for human pose estimation and tracking. In: ECCV (2018)
55. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: Sun database: Large-scale scene recognition from abbey to zoo. In: CVPR (2010)
56. Xu, Y., Zhu, S.C., Tung, T.: Denserac: Joint 3d pose and shape estimation by dense render-and-compare. In: ICCV (2019)
57. Yang, Y., Ramanan, D.: Articulated pose estimation with flexible mixtures of parts. In: CVPR (2011)
58. Yu, T., Kumar, S., Gupta, A., Levine, S., Hausman, K., Finn, C.: Gradient surgery for multi-task learning (2020)
59. Zanfir, A., Marinou, E., Zanfir, M., Popa, A.I., Sminchisescu, C.: Deep network for the integrated 3d sensing of multiple people in natural images. In: NeurIPS (2018)
60. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning requires rethinking generalization. In: ICLR (2017)
61. Zheng, Z., Yu, T., Wei, Y., Dai, Q., Liu, Y.: Deephuman: 3d human reconstruction from a single image. In: ICCV (2019)
62. Zhou, X., Zhu, M., Leonardos, S., Derpanis, K.G., Daniilidis, K.: Sparseness meets deepness: 3d human pose estimation from monocular video. In: CVPR (2016)
63. Zhou, X., Huang, Q., Sun, X., Xue, X., Wei, Y.: Towards 3d human pose estimation in the wild: a weakly-supervised approach. In: ICCV (2017)